

ABSTRACT

My thesis is about whether we are justified to apply accurate information about a group onto an individual in the group. More specifically, it's motivated by our different intuitions towards the following two cases. I have a justified belief that, statistically, slopes covered with ice are more likely to be slippery than flat ground covered with ice. All else equal, am I justified in believing that a particular icy slope is more likely to be slippery than a particular icy flat ground?

The answer seems to be yes here. Now consider a very similar case. I have a justified belief that, statistically, black people are more likely to commit some crimes than white people are. All else equal, am I justified in believing that a particular black man is more likely to commit such a crime than a white man is?

The answer seems to be no. Even though the two cases seem structurally similar, we have very different intuitions towards the two cases. In my thesis, I attempt to give a reasonable account for our different intuitions by considering two types of solution: moral encroachment, solutions that are compatible with strong evidentialism, and finally, my own solution, which is also compatible with strong evidentialism. I argue that the projected beliefs in both cases are both justified and morally ok. What our intuition identifies is in fact a moral wrong, and it's not the moral wrong of the projected belief itself. Instead, the moral wrong resides in the downstream acts or speech that may be licenced by the projected belief. These acts and speech are morally problematic because they reinforce the current social injustice in the world by delivering false implications. Since there is no

discrimination against sloped ground, there is no social injustice to be worried about in Slippery Slope, and therefore the actions are not morally wrong in that case.

Statistical Justification

Is racial profiling based on accurate statistics justified?

A Thesis Submitted to the Philosophy Department Faculty of Mount Holyoke
College in Partial Fulfillment of the Requirements for the Degree of
Bachelors of Arts with Honors.

Yuan Tian
Mount Holyoke College
2019

Advised by
Katia Vavova
Mount Holyoke College

ACKNOWLEDGEMENT

I would like to first thank my thesis advisor, Prof. Katia Vavova. I am forever thankful that I took my first-year seminar with her by chance. Through that class, I was introduced to the wonderful world of Philosophy. From then on, Prof. Vavova has been an inseparable part of my academic journey at Mount Holyoke College. Without her constant encouragement and reassurance, I would not have had the courage to major in Philosophy, to continue with Philosophy in graduate school, or to start and eventually finish this year-long project. This thesis would not have been possible without her.

I would also like to thank the rest of my thesis committee, Prof. Nina Emery and Peter Scotto. I got to know both of them in my senior year, and I only wish that I've taken classes with them earlier. I've only taken one class with each of them, but I've learned much more from them than from many other people. Nina has been an important inspiration for my learning experience at Mount Holyoke and my application to graduate school. She sets the standard of a calm, sharp, and kind scholar for me. Prof. Scotto specializes in Russian Literature. However, his questions during my thesis defence and his email afterwards have been very philosophical and helpful for my further study on the topic. Thanks to both of them, I had a great time during my thesis defence.

My gratitude also extends to my friends at Mount Holyoke College. I would like to thank Wenny Shen for keeping me accompany, especially through

the last few days of the writing process. The time we spent together writing in her apartment is invaluable. Special thanks to Anisha Pai and Fiona Morris for agreeing to do a last minute proofread for my thesis. This piece would be full of typos without them. I would also like to thank other Philosophy majors. Their friendship makes my years at Mount Holyoke easily one of the happiest periods of my life by far. The mental support they offer makes this tough thesis writing journey much more bearable than it would otherwise have been.

Finally, I would like to thank my parents, Lei Xie and Qiang Tian. I would like to thank them for giving me the privilege of attending such a great college, majoring in the subject I love, and choosing Philosophy to be my potential life-long career. The possibilities they created for me with their hard work are priceless.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION

CHAPTER 2: BLACK MAN AND SLIPPERY SLOPE

CHAPTER 3: ONE SOLUTION--MORAL ENCROACHMENT

- 3.1 Motivating the Intuition with a Pragmatic Encroachment Case
- 3.2 Moral Encroachment Doesn't Collapse into Pragmatic Encroachment
- 3.3 One Version of Moral Encroachment--Risky Belief
 - 3.3.1 The View
 - 3.3.2 Can This View Explain Black Man vs. Slippery Slope?
- 3.4 Second Version of Moral encroachment--Moral Demand
- 3.5 First Objection--Knowledge Loss and Knowledge Gain
- 3.6 Second Objection--Inconsistent Beliefs
- 3.7 Conclusion

CHAPTER 4: OTHER SOLUTIONS--SOLUTIONS COMPATIBLE WITH EVIDENTIALISM

- 4.1 Basic Evidentialism and Strong Evidentialism
- 4.2 First Solution--Understanding of a Belief
 - 4.2.1 The View
 - 4.2.2 Rejection--Absurd Consequences
- 4.3 Second Solution--Relation between Properties
 - 4.3.1 The View
 - 4.3.2 Rejection--Not Applicable to Black Man and Slippery Slope
- 4.4 Conclusion

CHAPTER 5: WHAT'S REALLY WRONG IN BLACK MAN?

- 5.1 Implication and Literal Meaning of What We Say
- 5.2 Distinct Moral Harm from False Implication in Black Man
- 5.3 Our Intuition Fails to Capture the Epistemic Problem in Black Man

CHAPTER 6: CONCLUSION

BIBLIOGRAPHY

CHAPTER 1: INTRODUCTION

This thesis is about whether we are justified to apply accurate information about a group onto an individual in the group. As a Chinese student in the US, I'm faced with many judgements based on stereotypes. Since I'm Chinese, people assume that I probably know how to play a musical instrument; my math is probably better than my white classmates; I probably get a high GPA, etc. For most of the time, I live up to these stereotypes—I fit the descriptions of my ethnicity group quite well. There are moments when my US friends say to me, "Well that's racist! How can they assume that you are good at math just because you are Chinese?" I do feel a bit uncomfortable about being judged based on my Chinese identity, but my gut reaction has been, "Yeah, but these assumptions are quite accurate. How can these accurate assumptions be racist? And by the way, don't we judge people based on the group they belong to all the time?"

My experience above motivates me to write this thesis. Now, let me spell out the puzzle in more detail. From my experience, I know that I'm more likely to slip and fall on slopes covered with ice than on flat ground covered with ice. When I walk on the street and see a slope covered with ice, should I believe that I'm more likely to slip on this slope than the flat ground nearby? From the

statistics of the US Department of Justice, I know that in 2008, the homicide offending rate for black people (24.7 offenders per 100,000) was 7 times higher than the rate for white people (3.4 offenders per 100,000) (Cooper and Smith 2011, 11). When I walk on the street and see a black man, should I believe that he is more likely to hurt me than the white man walking nearby? Intuitively, the answer is yes in the former case, but no in the second case. We feel uncomfortable about judging the black man based on his race, but we don't feel discomfort judging the ground based on its slope. But why does our answer differ given the similarity of the two cases? One quick answer is that I'm more likely to slip on a slippery slope, than to be harmed by a black man. However, to make the two cases more parallel, we can stipulate that the likelihood of danger in both cases is the same. Even with this stipulation, it still seems that the answers differ in these two cases.

My thesis will focus primarily on whether my belief about this black man is justified by the statistics. Here are two clarifications. First, it might seem morally wrong to judge people based on their race, but it's not morally wrong to judge the ground based on its slope, since the former are moral agents, whereas the latter is not. However, I'm not primarily interested in the moral issue here. Instead, I want to focus on the epistemic issue in our asymmetrical intuition in these two cases. Epistemic issue concerns the epistemic status of my belief about the individual incident, i.e. this slope or this black man I encounter. Epistemic status is an overarching term that measures features of a belief, such as how

justified it is, how rational it is, or whether the belief constitutes knowledge, etc. I'm going to focus on justification for the belief. In particular, I want to see if my belief about *this* black man is justified by the statistics of the racial group he belongs to. In places where other authors depart from talking about justification of belief, I will explain their views in their original terms, and then put them in terms of justification.

However, this is not to exclude the discussion about the moral issue in my thesis—I will touch on the moral issue when it may give answers to the epistemic issue in the first four chapters. Eventually, the moral issue will become my main focus in chapter 5. One may wonder why it is important to try to identify an epistemic issue here if there is already a moral issue. Here is the reason why: if it's morally wrong and yet epistemically permissible to have certain beliefs, it seems that we need to sometimes give up justified beliefs to remain morally good. This is by no means a desirable consequence.

Second, I want to primarily focus on the beliefs instead of the downstream actions that may result from the beliefs. Of course, if our belief that the black man is more likely to harm us is justified, it seems to give us reason to act on the belief by crossing the street to walk on the white man's side, and this action seems morally problematic. Although beliefs are importantly connected to actions, I'd like to set aside the question about the downstream actions based on racial profiling for the first four chapters, and just focus on whether the belief is justified. After that, I will return to the downstream actions in chapter 5.

Straightforwardly, we could dismiss the asymmetry in two ways. On the one hand, we could say that our intuitions are not tracking any epistemic mistake in the racial profiling case. The two cases are actually identical epistemically. This means that either we shouldn't feel uncomfortable in the racial profiling case, or that our discomfort in the racial profiling case is legitimate, but it's only tracking the moral wrong, not the epistemic wrong.¹ The only thing relevant to determining the epistemic status of a belief is evidence. Since the evidence, i.e. the statistics, in the two cases is the same, our beliefs about the individual incident should be equally justified. If our discomfort is rightly tracking the moral wrong, then it means that we need to give up justified beliefs to be morally good. If there isn't even a moral wrong, then our intuition is faulty, and our discomfort here is an unfortunate mistake that our limited brains make. The view that evidence is the only thing relevant to determining the epistemic status of a belief is held by evidentialists (this definition can be quick but misleading, and I will give a more accurate one in chapter 4), and this way of dismissing the issue is straightforwardly compatible with evidentialism. But it does come with a cost, namely it's morally wrong to judge people based on their race.

On the other hand, we could say that the two cases are neither morally nor epistemically identical. It's morally wrong to judge people based on their race, but

¹ One might say that there is another way in which the two cases are identical—that we should feel uncomfortable passing judgement about the ground based on its slope, just like we feel uncomfortable passing judgement about the man. I'm not considering it because the view is too radical. According to this view, lots of our predictions about the world are unjustified. We can't have weather forecasts, consulting for a company based on its past statistics would be unjustified etc.

it's not morally wrong to judge the ground based on its slope, since the former is a moral agent, whereas the latter is not. In addition, the moral wrong has epistemic consequences. Because of the moral wrong, our beliefs about the individual incident based on statistics are unjustified. This kind of view is called moral encroachment, which I will explain in more detail later.

My thesis is laid out in the following order. In chapter 2, I will spell out the puzzle in more detail by including further descriptions of the two cases I mentioned at the beginning of the chapter. In chapter 3, I will explore the solution based on moral encroachment. This solution is incompatible with evidentialism, at least on the surface. I will reject this solution since it either stops us from making obvious inferences, or makes losing and gaining knowledge too easy. In chapter 4, I will explore some solutions that are compatible with evidentialism. I will reject the solutions since they either don't apply to the cases I'm interested in, or they lead to absurd consequences. In chapter 5, I will spell out my own solution. I argue that our belief itself is justified in the racial profiling case. Our intuition does track a distinct moral wrong in the racial profiling case, but it's not that the belief itself is morally problematic. Rather it's the downstream actions or speech that our belief may give rise to that are morally problematic.

CHAPTER 2: BLACK MAN AND SLIPPERY SLOPE

To clarify the concern, let's spell out the cases in more detail.

Slippery Slope. I'm walking alone on a rather shabby street in Brooklyn, New York. It's late at night, and street lights are dim. I see that the ground in front of me is covered with thin ice. When I get closer, I see that the ground on my side is sloped, and the ground on the other side of the street is flat. Suppose that statistics show that in similar situations, sloped ground is 30% more likely to make me slip and break my leg than flat ground is.² Based on the statistics, I believe that this sloped ground is 30% more likely to make me break my leg if I walk on it.

Is my belief about this particular slope justified based on the statistics? Before answering, let's also see the following case.

Black Man. As before, I'm walking down the street in Brooklyn, New York. I then see two people walking towards me—one on my side of the street and one on the other side. Both of them wear a hoodie, and have their hands in their pockets. When I get closer, I see that the man on my side is black, and the other one is white. Suppose that statistics show that in similar situations, black men are 30% more likely to rob me than white men are.³ Based on the statistics, I believe that this black man is 30% more likely to rob me if I walk towards him.

Suppose, for the sake of argument, that being robbed causes the same amount of damage to me as breaking a leg on a slippery slope does. The two cases are structurally similar. In Slippery Slope, I know the statistics about sloped ground, and I make an inference about this particular sloped ground I'm about to

² I made up the statistics. For the sake of argument, let's assume that the statistics are accurate.

³ Again, the statistics are made up.

walk on based on the statistics. In Black Man, I know the statistics about black men, and I make an inference about this particular black man I'm about to walk towards based on the statistics. In both cases, I make an inference about the individual incident based on the statistics about the group that the individual incident belongs to. Now, let me introduce some terminology so that I can spell out the inference steps I go through in more detail.

Statistics. In situations that are similar to that in Black Man, black men are 30% more likely to rob me than white men are.

In Black Man, suppose there are 100 black men and 100 white men in the Brooklyn community. Forty out of 100 black men and 10 out of 100 white men will rob me under similar situations as in Black Man. Therefore, there is a 40% chance that the black man I encounter is one of the 40 black men that will rob me. There is a 10% chance that the white man I encounter is one of the 10 white men that will rob me. Thus, the black man I encounter is 30% more likely than the white man to be one of the men that will rob me, and therefore the former is 30% more likely to rob me in this situation. More abstractly, statistics are frequency claims concerning a certain property based on some purported set of samples. In Black Man, the sample includes black men in the Brooklyn community. The property is 'will rob me in certain situations.' The frequency that the property presents is 40 black men out of 100.

Next, I form a belief about the group based on the statistics. Let's call it statistical belief.

Statistical belief. “Under situations that are similar to that in Black Man, black men are 30% more likely to rob me than white men are.”

Statistical beliefs are beliefs based on statistics. The belief could be the same as the statistics in terms of content, but it could also take other forms. For example, based on the statistics that 80% of my pens are blue, I can have a statistical belief that 80% of my pens are blue, which mimics my statistics completely. I can also believe that most of my pens are blue, which is slightly different from the statistics I have in terms of the content.

Projected belief. “This black man I encounter is 30% more likely to rob me than this white man I encounter.”

Projected beliefs are beliefs about an individual that are based on the statistics about the group that the individual belongs to. To get from a statistical belief to a projected belief, we need an inference step called projection.

Projection. The inference step between statistical belief and projected belief.

More specifically, I’m performing the following inference, where E stands for evidence.

$E1_{BM}$ Under certain situations, black men are 30% more likely to rob me than white men are.

$E2_{BM}$ This man is black, the other man is white, and this is one of the certain situations.

-----support-----
This black man is 30% more likely to rob me than this white man is.

In other words, when I project my statistical belief onto a new individual, I’m using two pieces of evidence to support my projected belief. $E1_{BM}$ is the statistical belief about a group, i.e. the black people in the neighborhood. $E2_{BM}$ is that the

individual in fact belongs to this group that I have statistical belief. The statistical belief is further based on the statistics. In sum, I go through three steps in Black Man. A diagram will be helpful in showing the inference process:




Figure 1. The steps I go through in forming a projected belief in general.

The same process applies to Slippery Slope. In Slippery Slope, I go from

Statistics: Under situations that are similar to that in Slippery Slope, sloped ground is 30% more likely to make me break my leg than flat ground is.

to

Statistical belief: “Under situations that are similar to that in Slippery Slope, sloped ground is 30% more likely to make me break my leg than flat ground is.

to

Projected belief: “This sloped ground I see is 30% more likely to make me break my leg than this flat ground I see.”

The inference step from statistical belief to projected belief through projection is the following:

E1_{SS} Under certain situations, sloped ground is 30% more likely to make me break my leg than flat ground is.

E2_{SS} This ground is sloped, the other ground is flat, and this is one of the certain situations.

-----support-----

This sloped ground is 30% more likely to make me break my leg than this flat ground is.

Given this inference process, it seems that there are two ways for the projected belief to go wrong in theory. It could be that the statistics are inaccurate,

that the statistical belief is not supported by the statistics, or it could be that the projected belief is not supported by the statistical belief. In my thesis, I'm going to assume that the statistics are accurate in both *Black Man* and *Slippery Slope*, and that the statistical beliefs are well supported by the statistics. Here is the definition of accuracy:

Accurate. A statistic is accurate if it correctly reflects the frequency of a certain property presented among the sample.

In *Black Man*, the sample may include the black and white men that live in Brooklyn, NY. The property at issue is the tendency to rob me in certain situations. If the statistics are accurate, then the frequency of the property 'will rob me in certain situations' among black male residence is 30% higher than that among white male residence. By assuming the statistics are accurate, I'm assuming that there is no epistemic error in the first step, i.e. it's justified to form statistical beliefs based on the statistics. In my thesis, I will assume that nothing goes wrong with the statistics and statistical belief. I will mainly focus on the step from statistical belief to projected belief.

I want to first eliminate some more straightforward mistakes we might make for the projected belief. In *Black Man*, it could be that there are more white men in the community than black men. So, even though the ratio of robbers among black men is higher than that among white men, there are more white male robbers than black male robbers in the community. Therefore, I can't infer that this black man is more likely to rob me than this white man is. To be clear, I've

eliminated this possible mistake by stipulating the white male and black male population in the Brooklyn community to be the same (cf. pp 9).

Here is another straightforward way in which my projected belief goes wrong in Black Man. I might overlook some important evidence that shows that the black man is economically well off and would not have the motivation to rob me. Maybe the black man is wearing a new Nike hoodie, but I overlook the label because of my bias against black people. The case I'm considering is not like this. I'm assuming that in Black Man, I overlook no evidence. In other words, the only relevant thing that can help me to distinguish the black man and the white man is their race and the statistics.

Now, let's go back to comparing Black Man and Slippery Slope. It seems that I go through the same projection step in both cases. In addition, $E1_{BM}$ (statistics) and $E2_{BM}$ (belief about the individual incidence) in Black Man seem to be very similar to those in Slippery Slope. Traditionally, it is held that only evidence can determine the epistemic status of a belief. This kind of view is called evidentialism.⁴ Given my interest in justified belief, the view instructs me to take only evidence to be relevant in determining whether my projected belief is justified or not. Since both cases have the same kind of evidence, it seems that either my projected beliefs are not justified in both cases, or they are justified in both cases. If our distinct discomfort in Black Man captures anything at all, it

⁴ There are multiple versions of evidentialism. I'm going to introduce two versions in chapter 4.

should capture the distinct moral wrong in Black Man.⁵ Therefore, in Black Man, if I stick to the projected belief, my belief is justified but I'm doing something morally bad. If I don't stick to the projected belief, I'm morally good but need to give up a justified belief—sometimes we need to sacrifice justified beliefs in order to remain morally good. Neither direction seems satisfactory.

There are three ways to resolve this dilemma. One way is to argue that the moral difference in the two cases actually has epistemic consequences. In other words, whether a belief is justified not only depends on evidence, but also on moral features of the belief. This view is called moral encroachment, and it's directly against evidentialism, at least on the surface.⁶ What is at stake in Black Man is social justice, and nothing similar is at stake in Slippery Slope. When the stakes are higher, we need more evidence for a belief to be justified. Since the moral stake in Black Man is quite high (it's social justice after all), $E1_{BM}$ and $E2_{BM}$ are not enough for the projected belief to be justified. But $E1_{SS}$ and $E2_{SS}$ are enough for justifying projected belief in Slippery Slope, because the moral stake here is quite low. Following this line of thought, it will be both epistemically and morally good for me not to form the projected belief in Black Man. There shall be no conflict between moral and epistemic good, at least in this case.

The second way will agree with the first way in that the two cases are epistemically different, but it will argue that the moral difference shouldn't have

⁵ For similar views, see Louise Antony, "Bias: Friend or Foe?" 183-188 and Tamar Gandler, "On the Epistemic Costs of Implicit Bias," 33-63.

⁶ Some moral encroachers argue that moral encroachment is compatible with a weak version of evidentialism. I will talk about this in chapter 3.

epistemic consequences. Instead, the projected belief in Black Man is not supported by the same kind of evidence as that in Slippery Slope, even though the evidence looks the same on the surface. Note that this solution is compatible with evidentialism.

The third way is to argue that the projected beliefs in both cases are both justified and morally ok. What our intuition identifies is indeed a moral wrong, but it's not the moral wrong of the projected belief itself. Instead, the moral wrong resides in the downstream acts or speech that may be licenced by the projected belief. In addition, the moral wrong here should not have epistemic consequences. These acts and speech are morally problematic because they reinforce the current social injustice in the world. Since there is no discrimination against sloped ground, there is no social injustice to be worried about in Slippery Slope, and therefore the downstream actions and speech are not morally wrong in that case. This route is also compatible with evidentialism.

In chapter 3, I will consider the first solution. In chapter 4, I will consider the second solution. I will eventually reject both solutions. In chapter 5, I will argue for the my solution, which is the third solution I mentioned. Now, let's look at the second solution—moral encroachment.

CHAPTER 3: ONE SOLUTION—MORAL ENCROACHMENT

One way to explain our distinct discomfort in Black Man is by endorsing a view called moral encroachment, which holds that the moral features of a belief may affect the epistemic status of a belief. There are multiple ways to spell out what exactly the moral feature and epistemic status are. In this chapter, I'm going to consider two versions of moral encroachment.⁷ The two versions differ in terms of what the relevant moral feature is, but they both take the relevant epistemic status to be whether the belief constitutes knowledge.

Note that the epistemic status considered by moral encroachment views is different from the one I'm interested in, which is whether the belief is justified. However, the lesson from moral encroachment is still transferable to justified belief, because knowledge and justified belief are importantly connected, and here is how. How justified the belief is can be measured in degrees. When I have more support for a belief, the belief is more justified than if I have less support for the belief. For a belief to be justified, it needs to have a certain amount of support. For a belief to be knowledge, it needs to be not only supported, but also supported

⁷ This is by no means a comprehensive survey about moral encroachment. I only present the three most plausible versions of it. For more radical views, see Rima Basu, "What We Epistemically Owe to Each Other," ms.

to a very high degree.⁸ According to the two versions of moral encroachment, the moral feature of a belief will affect whether the belief is supported enough to be knowledge. Since justification is measured in degrees, we can also argue that the moral feature of a belief also affects whether the belief is supported enough to be justified. When I present moral encroachment, I will present it in terms of whether the belief constitutes knowledge. But I will discuss how they affect justified beliefs after I present the views.

It's worth noticing that moral encroachment is very similar to a nearby view called pragmatic encroachment. Pragmatic encroachment holds that our practical interest in a belief may make a difference to whether it constitutes knowledge. In short, both views hold that something non-epistemic about the belief can affect the epistemic status of the belief. Some people argue that the two views should be understood analogously (Moss ms). More specifically, they argue that practical interest and moral feature of a belief should affect the epistemic status of the belief following the same mechanism. In section 3.1, I will use a pragmatic encroachment case to motivate our intuition. I will then show that the same intuition should be extended to moral encroachment cases. In section 3.2, I'm going to clarify that moral encroachment cannot be reduced to pragmatic encroachment, despite their similarity. In section 3.3, I will present a version of moral encroachment that takes the riskiness of a belief to be the relevant moral

⁸ This is not the only way to understand knowledge. People disagree on whether we should set a threshold for justification, after which the belief suddenly turns into knowledge. Bayesians, for example, don't think that we should understand knowledge in an "on-and-off" way. For yet another way of understanding knowledge, see James, *Pragmatism*.

feature. I will argue that this version fails to apply to Black Man and Slippery Slope. In section 3.4, I will present a version of moral encroachment that takes the moral demand of our environment to be the relevant moral feature. I will argue that this version does apply to Black Man and Slippery Slope. However, this version has two unfavorable results. It either causes us to lose or gain knowledge too easily, or instructs us to have inconsistent beliefs. I will address these two consequences in section 3.5 and 3.6, respectively.

3.1 Motivating the Intuition with a Pragmatic Encroachment Case

Consider the following conversation:

*Bank.*⁹

Lisa: “Do you know whether the bank in our neighborhood is open next Saturday? I need to make a deposit.”

Jones: “Yes. It will be open. The banks in that area are generally open on Saturday.”

Lisa: “Are you sure? I really need to make this deposit on time, since otherwise I will be fined \$10,000.”

Jones: “That much?! Well, now that I know, I’ll double check.”

At the beginning, it seems that Jones knows that the bank will open next Saturday. However, after he learns that Lisa would be fined \$10,000 if the bank doesn’t open, Jones ceases to know about the bank’s Saturday schedule. In this case, it’s not that Jones gets any new evidence against his previous belief about the bank schedule. Rather, the only new thing that Jones learns is the \$10,000 potential fine for Lisa. Before learning about the fine, nothing much is at stake for

⁹ The case is adapted from the bank example in DeRose, “Contextualism and Knowledge Attributions,” 913-929.

Jones, so his practical interest in knowing about the bank's schedule is low. After learning about the fine, more is at stake for Jones, and his practical interest in knowing is therefore higher. In Bank, Jones's practical interest in the case seem to affect whether Jones has enough justification for his belief about the bank's Saturday schedule to constitute knowledge.¹⁰ The take-away from Bank seems to be:

Pragmatic Encroachment. What you know depends not only on facts traditionally treated as epistemic, but also on facts about your practical interests. (Moss 2018, 193)¹¹

This view is called pragmatic encroachment. In Bank, for example, the belief at issue for Jones is whether the bank will be open next Saturday. The belief both before and after learning about Lisa's fine is supported by the same amount of evidence, which is "traditionally treated as epistemic." And yet, Jones knows that the bank will open next Saturday before he learns about Lisa's fine, whereas he doesn't know after learning about Lisa's fine.

This principle can be easily extended to the moral realm by connecting "practical interests" with "moral features." The idea is that if something non-epistemic can affect one's knowledge, then why couldn't this non-epistemic thing be a moral feature? The principle, in terms of moral features, would then be:

¹⁰ One might think that in addition to the practical factor, there is also a moral burden here, namely, Jones might feel more morally obligated to give Lisa the right information given the amount of potential fine. So, the difference between pre-fine Jones and post-fine Jones is not just the practical factor, but also there is the moral factor. But let's set the moral factor aside for now for the sake of argument, since it seems that practical factor itself is enough to cause post-fine Jones to lose knowledge.

¹¹ For a more formal definition, see Stanley, *Knowledge and Practical Interests*, 2.

Moral Encroachment. What you know depends not only on facts traditionally treated as epistemic, but also on facts about the moral feature of the belief.¹²

In other words, the moral features of a belief can affect whether or not the belief constitutes knowledge. This view is called moral encroachment. Note that moral encroachment can explain our different intuition towards Black Man and Slippery Slope. Roughly, the idea is that the moral stakes are much higher in Black Man than in Slippery Slope. For example, there is historical discrimination against black people in the US, but there isn't such discrimination against sloped ground. To apply statistics in Black Man is to reinforce the morally wrong discrimination against black people. The moral stakes here are high. In Slippery Slope, however, there is no such moral wrong given the absence of discrimination, and therefore there is nothing much at stake. Therefore, according to moral encroachment, my projected belief in Slippery Slope constitutes knowledge, since the statistics (E1) and the slope of the ground (E2) are enough justification for my belief to constitute knowledge. My projected belief in Black Man does not constitute knowledge. This is because the moral stake is higher. At higher stake, I need more justification than the statistics (E1) and the race of the black man (E2) for my belief to be knowledge.

It's worth noting that pragmatic encroachment and moral encroachment are very similar. In fact, many suggest that the two should be understood

¹² The original wording is "The moral features of a belief may make a difference to whether it constitutes knowledge" (Moss 2018, 223). Here it is adapted to mimic the definition of pragmatic encroachment. For a more formal definition, see Moss, "Moral Encroachment."

analogously. However, moral encroachment cannot be reduced to pragmatic encroachment. I will argue for this point in the next section.

3.2 Moral Encroachment Doesn't Collapse into Pragmatic Encroachment

Before going into different versions of moral encroachment, I want to clarify that moral encroachment is not merely a version of pragmatic encroachment. Recall that pragmatic encroachment holds that my practical interests affect whether my belief constitutes knowledge. In the previous section, I mention that the moral stakes in Black Man are high, since it is morally wrong to reinforce the racial discrimination in society, while the moral stake in Slippery Slope is low.

Here is why one might think moral encroachment can be reduced to pragmatic encroachment. Instead of distinguishing the Black Man and Slippery Slope in terms of moral stake, it seems that we can also distinguish the cases in terms of practical interest. It could be that since the moral stake is high in Black Man, I have a high practical interest in not doing something morally wrong. In Black Man, for example, it could be that I will feel very bad if I do something morally wrong, and I really don't want to feel that bad. In this case, I have a high practical interest of keeping myself happy in Black Man. Since I won't feel bad if I don't do something morally wrong, I don't have the same practical interest in Slippery Slope. Following this line of thought, it seems that moral encroachment can be reduced to practical encroachment (Basu ms c, 8).

Although it seems that moral encroachment can be reduced to pragmatic encroachment, it actually cannot, and here is the reason. If the difference between Black Man and Slippery Slope can be fully explained by my different practical interests, then the following will happen. Suppose I have no practical interest in being morally good—I have no interest in making myself happy, or making my parents proud, or making society better. If this is so, then it seems that I have equally low practical interest in both Black Man and Slippery Slope. This means that my projected beliefs in both cases are equally justified.

Even if we assume that I lack practical interest in both cases, we still seem to feel uncomfortable about projection in Black Man, but not in Slippery Slope. Why? This is because moral codes bind us even if we don't have any practical interest in being bound by them.¹³ Even if I don't care about social justice, it seems that I still morally should not reinforce racial discrimination. This means that the moral stake is high in Black Man and low in Slippery Slope *independent* of my practical interest in following the moral codes. So, my projected belief doesn't constitute knowledge in Black Man, but it does in Slippery Slope. Therefore, we cannot fully explain the difference between Black Man and Slippery Slope with only pragmatic encroachment. Since moral encroachment is capable of explaining cases that pragmatic encroachment fails to, the former doesn't reduce into the latter.

¹³ See Pace, "The Epistemic Value of Moral Consideration" and Moss, "Moral Encroachment."

3.3 One Version of Moral Encroachment—Risky Belief

As I mentioned at the beginning of the chapter, there are different ways to spell out what exactly constitutes moral features. In this section, I'm going to present a version of moral encroachment that takes the riskiness of a belief to be the relevant moral feature. This version is proposed by Sarah Moss (Moss ms, 16). In section 3.3.1, I will present the view, and I will show how this view may identify the epistemic difference between Black Man and Slippery Slope. In section 3.3.2, I will examine whether the view in fact explains the difference successfully. I will argue that this version fails to provide an explanation for the cases I'm considering.

3.3.1 The View

Recall that moral encroachment holds that the epistemic status of a belief can depend on its moral features. Based on this idea, Sarah Moss proposes the Rule of Consideration.

Rule of Consideration. In many situations where you are forming beliefs about a person, you morally should keep in mind the possibility that they might be an exception to statistical generalizations. (Moss 2018, 221)

I interpret the “statistical generalizations” to be projection. The “many situations” is determined by the moral features of the beliefs. Now, it's natural to ask which moral features of the beliefs can make a difference to their epistemic status. Moss thinks that the moral features here should be understood analogously to what spikes my practical interest in the pragmatic encroachment case. In Bank, it's the

riskiness of the belief that spikes Jones's practical interest. Here is what Moss means by taking a belief to be risky:

Risky belief. A belief is risky just in case acting on that belief would lead to a significant harm if and only if the belief turned out to be false. (Moss, 20)

In Bank, Jones's belief that the neighborhood bank will probably open next Saturday is risky because Lisa will lose \$10,000 if and only if his belief turns out to be false; if his belief about the neighborhood bank turns out to be true, then it doesn't harm Lisa. Since the moral feature should be understood analogously to what spikes Jones's practical interest, the moral feature should also be the riskiness of belief. In Black Man, you will harm the black person morally if and only if your belief about him turns out to be false. To see this, let's first suppose that the black man is in fact an innocent college student who enjoys wandering down the streets of New York at night. Then imagine you tell this person: "Oh, I think you are a thief!" This is to blame the man for something he didn't do, and such mistaken blame carries harm. Let's then suppose that the black man is in fact going to rob you and the white man is not. Imagine you tell this robber: "I think you are a thief!" Intuitively, it seems that you don't harm the person in the same way that you do in the previous supposition. From the above analysis, it seems that the relevant moral feature in moral encroachment can be the riskiness of a belief, which is also the relevant factor in pragmatic encroachment.

My projected belief is not morally risky in Slippery Slope. Suppose my belief turns out to be false in Slippery Slope—I did not slip and break my leg by

walking on the slope. In this case, I don't seem to harm the slope morally. The slope is not a moral agent, and it's incapable of comprehending such harm at all. In this way, false belief about a person might harm that person, but a false belief cannot harm the slope. Therefore, the projected belief is risky in Black Man, but not in Slippery Slope.

To further illustrate, let's also consider a case where one's belief is not risky:

Brown Eyes. According to statistics, about 70% of the students in my high school have brown eyes. I encounter a person wearing high school uniforms in the hallway, and I believe that this person probably has brown eyes, even though I'm not close enough to tell her eye color.

In this case, even if my belief that the person has brown eyes turns out to be false, the belief doesn't harm the person. Imagine that you say to a green-eyed person: "Oh, I thought you had brown eyes!" Intuitively, the person is not morally harmed just because I made a false judgment about her eye color.

But this might be different if people's eye colors are considered to be related to their intelligence. It doesn't matter if they are truly related—maybe people just tend to perceive brown-eyed people to be intellectually inferior to green-eyed people for socio-economic reasons. Suppose that I falsely take a green-eyed person to be brown-eyed based on statistics. Now it seems that by saying, "Oh, I thought you had brown eyes!" I'm suggesting, "Oh, I thought you were not that smart!" Intuitively, the person is harmed if I make a false judgment about her intelligence. This example shows that not all false beliefs about people

will harm them—only false beliefs regarding certain traits harm people. And the tendency to rob people belongs to these traits.

We've seen that the relevant moral feature for Rule of Consideration is the riskiness of belief. Let's apply the riskiness of belief to Rule of Consideration, and see how the rule blocks the projection in Black Man. Recall that the projected belief in black man is supported by both the statistics about a group and the fact that the black man belongs to such a group. Let's call the former $E1_{BM}$, and the latter $E2_{BM}$, where E stands for evidence.

$E1_{BM}$ In certain contexts, black men are 30% more likely to rob you than white men are.

$E2_{BM}$ This man is black, the other man is white, and this is one of the relevant contexts.

-----support-----
This man is 30% more likely to rob you than this white man is.

Since the projected belief is risky in this case, as I've argued before, the Rule of Consideration applies. The Rule of Consideration will instruct me to keep in mind the possibility that the black man I encounter is unlike the average black man in this hypothetical Brooklyn community. This is to say that $E2_{BM}$ is not strong enough for me to know that the black man is like the average black man in Brooklyn. If I keep such a possibility in mind, I will need more justification to know that this man is in fact just like the average black man in Brooklyn. This way, the moral feature of the context, i.e. the risk of acting on the belief, raises the bar of justification that's required for projected beliefs.

Note that Moss thinks that I still know $E1_{BM}$ and $E2_{BM}$ —it's just that they together cannot support my projected belief enough for it to constitute knowledge. In Slippery Slope, the projected belief is not risky, so the Rule of Consideration doesn't apply. Since I can think of the slope I encounter as an average slope on the street, my projected belief in Slippery Slope constitutes knowledge.

Note also that even though Moss focuses on the question of whether a belief constitutes knowledge, her view can directly apply to the question of whether a belief is justified, because knowledge and justified belief are importantly connected. There are many ways to spell out the connection between justification, belief, and knowledge, and here is the view I take. I take that how justified the belief is can be measured in degrees. When I have more support for a belief, the belief is more justified than if I have less support for the belief. For a belief to be justified, it needs to have a certain amount of support. For a belief to be knowledge, it needs even more support. Since justified belief and knowledge differ only in terms of the amount of the justification the belief has, the lessons we learn from Moss's version and will learn from the other version of moral encroachment about knowledge can be directly applied to cases about justified belief.

Here is how. In Black Man, the moral feature of my projected belief makes it such that I need more evidence for my belief to constitute justified belief than in Slippery Slope. Therefore, even though I have the same amount of

evidence in both Slippery Slope and Black Man, my belief in Black Man is unjustified, while my belief in Slippery Slope is justified.

In the following section, I will still talk about moral encroachment regarding knowledge instead of justified belief. This is because both Moss and Rima Basu—who holds a different version of moral encroachment that I will present in a later section—present the view in terms of knowledge. But just keep in mind that moral encroachment can also apply to justified belief, in the same way it applies to knowledge.

3.3.2 Can This View Explain Black Man vs. Slippery Slope?

Although this way of understanding moral encroachment seems promising, it actually fails to explain the discrepancy in Black Man and Slippery Slope. This is because my projected belief in Black Man, that “This black man is more likely to rob me than this white man is” is not risky. Let me remind you how the statistics in Black Man are generated, and after that perhaps we can see more clearly why Moss’s solution fails to apply.

Recall that the statistics I have in mind are as follows. In Black Man, suppose there are 100 black men and 100 white men in the Brooklyn community. 40 out of 100 black men and 10 out of 100 white men will rob me under similar situations as in Black Man. Therefore, there is a 40% chance that the black man I encounter is one of the 40 black men that will rob me. There is a 10% chance that the white man I encounter is one of the 10 white men that will rob me. Thus, the

black man I encounter is 30% more likely than the white man to be one of the men that will rob me, and therefore the former is 30% more likely to rob me in this situation.

In this case, the fact that this black man will not rob me is still compatible with the fact that there is a 40% chance that he belong to the 40 people who will rob me. In other words, the fact that this black man is innocent doesn't falsify my projected belief that this black man is 30% more likely to rob me than the white man is. Recall that Moss calls a belief risky when acting on that belief would lead to a significant harm if and only if the belief turned out to be false. In Black Man, however, it seems that I can harm the black man even if I act on my projected belief and the belief turns out to be true. After all, my projected belief only says that the black man is more *likely* to rob me, not that he will definitely rob me, and it cannot be falsified by the fact that he in fact will not rob me. Since my projected belief in Black Man is not risky, the rule of consideration doesn't apply. Therefore, my projected belief should constitute knowledge in Black Man according to Moss's view.

Based on the analysis above, it seems that when we take the relevant moral feature to be the riskiness of the belief, moral encroachment fails to apply to Black Man. My projected beliefs in both cases should be equally well justified, and therefore both of them should constitute knowledge. In this way, Moss's version of moral encroachment fails to identify the epistemic difference between the two cases.

3.4 Second Version of Moral Encroachment—Moral Demands

Recall that Moss thinks that the Rule of Consideration applies when the relevant moral feature is the riskiness of a belief. In the section above, I showed that my projected belief in Black Man is not risky, so the Rule of Consideration doesn't apply. Therefore, Moss's version of moral encroachment fails to identify the epistemic difference between Black Man and Slippery Slope. However, we can perhaps endorse a different version of moral encroachment, which may successfully identify the epistemic difference between the two cases. This version is proposed by Rima Basu (Basu ms d, 3). It takes the relevant moral feature to be the moral demands of the environment, instead of the riskiness of the belief. In this section, I will first present the view, and then argue that this version applies to Black Man.

Let's start from understanding the moral demands of the environment.

According to Basu,

[we need to] be aware of the moral demands of one's environment. With regard to our epistemic practices, it is the demand to be aware of the moral stakes of our beliefs about one another. It is the demand to be aware of the background against which our epistemic practices exist, i.e., the unjust world we inhabit, and to ensure that our epistemic practices are not only responsive to unjust features of our environment but that they also do not themselves contribute to those unjust features of our environment. (Basu ms c, 9)

By moral stakes, Basu means the moral demands of one's environment. Given the unjust features of our environment, such as institutional discrimination against

black people, it's morally demanded that we counteract such injustice by being more careful with our race-based projected beliefs. In order to be more careful with our beliefs, we need more evidence for a belief to constitute knowledge.

I argue that the rule of consideration is one of the ways to be more careful with our race-based projected belief.¹⁴ In other words, the injustice in the society morally demands that we keep in mind the possibility that the black man I encounter might not be average black man in the hypothetical Brooklyn community. Here, the relevant moral feature is the moral demands from the environment, and it requires that I need more justification for my projected belief to constitute knowledge in Black Man than in Slippery Slope. This is because in Slippery Slope, there is no social injustice with respect to slope, and therefore I'm not morally demanded to be more careful with my projected belief. In this way, my projected belief constitutes knowledge in Slippery Slope, but not in Black Man.

Unlike Moss's version, Basu's version applies to Black Man. Moss's view fails to apply because the fact that this black man I encounter is innocent doesn't falsify my projected belief, and therefore the belief is not risky in Black Man. According to Basu, even if our projected belief about the black man remains true, the moral demands of our environment still binds to us. Since moral demands of our environment requires us to be more careful with our belief, we need more

¹⁴ Note that the Rule of Consideration specifically applies to projected beliefs based on statistics. There might be other ways to be careful with other kinds of belief, but I'm only considering a specific kind of belief, i.e. the projected belief based on statistics.

justifications for our belief to constitute knowledge in Black Man, despite the belief being true. In Slippery Slope, however, there is no such moral demand from the environment. First, slopes are not moral agents. Second, there hasn't been historical discrimination against slopes. Since there is no such moral demand, we don't need to exercise the Rule of Consideration in Slippery Slope, and therefore my belief is justified enough to be knowledge in Slippery Slope.

In conclusion, Basu's version of moral encroachment can explain the epistemic difference between Black Man and Slippery Slope. However, I think that we should reject her view because it leads to two absurd consequences. It either makes us lose or gain knowledge in a strange way, or it results in inconsistencies in our beliefs. In the following two sections, I will raise these two rejections respectively.

3.5 First Objection—Knowledge Loss and Knowledge Gain

In this section, I'm going to show that Basu's version of moral encroachment allows a kind of knowledge loss.¹⁵ By explaining this consequence, I want to point out one of the major costs of endorsing moral encroachment.

To illustrate how the moral demands of one's environment, i.e. the relevant moral feature, affect whether our beliefs constitute knowledge, consider the following scenarios:

¹⁵ Other versions of moral encroachment may also allow the same kind of knowledge loss. However, I'd like to just focus on Basu's version here.

Sexist World. John is a hiring manager in an office. As a result of gender discrimination, most women in John's office are secretaries. John knows the proportion of women in the building, but is ignorant to the gender discrimination that causes it. John encounters Jane in the hallway in the morning. He doesn't know Jane, but he thinks that Jane is probably a secretary.

Fair World. Everything is initially identical to the sexist world, except a sudden feminist revolution occurs when John falls asleep during a lunch break. When he wakes up, unknown to John, all the sexists are brainwashed out of their sexist beliefs, and the employment rates for most companies suddenly reach gender equality.¹⁶ John encounters Jane again in the afternoon. He forgot he encountered her before, and he thinks that Jane is probably a secretary.¹⁷

Basu's version of moral encroachment says that the moral demands of our environment requires we have more evidence for certain beliefs for them to constitute knowledge. When combined with the Rule of Consideration, it means that we need to keep in mind that the individual may be an exception to the group she belongs to. In Sexist World, the sexist environment morally requires John to keep in mind that Jane may be an exception to her gender. Therefore, he needs more than the statistics and Jane's gender for his belief that Jane is probably a secretary to constitute knowledge. In Fair World, since the environment is not sexist anymore, John is not required to follow the Rule of Consideration, so the statistics and Jane's gender are enough justification for him to know that Jane is probably a secretary. Therefore, even though John has the same evidence in both worlds, he doesn't know that Jane is probably a secretary in the Sexist World. In

¹⁶ You may wonder why the feminists didn't wake John up and brainwash him. Note that John has been ignorant to sexism in the world. He notices that the proportion of female secretaries is extremely high in the building because knowing this is required for being a responsible hiring manager.

¹⁷ Thank you to Prof. Katia Vavova for suggesting this example.

Fair World, however, he gains this new piece of knowledge by just taking a nap during lunch break.

Fair World and Sexist World demonstrate a case of strange knowledge gain. In addition, Basu's version of moral encroachment seems also lead to strange knowledge loss. To see it better, let's compare this case to the Bank case, which is one of the paradigm cases used to illustrate the plausibility of knowledge loss. To repeat, the Bank case looks like this (cf. pp18):

Lisa: "Do you know whether the bank in our neighborhood is open next Saturday? I need to make a deposit."

Jones: "Yes. It will open. The banks in that area generally open on Saturday."

Lisa: "Are you sure? I really need to make this deposit on time, since otherwise I will be fined \$10,000."

Jones: "Well, now that I know, I'll double check"

Sexist World and Fair World could be put in conversation format:

Jane: Do you think I'm probably a secretary?

John: No. I don't want to be sexist.

Jane: You are in the post-revolution world now. Gender discrimination at workplace is not a thing.

John: "..."

Here, it seems that John should reply "Well then, I know you are probably a secretary," given his previous reason for not knowing is not wanting to be sexist. John seems to gain a piece of knowledge about Jane after his lunch break by just taking a nap. The case can also work in reverse—if somehow a fair world turns into a sexist world again, John will lose this piece of knowledge. To sum up, the change in the moral demand from the environment will lead to both strange

knowledge loss and strange knowledge gain according to Basu's moral encroachment thesis.

3.6 Second Objection—Inconsistent Beliefs

In addition to the strange knowledge gain and knowledge loss, moral encroachment may force me to have inconsistent beliefs. To illustrate this consequence, let's suppose that in Black Man, there are 100 black men and 100 white men in the Brooklyn community, and that's everyone in the community. 40 out of 100 black men and 10 out of 100 white men will rob me under similar situations as in Black Man.

Now, moral encroachment will instruct me that my projected beliefs about each individual black man doesn't constitute knowledge. According to Basu, this is because I should respond to the moral demands of my environment. Since there is racial discrimination in the social context, I should be more careful with my projected belief. Therefore I will need justification beyond statistics for my projected belief to constitute knowledge. As a result, when I encounter black man #1, I can't know that he is 40% likely to rob me. When I encounter black man #2, I can't know that he is 40% likely to rob me. Same thing happens when I encounter black man #3, and so on. For each individual, I can't know their likelihood of robbing me. However, when all black men stand together, I know

that 40% of them will rob me in certain context. How can I know about the group, when I know nothing about the individual?

Even if moral encroachers can give an account for the strange result above, the view may result in yet another inconsistency in my belief. Since there is no discrimination against white people in terms of violence in society, my environment doesn't morally demand me to be more careful with my projected belief about white men. Therefore, my projected belief about an individual white man constitutes knowledge. Therefore, for half of the population in the Brooklyn community (recall that we assume there are only 100 black men and 100 white men in the community), I can know that their likelihood for robbery is less than average. However, for the other half of this fictional community, I cannot know that their likelihood for robbery is higher than average. It seems outright irrational to know the former but not the latter.¹⁸

3.7 Conclusion

Moral encroachment says that the epistemic status of my belief depends partly on the moral feature of my belief. More specifically, the epistemic status here is whether my belief constitutes knowledge. The relevant moral feature is either the riskiness according to Moss or the moral demand of the environment according to Basu.

¹⁸ For similar rejection, see Gardiner, "Evidentialism and Moral Encroachment," 14. In addition, some people also argue one might lose knowledge by gaining evidence according to moral encroachment. For an argument following this line, see Eaton and Pickavance, "Evidence against pragmatic encroachment."

The epistemic status here is different from the one I'm interested in, which is whether the projected belief is justified. However, the lesson from moral encroachment is still transferable to my interest, because knowledge and justified belief are importantly connected. Given the high moral stake in Black Man, my belief needs more evidence to constitute justified belief than in Slippery Slope. Therefore, there is an epistemic difference between Black Man and Slippery Slope because there is a moral difference.

I've argued that Moss's version of moral encroachment fails to apply to Black Man and Slippery Slope. In addition, although Basu's version does apply, it leads to two absurd consequences. One is that we will gain or lose knowledge very easily based on this view. The other is that we will be forced to have some kind of inconsistency in our beliefs. Because of the absurd consequences, I argue that we should reject Basu's version of moral encroachment, even though it can successfully explain the epistemic difference between Black Man and Slippery Slope.

Note that moral encroachment seems incompatible with evidentialism, at least on the surface. In the next chapter, I will present two solutions that are compatible with evidentialism.

CHAPTER 4: OTHER SOLUTIONS—SOLUTIONS COMPATIBLE WITH EVIDENTIALISM

To seek an explanation for our distinct discomfort in *Black Man*, we considered the moral encroachment solution and rejected it. Recall that moral encroachment is incompatible with evidentialism, at least on the surface (cf. pp 6). Now, let's move on to solutions that are compatible with evidentialism. Broadly, evidentialism is the view that the only thing relevant to determining the epistemic status of a belief is evidence. In my thesis, I take the epistemic status of a belief to be whether the belief is justified or not. In section 4.1, I will present two versions of evidentialism—basic evidentialism and strong evidentialism—and show that although moral encroachment is compatible with the basic version, it's incompatible with the strong one. Therefore, if a solution is compatible with strong evidentialism, it cannot be used to defend or revise the moral encroachment thesis. In section 4.2 and 4.3, I will present two solutions that are compatible with strong evidentialism. In 4.2, I will present a solution suggested by Georgi Gardiner that concerns the understanding of a belief. In 4.3, I will present a solution suggested by Jessie Munton that concerns the relation between the tracking property and tracked property. In *Black Man*, for example, we use the property of being black to track the property of being likely to rob people. The

former property is the tracking property and the latter is the tracked property. I will eventually reject both solutions—I argue that Gardiner’s suggestion make our beliefs too falsifiable, and Munton’s suggestion fails to apply to Black Man and Slippery Slope.

4.1 Basic Evidentialism and Strong Evidentialism

In this section, I will first introduce two versions of evidentialism. I will then present the argument that moral encroachment is compatible with one, but not the other.¹⁹

Imagine a case where I want to know whether the bank opens on Saturday.

After doing the research, here is the information I find:

- (1) I care a lot about whether the bank opens on Saturday.
- (2) The bank website says it will.
- (3) My friend who works at the bank says it will.

With these three things, am I justified in believing that the bank opens on Saturday? To answer this question, I first need to answer two other questions. The first one is: can all three things serve as justification for my belief about the bank? The second one is: am I justified enough to believe that the bank will open on Saturday with these three things? I call the first one threshold question, and the second one justification question:

¹⁹ By moral encroachment, I mean Basu’s and Moss’s version. See Basu, “Moral Stakes” and Moss, “Moral Encroachment.”

Threshold question. What settles the required level of justification one's belief to be counted as justified?

Justification question. What can justify one's belief?

Let's first look at the justification question. You might think that (1) is not evidence, while (2) and (3) are. It seems irrational to believe that the bank will open just because I really want it to open. Although (1) is important, it cannot serve as justification because it's not evidence, while (2) and (3) can because they are evidence. In other words, you might think that only evidence can serve as justification for knowledge. If you hold this view, then you are at least a basic evidentialist. Here is the definition provided by Richard Feldman:

Basic evidentialism. Doxastic attitude D toward proposition p is epistemically justified for S at [time] t if and only if having D toward p fits the evidence S has at t (Feldman & Conee 15).

In other words, the epistemic justification of a belief depends only on evidence.²⁰

Note that basic evidentialism only give answer to what constitutes justification, but it doesn't give answer to how justified you will be with the evidence you have, or what settles the required justification level for your belief to be counted as justified. In other words, it only answers the justification question, but doesn't answer the threshold question.

Now, in addition to thinking that (1) is not evidence, and that (2) and (3) are, you might think that the level of interest you have in the case shouldn't affect whether you have enough justification for your belief about the bank based on (2)

²⁰ The principle is suppose to extend to all doxastic attitudes, including belief, withholding of belief, and knowledge etc. Here I'm only considering belief.

and (3). In other words, the amount of justification you need for a belief to be justified doesn't depend on a non-evidential factor. If you hold this view, then you are a strong evidentialist:

Strong evidentialism. The justificatory status of a belief depends only on evidential factors.²¹

The justificatory status here is the measurement of how justified one's belief is. In other words, evidential factors not only serve as justification, but also settle the required level of justification for one's belief to constitute being justified. In this way, strong evidentialism gives answer to both threshold question and justification question.

Let's see if moral encroachment is compatible with strong or basic evidentialism. It's actually compatible with basic evidentialism. When the moral stake is higher, the required level of justification for justified belief is raised by moral factors. The moral stake, i.e. the non-evidential factor, doesn't serve as justification. Rather, it settles the required level of justification for belief to constitute justified. Basic evidentialism is silent about what settles the required level of justification—it can allow such a level be settled by non-epistemic factors. Therefore, moral encroachment is compatible with basic evidentialism.²²

Moral encroachment is incompatible with strong evidentialism. Strong evidentialism says that the required level of justification for belief can be settled

²¹ For an even stronger version of evidentialism, see White, "Evidence Cannot Be Permissive."

²² The view is expressed in Gardiner "Evidentialism and Moral Encroachment", Basu "The Moral Stakes of Racist Beliefs", Basu and Schroeder "Doxastic Wronging."

only by evidence. Moral encroachment requires the level to vary based on what is at stake, and therefore the two theories are incompatible.

Note that the fact that the moral encroachment thesis is compatible with the basic evidentialism is by no means a sufficient reason for basic evidentialists to endorse it. Since basic evidentialism is silent to the threshold question, basic evidentialists might commit to different answers to the threshold question, and a moral encroacher's way of answering the threshold question is just one of them.

4.2 First Solution—Understanding of a Belief

In this section, I'm going to present a solution proposed by Georgi Gardiner.

Briefly, Gardiner thinks that a belief can be epistemically wrong if the understanding of the belief is wrong. Gardiner doesn't specify what she means by epistemically wrong, so I'm interpreting such wrong to include that the belief is unjustified. I will interpret Gardiner's view to be focusing on justification of belief throughout the chapter. In section 4.2.1, I'm going to present Gardiner's view and show how her view might explain the discrepancy between Black Man and Slippery Slope. In section 4.2.2, I'm going to argue that Gardiner's concept of understanding is so broad that it makes her view impossible.

4.2.1 The View.

Gardiner thinks that my projected belief is unjustified if I have an unjustified understanding of it (Gardiner Forthcoming, 20). Let's suppose that I project in

Black Man. From the statistics about black people, I believe that this black man is 30% more likely to rob me than this white man is. To show what Gardiner means by unjustified understanding of a belief, here is an example. My projected belief reflects the statistics about black men accurately. In addition, there is no further evidence available for me to tell whether the black man is more likely to rob me than the white man is. There is no substantial difference between two men's appearance, and I overlook nothing. However, my understanding of the statistics and the projected belief goes astray: I think that the black man is more likely to rob me because black men are more likely to do so by nature—the violence is in their genes. This understanding is unjustified because there is in fact plenty of evidence against this understanding.

In this case, my projected belief reflects evidence well, but the understanding of the projected belief reflects evidence poorly. In this case, the belief is epistemically wrong because it is embedded in faulty understandings of the world, and these understandings are faulty because they fail to reflect evidence correctly.²³ To put it in terms of justification, Gardiner thinks that beliefs should be justified by our understanding of it. If our understanding is unjustified, then the belief is unjustified. In other words, Gardiner thinks the following evidence 1 and 2 are not enough justification for the conclusion:

²³ See footnote 52 of Gardiner, "Evidentialism and Moral Encroachment."

E1_{BM} In certain situations, black men are 30% more likely to rob you than white men are.

E2_{BM} This man is black, the other man is white, and this is one of the relevant situations.

-----support-----
 This black man is 30% more likely to rob you than this white man is.

For the projected belief to constitute justified belief, there needs to be a third piece of evidence, namely,

E3 Our understanding of the projected belief is justified.

In other words, even if one's projected belief reflects E1 and E2 correctly, the belief can be unjustified since E1 and E2 are not enough to justify the belief. One needs to further justify her belief through her understanding of the belief, i.e. E3.

To further understand Gardiner's notion of understanding, consider the following two different understandings of the projected belief:

(1) Black people are more violent than people of other races by nature.

(2) Black people have harder time finding proper jobs to support their family due to social oppression.

Gardiner thinks that (1) could be a racist understanding in which the belief that "This black guy will probably rob me" is embedded. The belief is racist partially because it's unjustified (Gardiner Forthcoming, 22).

In addition, Gardiner thinks that I don't need to take (1) to be my reason for having the projected belief in order for (1) to be the understanding of the very

belief. (1) could just be something I also believe in addition to my projected belief. It might be helpful to think of our beliefs as forming a spider web in order to understand Gardiner's view. The beliefs are the nodes on the web, and each belief is connected to its adjacent beliefs, just like the nodes on a web are connected to adjacent nodes. All beliefs that are adjacent to the projected belief are understandings of the belief, and they are connected to the projected belief in that they serve as justification for the belief. In addition, the beliefs are connected independently of how I think they are connected. In other words, even if I don't take certain understandings to be the reasons for a belief, the understanding can still affect the epistemic status of the belief. Suppose I have the projected belief that this black man I encounter is more likely than the white man to rob me. My understandings of the belief may include the adjacent belief that black men are more violent by nature, the statistics about black men, and the appearance of this black man I encountered, etc. Even if I don't take any of them to be the reason for my projected belief, they still affect whether my projected belief is justified or not.

Let's apply the view to Black Man and Slippery Slope.²⁴ Gardiner thinks that my projected belief is unjustified if the belief is embedded in unjustified understandings of the world. Note that this view can only account for the

²⁴ Note that Gardiner doesn't apply her view to explaining difference between cases like Black Man and Slippery Slope. However, one of the main targets in her paper is moral encroachment. Therefore, I think it's reasonable to assume that Gardiner thinks that her view can do the job that moral encroachment can do. Since moral encroachment can account for the epistemic difference between Black Man and Slippery Slope, I assume that Gardiner will think that her view is also capable of doing so.

epistemic difference between Black Man and Slippery Slope under a special circumstance. It can give an account only when my projected belief in Black Man is embedded in unjustified understandings of the world whereas my projected belief in Slippery Slope is not. Let me spell out the circumstance in more detail. In Black Man, in addition to the projected belief, I may also believe that black men are by nature lazy, and that they want to get money without working hard. In this case, I have some unjustified understanding about the world. Since my projected belief about this black man is embedded in these unjustified understanding about the world, my projected belief is unjustified. In Slippery Slope, it might be the case that my projected belief about the slope is not embedded in any unjustified understanding about the world. All my beliefs about slopes and ground are justified. In addition, I have some justified understandings regarding the slope. Since my projected belief about the slope is embedded in justified understanding about the world, my projected belief in Slippery Slope is justified. Therefore, my projected belief is justified in Slippery Slope but not in Black Man, given the different understandings the projected beliefs embedded in. In this way, it seems that Gardiner's view can argue that there is an epistemic difference between Black Man and Slippery Slope.

4.2.2 Rejection--Absurd Consequences

Although Gardiner's view may explain our different feelings in Black Man and Slippery Slope to some extent, I think her view is much more radical than she

might think. In this section, I argue that her view may lead to massive loss in justified belief. Gardiner may be able to qualify her view by limiting the epistemic influence of understandings on beliefs. However, I argue that even this qualified version doesn't seem right.

Gardiner thinks that a belief is unjustified if it's embedded in unjustified understandings of the world. To see an absurd consequence of the view, consider the following model. Let's suppose that belief *A* is embedded in understanding *B* and *C*, and understanding *C* is further embedded in understanding *D* and *E*, as shown in the following diagram.




Figure 2. A web-of-belief model for beliefs A - G and more.

Suppose belief *D* is unjustified since it lacks evidential support. Belief *D* can serve as the understanding for belief *C*. Since belief *C* is embedded in an unjustified understanding, belief *C* is unjustified. Belief *C* can serve as the understanding for belief *A*. Since belief *A* is embedded in unjustified understanding *C*, belief *A* is unjustified. The series goes on. Eventually, it seems that most, if not all, of our beliefs are unjustified just because one belief is unjustified. Using the web-of-belief model, it seems that as long as one node is

tainted, then the node will taint its adjacent node. And eventually the whole web is tainted because of this one node.

It looks like our beliefs can easily be tainted by their surrounding beliefs. This might result in too many of our beliefs being unjustified because of one unjustified understanding. In response, Gardiner may qualify her definition of understanding. One way is to say that for an understanding to affect my belief, I need to actively take the understanding to be the reason for my belief. However, this qualification cannot savage her view.

To clarify how beliefs are related after such qualification, and to see the strange consequence of Gardiner's view even after the qualification, consider the following three scenarios.

Chemistry 1. You are calculating the standard deviation of ten measurements. You messed up several steps in the calculation and get the wrong answer.

Chemistry 2. You perform the same calculation as that in case 1, except this time you get the right answer out of luck even though you messed up several steps in the calculation.

Chemistry 3. You perform the same calculation as in case 1. This time you did every step correctly and get the right answer. The question then asks you to account for the standard deviation of the measurements. You give an account but the account is wrong.

Chemistry 4. You perform the same calculation as that in case 1. This time you get the right answer through correct calculation. Your account for the standard deviation is also correct.

Now, in which case is your standard deviation result justified? Intuitively, your result in case 1 and 2 are not justified since your calculations are wrong. In

other words, your result is not well supported because of errors in your intermediate inference steps. In case 3 and 4, your result in case 3 and 4 are equally justified, since you are right in every step of your calculation. However, according to Gardiner's view, your result is justified only in case 4. In case 3, your understanding of your result is wrong, and therefore there is an epistemic error in your standard deviation result. But this seems strange—it seems that in this case, even if your understanding of the calculation result is wrong, the calculation result itself should still be justified. The epistemic status of the calculation result should only be affected by the initial measurements and the calculation steps. It should not be affected by other beliefs that are related to the result.

My thought can also be explained using an analogy in ethics. Intuitively, we think that a bad person can still do something good. For example, suppose a serial killer chooses to save someone's life out of sympathy. Her saving someone's life is a good deed because she does so freely and out of sympathy. The action doesn't cease to be good because she has been doing wrong things in the past, nor does saving one person override her murder and make her a good person in general. In other words, it seems that we should be able to evaluate someone's action independent of the person's other actions. Similarly, we should be able to evaluate a belief independent of the web of belief that it fits in. A belief is epistemically good because itself is justified by evidence. It shouldn't cease to be epistemically good because some nearby beliefs are unjustified.

4.3 Second Solution—Relation between Properties

In 4.3, I will present a solution suggested by Jessie Munton that concerns the relation between the tracking property and tracked property. In *Black Man*, for example, we use the property of being black to track the property of being likely to rob people. I call the former property *tracking* property and the latter *tracked* property. Briefly, Munton argues that for the projected belief to be justified, we need to be able to use the tracking property to explain the tracked property, at least to some extent. In many cases, we project without having such explanatory relation between the tracking and tracked properties, so our projected beliefs are unjustified. In section 4.3.1, I'm going to present Munton's view. In section 4.3.2, I'm going to argue that although Munton identifies a common mistake that people make in projection, such a mistake doesn't apply to *Black Man* or *Slippery Slope*, and therefore her view fails to explain the discrepancy between the two cases.

4.3.1 The View

Let's start with a thought experiment.

Chinese Elm: According to accurate statistics, 98% of Chinese Elms are less than 18 inches in height. The remaining 2% can reach up to 15-20 meters. I come to believe this from a reliable source. One day, I find a Chinese Elm sprout in my yard, and I don't want it to block the sunlight.²⁵

²⁵ This example is adapted from Munton, "The Scope for Epistemic Flaws with Statistical Generalizations." In addition, the statistics about Chinese Elms are made up.

For the sake of argument, let's assume that the statistics are accurate. Recall that a statistic is accurate if it correctly reflects the frequency of certain property presented in the sample (cf. pp 12). In the case above, the sample consists of Chinese Elms in the world. Let's suppose that there are 1,000 of them. And let's suppose that 980 out of 1,000 Chinese Elms are less than 18 inches in height. The statistic is accurate because the property of less than 18 inches actually presents in 98% of Chinese Elms in the world.

So, should I believe that the sprout will likely grow tall or not? Just from the numbers, the answer seems to be clearly yes given the following inference:

E1_{CE} 98% of the Chinese Elms are below 18 inches.
 E2_{CE} The sprout in my yard is a Chinese Elm sprout.

-----support-----

The sprout in my yard is 98% likely to stay below 18 inches.

E stands for evidence. Here, it looks like the projected belief is well supported by the evidence. In addition, since we use the property of being elms to track the property of being short, let's call the former tracking property, and the latter tracked property. Even though the answer seems to be clearly yes, please hold on to your answer for now, and let me fill in more details.

Bonsai Elm: The statistics are the same—98% of Chinese Elms are less than 18 inches in height. Chinese Elm trees are native to eastern Asia. When left to grow naturally, they can reach 15 to 20 meters in height, but when they are cultivated as bonsai trees, they only grow to 18 inches tall. It happens that 98% of Chinese Elms are cultivated as bonsai trees, and are therefore under 18 inches in height. These cultivated elms are not genetically different. If left to grow naturally again, they would still reach 15 to 20 meters.

Intuitively, knowing only that 98% of Chinese Elms in the world are less than 18 inches tall doesn't allow me to project in Bonsai Elm. This is because the *real reason* why Chinese Elms in the world are short is that they are cultivated as Bonsai trees. We cannot predict the tree's height only based on E1 and E2.

Although the trait of being short correlates well with the trait of being an elm, the trees are not short because of the character of the species. In other words, although the tracking property correlates well with the tracked property, we can't use the tracking property to explain the tracked property. Given that no one will cultivate the sprout in my yard, the sprout will likely grow tall. Projecting simply based on E1 and E2 in Bonsai Elm will be a mistake. This is because the relation between the tracking and tracked properties is mere correlation. Projection, however, requires not only a correlation between two properties, but also a more robust one. Therefore, E1 and E2 in Bonsai Elm are short for projection (Munton ms, 13). This means that in addition to E1 and E2, we need E3 to justify our projected belief. In Bonsai, E3 will be:

E3: The relation between being elm and being short is more robust than mere correlation.

When put in terms of tracking and tracked properties, E3 should be:

E3: The relation between tracking and tracked properties is more robust than mere correlation

In the next section, I will look at a version of the Chinese Elm case and see what the right kind of relation will look like. A spoiler: there is more than one

kind. Then, I will go back to Black Man and Slippery Slope and see if both cases have the right kind of relation between tracking and tracked properties.

4.3.2 Rejection—Not Applicable to Black Man and Slippery Slope

Although I agree that Munton identifies a common mistake people make when projecting, this is not the error that we make in Black Man or Slippery Slope.

According to Munton, we may make the mistake of assuming that statistics show more than mere correlation between tracking and tracked properties in Bonsai Elm. Now, let me spell out more details about the statistics in Black Man, and it will be clear that the mistake Munton identifies doesn't apply to Black Man or Slippery Slope. Black people are historically discriminated against in the US.

They have been subject to a kind of social oppression that hinders them from being successful in life. They are assumed to be less promising in school; they are stopped by the police when driving more often than their white friends, etc.

Arguably, all these implicit or explicit discriminations that black people experience in their life make them more likely to fail to get a good education or a successful career. As a result, they are more likely to become robbers in the street than an average white man is. If we have to draw an analogy between the black man I encounter in Black Man and the sprout in Bonsai Elm, then the black man is not the sprout that I find in the yard that is not going to be cultivated as a bonsai tree. Instead, the black man is one of the trees that's subject to cultivation—he is already subject to the influence in his surrounding environment that stunts his growth as a human being.

Here is a case that's similar to the black man's situation.

Climate Elm: The statistics are the same—98% of Chinese Elms are less than 18 inches in height. When Chinese Elm trees grow in an eastern Asian climate, they can grow to 15-20 meters. Otherwise they stop growing at 18 inches. The climate in which the Chinese Elm trees grow doesn't change their genes. If a short elm is moved back to eastern Asia, it can still reach 15-20 meters. I do not live in eastern Asia. One day I saw a Chinese Elm sprout in my yard.

The correlation between being Chinese Elms and being short is not accidental.

The climate outside of eastern Asia hinders the growth of Chinese Elms, and therefore Chinese Elms outside of eastern Asia are largely short. Unlike the Bonsai Elm case, the relation between the two properties here is robust enough for projection. Intuitively, it seems legitimate for me to make the following inference in Climate Elm:

- E1_{CE} 98% of Chinese Elms outside of eastern Asia are short.
- E2_{CE} The Chinese Elm sprout in my yard is growing outside of eastern Asia.
- E3_{CE} The relation between being an elm and being short is more robust than mere correlation.

The Chinese Elm sprout in my yard is 98% likely to be short

The contrast between Bonsai Elm and Climate Elm suggests that if the relation between two properties is of the right kind, our projected beliefs can be justified.

Similarly, the correlation between the man's skin tone and his likelihood to commit a crime is not accidental. Black people are discriminated against because of their different skin tone, and this discrimination further results in their relatively low level of education, which further leads to their greater tendency of

becoming robbers. Since the causal relation between climate and Chinese Elm's height is robust enough for projection, it seems that the causal relation between the man's race and his tendency of being a robber should also be robust enough for projection. Therefore, although Munton's robustness thesis may explain why projection is not allowed in some cases, it cannot justify our discomfort in Black Man. If my projected belief is unjustified Black Man, it can't be because I project without E3.

4.4 Conclusion

In this chapter, I examined two solutions to the Black Man and Slippery Slope puzzle. These two solutions are compatible with strong evidentialism. Both solutions propose that in addition to E1 (the statistics of a group) and E2 (the fact that the individual is in the group), we need further evidence E3 so that my projected belief is justified. Gardiner argues that E3 is that my understanding of the belief is justified (Gardiner Forthcoming, 20). Munton thinks that E3 is the relation between tracking and tracked properties is more robust than mere correlation (Munton ms, 13).

I've rejected both proposals. I argued that Gardiner's proposal will make losing justified beliefs too easy. Even if she qualifies her view, her view is still subject to such critique. I agree with Munton that she identifies a common mistake that people make through projection. However, I argued that such a mistake doesn't apply to Black Man and Slippery Slope.

Therefore, I conclude that both solutions fail to show that there is an epistemic difference between Black Man and Slippery Slope. In the next chapter, I will present what I identify to be the difference between the two cases. Although I argued that Munton's view fails to apply to the cases I'm interested in, her idea that there needs to be a robust enough relation between tracking and tracked property for projection largely motivates my solution.

CHAPTER 5: WHAT'S REALLY WRONG IN BLACK MAN?

In chapter 3 and 4, I considered several solutions that attempt to identify epistemic mistakes we might make when we project in *Black Man* but not in *Slippery Slope*. I rejected all of them, because none of them explain our different intuitions towards the two cases successfully—the solutions either fail to apply to the cases I have in mind, or entail absurd consequences. In this chapter, I want to switch gears a bit. Instead of evaluating the solutions in the literature, I will argue for my own solution. I will argue that the discomfort we feel in *Black Man* doesn't come from the epistemic error of our projected belief. Instead, our discomfort captures the morally problematic downstream acts and speech that the projected belief may give rise to. In other words, although the projected belief itself is morally and epistemically fine, it may support some morally problematic speech or acts.²⁶ In *Black Man*, for example, my projected belief gives me a reason to tell my friend “Let's cross the street. That black man might be dangerous.” Despite being literally true, what I say is morally problematic because it has a false implication. I will explain how the literal meaning and

²⁶ Some people hold that belief itself can morally wrong people. For this kind of view, see Basu, “The Wrongs of Racist Beliefs,” Basu, “What We Epistemically Owe to Each Other,” and Basu and Schroeder, “Doxastic Wroning.” I'm not considering this kind of view here.

implication of what I say can differ later. For now, I want to give a clearer sense of how the concepts mentioned above are related by attaching the following diagram:




Figure 3. The relation between projected belief, morally problematic speech and false implication. The speech is morally problematic specifically because it has a false implication.

In section 5.1, I will distinguish between the implication and the literal meaning of what we say. I will then consider the specific implication that speech may carry in cases like *Black Man*, and how it can be literally accurate but have a false implication. In section 5.2, I will consider how the false implication makes what I say morally problematic at both an individual and societal level. In section 5.3, I will clarify that it's not the false implication itself that gives rise to our discomfort. In other words, our discomfort only captures the moral wrong of what we might say, but fails to capture the epistemic wrong of what we might say.

5.1 Implication and Literal Meaning of What We Say

Briefly, my solution relies on the distinction between the literal meaning of what we say and its implication. The implication of what we say can be different from its literal meaning. When they are different, what I say may be morally problematic even if it's literally accurate.

To see how the implication of what we say can be different from its literal meaning, consider the following scenario. Suppose a man has been meeting with his wife every day after work. One might say, “That married guy has been meeting the same woman every day after work.” The statement is literally true, since the guy is in fact meeting the same woman every day. The literal meaning of what we say is true as long as it corresponds with facts. What the person says also has a false implication, namely that the man is developing an unallowed romantic relationship. The implication of what we say is the meaning that we actually communicate, and what ends up being communicated might be different from the literal meaning. In this case, even though what the person says is literally true, it has false implication. In addition, we can also block the suggested implication by saying, “He’s not meeting a woman, he’s meeting his wife!” By saying this sentence, we are not rejecting the fact that the man has been meeting a woman. We are denying that the man is meeting with a woman that’s not his wife. It’s clear to us that the sentence has implications beyond its literal meaning, and the denial of the sentence targets its implication instead of its literal meaning (Haslanger 2011, 14).

Now, there might be the following question: if the implication of what we say is the meaning that we actually communicate and is not completely tracked by the literal meaning, then who is to decide what’s getting communicated? Does the speaker get to decide, or does the listener?

The implication is collectively determined by the members involved in the conversation, as well as the social context. “In order for communication, we must take certain things for granted as background to our conversation, i.e. we must presuppose certain things as common ground” (Stalnaker 2002, 701). “To presuppose a proposition ... is to take its truth for granted, and to assume that others involved in the context do the same... . Presuppositions are propositions implicitly supposed before the relevant linguistic business is transacted” (Stalnaker 2002, 279-80). In the conversation about the married man above, we listeners presuppose that if one says that a man has been seeing the same woman every day for a week, the person must mean that the man is seeing someone other than his wife. Otherwise, what’s the significance of pointing out that he’s been meeting with the same woman everyday? Even if the speaker doesn’t mean to deliver the false implication about the man, she does deliver such implication. This is because the implication of what she says is collectively determined by the speaker and listeners involved in the conversation, as well as the social context.²⁷ In other words, the speaker has some control over the implications of what she says, but she cannot have full control of the implication, since it is also decided by social context and other members involved in the conversation.

In *Black Man*, my projected belief is justified by the relevant statistics about how black men may act in certain situations. As I argued before, this is

²⁷ One may wonder what is social context doing here. Here is a suggestion. We understand the suggested romantic relation to be *problematic* because the society we live in allows only monogamous marriage. Therefore, social context does play a role in shaping the implication in this case.

because the relation between being black and the tendency for robbery is beyond mere correlation. In fact, the relation seems robust enough for justifying projected belief. The projected belief may then give me a reason to say things. Based on the projected belief, I might tell my acquaintance beside me (suppose there is one) “Watch out for the black guy”.²⁸ I will argue that in *Black Man*, the moral wrong of this kind of downstream speech resides in the false implication it has. Note that the same wrong applies to downstream acts as well, but in this chapter I will focus on only the downstream speech.

Now, I will talk about what implication my speech carries and how it is false in *Black Man*. To see this, we first need to return to the Chinese Elm thought experiments. What we learned by comparing Bonsai Elm and Climate Elm is this: there needs to be a robust enough relation between tracking and tracked properties for the projected belief to be justified. In Climate Elm, for example, the elms are sensitive to certain features of the climate, and the particular climate of where I live stunts their growth. The relation here is stronger than mere correlation and is robust enough for projection. There could, however, be a relation that’s more robust than that in Climate Elm:

Nature Elm: The statistics are the same—98% of the Chinese Elms are less than 18 inches. This time, they are short by nature. This means that climate, fertilizer, etc. don’t affect their growth a lot. Occasionally some elms undergo a genetic mutation and can grow up to 15-20 meters.

²⁸ Note that this is not to say the justified belief is sufficient in issuing any speech or act. I might have reasons not to say the sentence that overrides the reason to say so, and thereby end up not saying so.

In this case, the elms are short only because they are Chinese Elms. This is different from Climate Elm. Recall that in Climate Elm, the climate is what stunts the elms' growth. Here, it's in the elms' genes that they cannot grow tall. In other words, there is something in and of the elms' nature that causes them to stay short. Although the relation between being a Chinese Elm and being short is quite robust in Climate Elm, the relation is even more robust in Nature Elm. I will call the relation in Climate Elm *environmental*, and that in Nature Elm *natural*, based on what makes the tree short in each case. To repeat, the environmental and natural relation in the elm tree cases are:

Natural: The trees are short just because they are Chinese Elms.

This means that something in the tree's nature is making the trees likely to be short.

Environmental: The trees are short because of the climate.

This means that something that's external to the trees is making the trees likely to be short.

For my projected belief about the elms' height to be justified based on the statistics of elms, there needs to be a robust enough relation between the property of being an elm and the property of being short. Let's go back to the original Chinese Elm case, where I only know the statistics and have no information about the exact relation between being an elm and being short. If I do project in Chinese Elm, then I need to assume that the relation between the two properties is beyond mere correlation. I might assume that there is either a natural or environmental

relation between the two properties. It might not be clear exactly which robust relation I'm assuming, but I need to assume that there is a robust enough one.

Now, let's go back to Black Man and Slippery Slope. In Black Man, the two properties at issue are the property of being black and the property of being likely to rob me in certain contexts. Below are examples of potential environmental and natural relations between the two properties at issue:

Environmental: Black people are oppressed by society. The majority of society thinks that black men are prone to be more violent. As a result, black men are treated differently by their friends, teachers, and strangers on the street, etc. Because of the discriminative environment, black men are more likely to rob me than white men are in the relevant situation.

Here, the relation between being black and being likely to rob me is environmental. This means that something that's external to the black man is making him more likely to rob me than the white man is.

Natural: The genes of being black are biologically connected to the genes of robbery, so it's in black men's genes that they are more likely to rob me in the relevant situation than white men are.

To clarify, I'm not saying that there are in fact genes for robbery, or if there are, they are biologically connected to genes for skin tone. The point of this hypothesis is to show what might be at work if there is something in and of black men's nature that makes them prone to robbing people. In fact, it's not uncommon for people to have held this point in the past, or even in current society. There are many other ways to spell out the natural relation, and I'm spelling it out in terms of genes. I'm only using it to illustrate the hypothesis that there is something in black men's nature that makes them prone to robbery.

Either a natural or environmental relation is robust enough as a justification for my projected belief. Since the actual relation in Black Man is the environmental one, my projected belief is justified. However, the fact that the belief is justified is compatible with the fact that it leads to downstream morally problematic speech. Roughly, the speech is problematic because it may send out false implications. I going to spell this out in more detail below.

Since my projected belief is justified, the belief provide a reason for downstream speech. Suppose I'm walking with an acquaintance when I encounter a black man in the street. In this case, my justified belief gives me reason to tell her "Watch out for that black guy. He's more likely than the white guy to rob you." What I say may have false implication, and here is the reason why. For example, from my acquaintance's standpoint, me saying so shows that I formed the projected belief that the black man is more likely to rob me. Recall that for my projected belief to be justified, I need three things:

- E1 Under certain situations, black men are 30% more likely to rob you than white men are.
- E2 This man is black, the other man is white, and this is one of the relevant situations.
- E3 The relation between being black and being likely to rob people is beyond mere correlation.

Since I form the projected belief, I must have all three things in my mind. For E3, there are at least two ways to fill out the relation between the two properties. The relation could be environmental, or it could be natural. The projection itself is neutral with respect to which relation I take it to be. But if my acquaintance, who

doesn't know me very well, has to guess which one relation I'm going for, which one will she pick?

It seems that she is much more likely to think that I'm taking the natural relation to be the right one than to think that I go for the environmental one. To motivate your intuition more, just imagine someone you don't know very well tells you in a case like Black Man that you should watch out for the black man instead of the white man. What would be your reaction? For me at least, I tend to think that the person is assuming the natural relation. I tend to think so because of the discrimination against black people in the US social context. Historically, there were experiments aiming to show that black people are by nature inferior. In the 19th century, some western researchers believe that comparisons of skulls of different races justified ranking of races from least to most evolved.²⁹ This kind of scientific racism was developed and gained its popularity from the 1600s to the end of World War I. Even though it was largely criticized by the scientific community since the second half of the 20th century, the residue of its strong influence in society still exists nowadays. It's not uncommon for people to think that black people are poor because they are lazy and that they are not hardworking; they might get into prestigious colleges, but that's mostly because of the politically correct college admission policy that lowers the standard for their entrance. The assumption behind these thoughts is that there is something in and of black people's nature that's the reason for their poor behavior. Because of

²⁹ See Nott and Gliddon, *Types of Mankind*.

the racist social context, what I may say in Black Man has a racist implication, namely, this black man is by nature more likely to rob me than the white man is.

To see further how social context restricts the implication of what we say, consider the following example about gender discrimination. Just imagine someone saying, “You are probably bad at math because you are a woman.” Here, the person forms projected belief about an individual based on the statistics about female math students. Such a statement carries implications. Most of the time, the implication is perhaps that women’s brains are built for emotion, not for numbers. In other words, the implication is that there is a natural relation between being a woman and being bad at math. Note that the following may be true: women in fact do worse at math than men. But the implication picks up a false relation between being a woman and being bad at math—women are worse at math than men are because of the inequality in the education system, not because of their difference in nature. In other words, the true relation between being a woman and being bad at math is environmental, not natural. Therefore, by saying the sentence “You are probably bad at math because you are a woman,” the person delivers a false implication most of the time. Since the projected belief gives a reason for the person to say the sentence, the belief gives the person a reason to deliver a false implication.

Note that the person can make this mistake even though she doesn’t intend to deliver the implication with what she says. This is because implications are meanings of what we say that end up being communicated, and what ends up

being communicated is determined by the social context of the conversation, which is beyond the speaker's control. Recall the case where one says that the man has been meeting the same woman everyday. Even if the speaker doesn't intend to deliver a negative picture of the man, the social context of the conversation determines the way in which people interpret the sentence. Since the social context is formed based on the agreement among most people in the community, it cannot be controlled by the speaker alone. In other words, the speaker doesn't have much control of the implication of the sentence.

To clarify, when I say that the implication is not under the speaker's control, I'm only referring to the implication of the projected belief itself. The speaker can of course clarify the implication with further explanations. Suppose that after saying "You are probably bad at math because you are a woman," the person follows up with an explanation: "I don't mean that women's brains are not built for math. I mean that the discrimination in the education system makes it harder for women to succeed in math, and I think you might need more help because of that." This explanation seems to at least weaken our discomfort for the projection by clarifying the implication of the projection. Since the only thing the clarification does is change the implication, there is more of a reason for us to think that our discomfort for projected belief comes from the false implication of its downstream speech.

To sum up, in Black Man, my projected belief gives me a reason to say things with false implications. Suppose in Black Man, I say to my acquaintance

“Let’s cross the street. That black man might be dangerous” based on my projected belief. As I’ve argued before, what I say has a false implication, namely, this black man is by nature more likely to rob me than the white man is. What I say has such implication because of the US social context.

Note that our acts can deliver implications in a similar way as our speech does. Therefore, although my argument has been about speech, it also applies to the acts that our projected beliefs may give rise to.

Now, one may have the following question: what if I also deliver false implications about the slope in Slippery Slope because of some kind of social context? Suppose in Slippery Slope, I tell my acquaintance that “Let’s cross the street. That slope might be too slippery.” based on my projected belief about the slope. In addition, let’s suppose that the social context makes it such that my sentence implies a natural relation between being sloped and being more slippery, while the relation is in fact environmental. In this case, even though I deliver a false implication in both Slippery Slope and Black Man, it still seems that the flaw in Black Man brings about a distinct discomfort. Therefore, our original discomfort in Black Man cannot be explained by just the false implication.

I agree. The false implication alone cannot explain the distinct discomfort we feel in Black Man. I will present a more detailed argument for this in section 5.3, in case you are not convinced. For now, I will claim that there must be something more than just false implication that contributes to our discomfort in Black Man. In the next section, I will argue that the false implication we may

deliver through speech in *Black Man* makes the speech morally problematic, while the false implication in *Slippery Slope* does not. It's the moral wrong resulted from the false implication that distinguishes the two cases.

5.2 Distinct Moral Harm from False Implication in *Black Man*

The act or speech that sends out a false implication is morally problematic in two ways. Firstly, it harms the black man who may overhear and understand the implication of my conversation.³⁰ Secondly, it harms the society by reinforcing the social injustice that already exists, i.e. the discrimination against black people. More specifically, there are two ways in which acts or speech can reinforce the existing social injustice in *Black Man*. One is through reinforcing the implication in the black man's mind (and maybe also the white man's mind, if he recognizes it). The other is through reinforcing the implication in my mind. In this section, I will first argue that the speech has an implication that harms the black man, and also reinforces the implication in his mind. I will then argue that when I repeat the projected belief to myself, the sentence reinforces the implication in my mind. It should be obvious that the reinforced false implication in our mind is bad for social justice.

Let's start by looking at how the implication harms the black man.

Imagine you tell the black man, "You are more likely to rob me because you are

³⁰ In this case, realistically it seems more likely that my action of crossing the street will make my implication salient to the black man. But I will just focus on the speech for now.

genetically predisposed toward robbery.” When the black man asks me, “What about all this discrimination I’ve gone through? White men don’t get that when they grow up.” I answer “Well, those don’t really matter. You have the robber gene, so you are just more likely to rob someone sooner or later.” By saying so, I ignore the distinct difficulties the black man goes through. It seems that the black man is harmed in a straightforward way.

In addition, what I say may reinforce the social injustice by emphasizing the false implication in the black man’s mind through stereotype threat. Here is what I mean by stereotype threat. Researches show that when people recognize the stereotypes about their group, they are more likely to fulfill the stereotype. For example, when black people recognize the negative racial stereotype about their intellectual ability, they perform worse in verbal tests than they would have if the stereotype weren’t made salient to them (Steele and Aronson 1995, 799-801). The false implication of what I say in *Black Man* makes the stereotype that black men are more violent than white men salient in the black man’s mind, and thereby subjects the black man to stereotype threat. Suppose that the black man is not in a good economic standing and needs money for emergent reasons, he may be more likely to live up to the negative stereotype of his race. The fact that he lives up to the false implication will be used as misleading evidence to justify the false implication. This is not conducive to social justice. But do note that such stereotype threat in *Black Man* may seem unrealistic in some cases.³¹ Suppose the

³¹ Thank you to Prof. Katia Vavova for bringing up this point.

black man I encounter is in fact Barack Obama, but I fail to recognize him because he covers up his face. Obama is not going to be more likely to live up to the violent stereotype even if I make the stereotype salient to him by delivering the false implication. It seems that people with a good moral standing and strong will are not affected in Black Man by stereotype threat. However, I still want to point out that there are many people who are vulnerable to this kind of stereotype threat.

Now, let's look at how the projected belief can reinforce the false implication in my mind. Recall that a sentence carries an implication with it. Suppose I form the projected belief that the black man is more likely to rob me than the white man is, and recognize the belief in my mind. This is as if I repeat the sentence "The black man is more likely to rob me than the white man does" to myself. By repeating the sentence to myself, I'm repeating the false implication of the sentence to myself. Even if I believe explicitly that the natural relation is the false implication, while the environmental relation is the right one, I'm still subject to the effect of implicit bias. Implicit bias is the unconscious attribution of particular qualities to a member of a certain social group (Greenwald & Banaji 1995, 11). One could have implicit bias against members of a certain social group even if she explicitly disbelieves such bias. This means that even if I explicitly believe that the relation between being black and the tendency to rob is an environmental one, I still risk being influenced implicitly by the false implication that the sentence may deliver. This means that it's very difficult to counteract the

implicit influence of the false implication that the social context of the sentence forces on me. By repeating the sentence, I risk being influenced implicitly by the false implication. Having such implicit bias in my mind is by no means good for social injustice.

Now that we've seen that the false implication of our actions and speech gives rise to moral problems in Black Man, let's see if the false implication gives rise to the same moral problems in Slippery Slope. I think the answer to that is no. First, slopes are not moral agents, and we cannot possibly harm them morally. Second, there hasn't been historical discrimination against slopes, so we are not reinforcing any social injustice by delivering false implications about slopes. Therefore, the moral problem of our downstream actions and speech based on projected belief is what distinguishes Black Man and Slippery Slope.

5.3 Our Intuition Fails to Capture the Epistemic Problem in Black Man

By epistemic problem, I mean the false implication of what we say. Recall that at the end of section 5.1, I promised to give a more detailed account for why I think the false implication alone cannot explain the distinct discomfort we feel in Black Man. In this section, I will give a more detailed answer than what I gave in section 5.1. In addition, I will argue that the implication needs to be false in a certain way to make our action or speech morally problematic. Both points will be argued with one thought experiment.

In *Black Man*, the implication we deliver through speech is that there is a natural relation between tracking and tracked properties. The implication is false because the actual relation between the two properties is an environmental one. We feel uncomfortable about projection in *Black Man* because such false implication gives rise to the moral problem of our downstream action based on the projected belief. Now, there is another way in which the implication can be false. It could be that the actual relation between two properties is natural, but the implication says the relation is environmental. For example, suppose that black men are in fact genetically disposed to robbery under certain situation, while white men are not. However, the majority of society believes that black men are more likely to rob people because of the influence from the environment, such as the discrimination they undergo. Because of the change in social context, what I say about the black man has different implication than that in the original *Black Man*. Suppose I tell my acquaintance that, “Watch out for the black man. He might rob us.” Based on the social context, what I say implies that the black man is more likely to rob me than the white man is because of environmental effect. This implication is false, since black men are in fact by nature more likely to rob me than the white men are. In this case, do we feel the discomfort that we feel in the original *Black Man*?

My intuition is that the answer is no. Even though I do send out a false implication through what I say, I don't feel bad when the implication mistakes a natural relation to be environmental. If our discomfort captures the epistemic

wrong of the implication, we should feel the discomfort whenever the implication is different from facts. In the reversed Black Man case, however, even though the implication is false, we don't feel the discomfort. This shows that our discomfort in Black Man doesn't come from the false implication itself. Instead, it comes from the moral problem of our speech and actions that the false implications give rise to.

To clarify, this is not to say that it's ok to send out false implications. I just want to point out that our intuition only captures the moral problem that the false implications give rise to, and that it ignores the epistemic wrong of implications.

Our indifference in the reversed Black Man case shows another thing. It shows that we harm people when we mistake the environmental relation to be the natural relation. Suppose black men are in fact disposed to robbery by nature. It doesn't seem that one can harm a black man by saying, "I understand that you've gone through more difficulties than white people have done, and that's probably why you are more likely to rob me." Therefore, it seems that the downstream action or speech based on projected belief is morally wrong only when the implication of the belief mistakes the environmental relation to be natural relation.

CHAPTER 6: CONCLUSION

My thesis is motivated by our different intuitions towards the following two cases. I have a justified belief that, statistically, slopes covered with ice are more likely to be slippery than flat ground covered with ice. All else equal, am I justified in believing that a particular icy slope is more likely to be slippery than a particular icy flat ground? The answer seems to be yes here. Now consider a very similar case. I have a justified belief that, statistically, black people are more likely to commit some crimes than white people are. All else equal, am I justified in believing that a particular black man is more likely to commit such a crime than a white man is? The answer seems to be no. Even though the two cases seem structurally similar, we have very different intuitions towards the two cases. In my thesis, I attempt to give a reasonable account for our different intuitions by considering two types of solution: moral encroachment, solutions that are compatible with strong evidentialism, and finally, my own solution, which is also compatible with strong evidentialism.

Moral encroachment holds that the moral difference in the two cases actually has epistemic consequences. In other words, whether a belief is justified

not only depends on evidence, but also on moral features of the belief. Moss takes the moral feature to be the riskiness of the belief. This version of moral encroachment fails to apply to Black Man. Basu takes the moral feature to be the moral demands of our environment. Although this version applies to Black Man and is able to identify the epistemic difference between Black Man and Slippery Slope, I argue that this view will lead to absurd consequences, which includes easy knowledge loss and inconsistency in our belief. Note that moral encroachment is compatible with basic evidentialism, but is incompatible with strong evidentialism.

I also considered two solutions that are compatible with strong evidentialism by Gardiner and Munton. These solutions agree with moral encroachment that there is an epistemic difference between Black Man and Slippery Slope. However, the two solutions think that the epistemic difference doesn't result from their moral difference. Instead, the projected belief in Black Man is not supported by the same kind of evidence as that in Slippery Slope, even though the evidence looks the same on the surface. Both solutions propose that in addition to E1 (the statistics of a group) and E2 (the fact that the individual is in the group), we need further evidence, E3, so that my projected belief is justified. Gardiner argues that E3 is "my understanding of the belief is justified" (cf. pp 53). Munton thinks that E3 is "the relation between tracking and tracked properties is more robust than mere correlation" (cf. pp 53). I reject both proposals. I argue that Gardiner's proposal will make losing justified beliefs too

easy. Even if she qualifies her view, her view is still subject to such critique. I agree with Munton that she identifies a common mistake that people make through projection. However, I argue that such a mistake doesn't apply to Black Man and Slippery Slope.

After rejecting the two kinds of solutions, I propose my own solution. I argue that the projected beliefs in both cases are both justified and morally ok. What our intuition identifies is in fact a moral wrong, and it's not the moral wrong of the projected belief itself. Instead, the moral wrong resides in the downstream acts or speech that may be licenced by the projected belief. These acts and speech are morally problematic because they reinforce the current social injustice in the world by delivering false implications. Since there is no discrimination against sloped ground, there is no social injustice to be worried about in Slippery Slope, and therefore the actions are not morally wrong in that case.

This solution has interesting applications in many other difficult cases. For example, suppose that white people have a kind of genetic deficiency—they are more likely to catch a certain disease than people of another race. One day a patient comes to a doctor, and the patient is white. Should the doctor believe that this patient is more likely to catch this very disease than her black patients, and suggest that this patient do an extra test to see if she has actually caught this disease?

For another example, suppose we see a person with a disability. In general, it seems that disabled people are more likely to need more help in life

than people without disability are. Are we thereby justified in believing that this person with disability will need more help in life? Will it be condescending when we offer her help?

For yet another example, consider a case proposed by Renée Bolinger. Suppose Blue Bus company and Green Bus company are the only two bus companies in the city. As their names suggest, Blue Bus company owns only blue buses, and Green Bus company owns only green buses. In addition, 98% of the buses in the city are blue, and the remaining 2% are green. Now, I know that a person is hit by a bus. Based on the statistics, am I justified in believing that the person is 98% likely to be hit by a blue bus (Bolinger Forthcoming, 4)?

All these are interesting cases. However, since I'm running out of both time and space here, I'm going to leave the discussion to the future. But before I end my thesis, I'd like to consider one final doubt. My solution may not seem very cool to you. It might even seem too straightforward—so straightforward that you may wonder why I need to spend nearly one year on it. To this I reply: sometimes the true answer is just the straightforward one, and it is never a waste of time if the time is spent on defending the truth.

BIBLIOGRAPHY

- Antony, Louise. 2016. "Bias: Friend or Foe? Reflection on Saulish Skepticism." *Implicit Bias and Philosophy*. Ed. Michael Brownstein, Ed. Jennifer Saul. Oxford: United Kingdom, 157-190.
- Basu, Rima. ms a. "The Wrongs of Racist Beliefs."
- . ms b. "What We Epistemically Owe to Each Other."
- . ms c. "The Moral Stakes of Racist Beliefs"
- . ms d. "Moral Stakes."
- Basu, Rima and Schroeder, Mark. 2019. "Doxastic Wronging." *Pragmatic Encroachment in Epistemology*. Routledge, 181-205.
- Bolinger, Renée Jorgensen. Forthcoming. "The rational impermissibility of accepting (some) racial generalizations." *Synthese*.
- Cooper, Alexia and Smith, Erica. 2011. "Homicide Trends In The United States, 1980-2008." *Bureau of Justice Statistics*.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research*. 52(4): 913-929.
- Eaton, D and Pickavance, T. 2015. "Evidence against pragmatic encroachment." *Philosophy Study* 172: 3135.
- Feldman, Richard and Conee, Earl. 1985. "Evidentialism." *Philosophical Studies* 48 (1): 15-34.
- Gandler, Tamar Szabó. 2011. "On the Epistemic Costs of Implicit Bias." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 156(1): 33–63.

- Gardiner, Georgi. Forthcoming. "Evidentialism and Moral Encroachment." *Believing in Accordance with the Evidence: New Essays on Evidentialism*.
- Greenwald, A. G. and Banaji, M. R. 1995. "Implicit social cognition: Attitudes, self-esteem, and stereotypes." *Psychological Review*. 102 (1): 4-27.
- Haslanger, Sally. 2011. "Ideology, Generics, and Common Ground." *Feminist Metaphysics*, 179-207.
- James, Williams. 1907. *Pragmatism*. Cambridge, MA: Harvard University Press, 1979.
- Moss, Sarah. ms. "Moral Encroachment."
- . 2018. *Probabilistic Knowledge*. Oxford University Press, New York.
- Munton, Jessie. ms. "The Scope for Epistemic Flaws with Statistical Generalizations."
- Nott, Josiah Clark and Gliddon, George Robins. 1854. *Types of Mankind*. Creative Media Partners, LLC. 2015
- Pace, Michael. 2011. "The Epistemic Value of Moral Consideration." *Noûs* 45(2): 239-268.
- Stalnaker, Robert. 2002 "Common Ground." *Linguistics and Philosophy* 25: 701-721.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press, Oxford.
- Steele, Claude M and Aronson, Joshua. 1995. "Stereotype Threat and the Intellectual Test Performance of African Americans." *Journal of Personality and Social Psychology* 69(5): 797-811.
- White, Roger. 2013. "Evidence Cannot Be Permissive." *Contemporary Debates in Epistemology*. Blackwell. 312.