Morality Does Not Encroach

Zachary Goodsell and John Hawthorne

June 10, 2024

Abstract

Moral encroachment is the thesis that morality has an effect—unrecognized by traditional epistemology—on which doxastic states are epistemically appropriate. The thesis is increasingly popular among those who, in opposition to Gendler (2011), desire harmony between epistemic and moral demands on belief. This paper has three main goals. First, drawing on attractive structural principles concerning belief and justification, it is shown that a thoroughgoing harmony between moral and epistemic demands is implausible. This weakens the motivation for positing moral encroachment, but a proponent might still hope that moral encroachment can have a significant effect, so that many of the central cases of apparent disharmony can be avoided. The next goal of the paper is to argue that any significant effect of moral encroachment would manifest itself in the credences which are most epistemically appropriate. Finally, it is shown that even a small effect of moral encroachment on credence is in tension with widely accepted basic constraints on epistemically appropriate credence.

1 A Conflict Between Epistemology and Morality

Forming beliefs about people based on their race or gender can very often result in harmful dispositions to behave. Yet, as Gendler (2011) has argued, in some cases there seems very little to criticize about these attitudes from an epistemic perspective (in many other cases there is much to criticize: racism and sexism are not supported by traditional epistemology).

One example Gendler focuses on, and which has dominated recent literature on the topic, is the case of John Hope Franklin, a celebrated African American historian, at the exclusive Cosmos Club in Washington DC. Franklin was hosting a party at the club to celebrate his being awarded the Presidential Medal of Freedom. A woman clubgoer assumed he was staff on the basis of his race, and demanded

he fetch her coat. Her behaviour seems racist, but as Gendler points out, her belief that Franklin was staff was plausibly well-supported by the evidence, especially if, at that time, the vast majority of African American men at the Cosmos Club on a usual day were staff rather than members. Thus, depending on the details of the case, the woman's belief may well have been epistemically impeccable despite leading to racist behaviour.

Examples of this phenomenon are easy to multiply. One's background knowledge can lead to the apparently epistemically impeccable adoption of doxastic states that are distasteful in a similar way.

Handbag

Sally walks down the street carrying her handbag. She sees a stranger approaching. The stranger is of a marginalized ethnic group, which Sally knows commits robbery at a significantly higher rate than unmarginalized ethnic groups. As a result, she grips her handbag more tightly and is keenly aware of the stranger as they walk by each other. In fact, the stranger is an upstanding member of society. If the stranger had been of an unmarginalized ethnic group but otherwise similar, Sally would not have put up her guard.[1]

Assault

Miss X has been sexually assaulted, and she reports this to her friend, Mr Y. However, Mr Y knows that Miss X has recently received a treatment which he has been credibly informed can result in false memories of this nature. As a result, Mr Y fails to believe that Miss X was really the victim of sexual assault. In fact he is mistaken: the treatment has no such effect.

Eye Contact

Mary does not make eye contact with her teacher. As a result, the teacher writes her off as disinterested and unfocused. In fact, Mary is both interested and focused, but avoids eye contact because she is autistic.

Workplace

Victor is a visitor to a company where he is eager to impress people in management. He also knows that of the many women who work in the company, very few are in management. Sitting in an office, he hears

---

[1]Cf. Moss 2018a, p. 178.

two women speaking outside. He relaxes, giving little credence to their being in management. In fact these two women are two of the few who are in management.[2]

Each case is a plausible illustration of Gendler's tragedy: that epistemic demands require the adoption or non-adoption of doxastic states that one would, from a moral perspective, rather avoid. The relevant doxastic attitudes—fearing that one will be robbed on the basis of someone's ethnicity, disbelieving a genuine victim of sexual assault on faulty medical advice, writing off an autistic person as a bad student, and giving little credence to someone's being in management on the basis of her gender—clearly have something quite distasteful about them. It is unpleasant to live in a world where these doxastic attitudes are epistemically justified. However, there seems to be nothing to criticize about them, from an epistemic perspective.

Recent literature has pushed back on Gendler's claim that the demands of epistemology and morality can pull apart in this way. Opponents of Gendler posit a harmony between moral and epistemic demands that they cash out in a variety of ways.

> No Conflicts: If an epistemic attitude is epistemically impeccable, it must be morally permissible. (Bolinger 2020, p. 10)

> I believe that we should also seek to avoid [. . .] cases in which it is morally wrong to believe what is known or epistemically impeccable. It is much more plausible that beliefs wrong, if the beliefs that wrong are never epistemically impeccable, than that we can wrong one another even with epistemically impeccable beliefs. (Schroeder 2018, p. 116)

> [. . .] there cannot be beliefs that are both rationally epistemically permissible and also constitute doxastic wrongs. That is, doxastic wrongs are all epistemically impermissible. (Basu and Schroeder 2019, p. 199)

> There is no tension between the requirements of epistemic rationality and those of morality. (Johnson King and Babic 2020, p. 103)

> **No Tension:** A belief that morally wrongs a person cannot be an epistemically rational belief. (Jackson and Fritz 2021, p. 1393)

These authors impugn the woman's assumption that John Hope Franklin was staff as racist, and therefore deem it epistemically unjustified. And in related cases that

---

[2]Cf. Moss 2018b, pp. 220-229; Gardiner 2018, p. 174.

they present (and so presumably in the Handbag, Assault, Eye Contact, and Workplace cases), they impugn the relevant doxastic state as having some moral failing, so find it epistemically unjustified for the same reason. Asserting that traditional epistemology seems to speak in favour of the belief, they posit a previously unrecognized moral factor in epistemology that restores harmony between the moral and the epistemic. This factor is called *moral encroachment*.[3]

This paper has three main goals. First, drawing on attractive structural principles concerning belief and justification, it is shown that a thoroughgoing harmony between moral and epistemic demands as suggested in the above quotations is extremely unpromising, insofar as we adhere to typical judgments in the literature about which beliefs wrong (e.g., so that the harmony principles apply to the case of John Hope Franklin and similar cases). This weakens the motivation for positing moral encroachment, but a proponent might still hope that moral encroachment can have a significant effect, so that many cases of apparent disharmony, such as that in the John Hope Franklin case, are avoided. The next goal of the paper is to argue, in agreement with Johnson King and Babic (2020) and Jackson and Fritz (2021) that any significant effect of moral encroachment would manifest itself in the credences which are most epistemically appropriate. However, it is then argued in contradiction with these authors that even a small effect of moral encroachment on epistemically appropriate credence is implausible given widely accepted principles governing credence, such as that credence should be probabilistically coherent. Morality does not encroach.

The authors do not find it important to locate the discomfort we feel in the above

---

[3]We also note two quite different paths to the conclusion that in various of these cases the relevant doxastic state is epistemically unjustified. Many of the cases discussed in the literature involve outright beliefs based on statistical evidence, but it is not uncommon to think that one cannot know solely on the basis of statistical evidence, nor is it uncommon to think that an outright belief is justified only if it is formed in a way that is apt to produce knowledge. Readers should be careful to test whether this quite separate set of considerations is what is driving their judgment about those cases. To probe this, we recommend readers (a) look at cases where the relevant attitude is not one of outright belief but is something weaker (such attitudes will be important to our later discussion); and (b) they should reflect on whether their sense of epistemic impropriety remains when the method of belief formation is not statistical. For example, one could consider a variant of the John Hope Franklin case where we imagine Franklin resembles a member of staff that is known to the woman, and the woman believes Franklin is staff on the basis of this uncanny resemblance.

More simply, one might deny that any belief in a falsehood can be epistemically justified (for a sympathetic discussion of this idea see Williamson forthcoming). Since most paradigm cases in the moral encroachment literature involve false beliefs, this position would be sufficient to deny the justifiedness of the beliefs without having to wheel in moral encroachment at all. However, even someone who took this view about justified belief may turn to moral encroachment in other cases, where either (a) the attitude is less demanding, or (b) the content of the belief is true but still somehow morally problematic.

cases in the epistemic justification of the doxastic state. Most straightforwardly, the attitudes might trigger behaviour that is morally wrong, because the agent holding the attitude might inappropriately evaluate the consequences of their actions. In the John Hope Franklin case, even if the woman's belief he was staff was justified it seems clear that a suitable moral sensitivity would disincline her from presumptuously demanding her coat from a man not even in a staff uniform. And in the Eye Contact case, one would hope that the teacher would not neglect Mary even if her belief that she was unfocused and disinterested was justified. Racism and ablism raise the stakes for behaviour in these cases, rendering them similar to other cases where a justified belief does not make it reasonable to act as if that belief were certain. Moreover, even if the cases are modified so that the believer in fact knows with certainty, the behaviour would not be clearly permissible, since a clubgoer should be polite to staff and a teacher should not neglect students even if they are disinterested.

There are also a range of ways in which these attitudes could have something wrong with them without being epistemically unjustified. For example, they might be formed by problematic dispositions: if Sally thinks of crime every time she sees a member of the marginalized ethnicity in the Handbag case, she might be manifesting a morally and epistemologically problematic disposition when she forms her suspicion in a particular case. But if so then what is criticizable about Sally is her general disposition, and not necessarily her suspicion in the particular case. Similarly, the woman in John Hope Franklin's case might be disposed to, upon believing someone to be staff, think also of them as lesser. Thus her believing Franklin to be staff may have been problematic, but because of her disposition to thereby think of him as lesser, rather than because of a lack of justification.

Proponents of moral encroachment will insist that some harm can be blamed on the attitude itself, rather than the dispositions that prompt the attitude or the behaviour that flows from it. But we will argue that no matter where the wrong is located, there is no promising path to moral encroachment.

2 Moral and Pragmatic Encroachment

Moral encroachment is often taken to be a species of the longer-discussed phenomenon of *pragmatic* encroachment. According to pragmatic encroachment, considerations of practical action have an otherwise unrecognized effect on which doxastic states are epistemically justified or constitute knowledge. So-construed, the existence of pragmatic encroachment would strongly suggest the existence of moral encroachment, especially insofar as moral considerations are practical, as they are for morally well-adjusted agents. Despite this similarity, recent literature on moral

encroachment raises significantly different issues from those raised by pragmatic encroachment, causing novel problems for moral encroachment which will be discussed in the remainder of this paper. This section will also show how the differences between moral and pragmatic encroachment make it less clear how previous arguments against pragmatic encroachment, including those due to Cohen, are to apply to moral encroachment.

The most widely countenanced form of pragmatic encroachment—*stakes-based* pragmatic encroachment[4]—asserts that the standards required for a belief to constitute knowledge are raised when the agent of the belief has the opportunity to risk a lot on the basis of that belief.[5] It is easier to know, on this view, that the bank will be open tomorrow when one is not in a position to risk one's livelihood on that being the case. Proponents of stakes-based pragmatic encroachment are motivated by principles such as the following:

Knowledge-Action Link If an agent knows one available action would result in a better outcome than any other available action, then they should perform that action.

Knowledge-Action Link supports pragmatic encroachment because, usually, we are not inclined to say that an agent should risk it all on the bank being open tomorrow, even in circumstances where we would normally say the agent knows the bank will be open (e.g., if they remember the hours correctly listed on the bank's website). Suppose someone has the option of risking it all on the bank being open and shouldn't. Then by Knowledge-Action Link, they do *not* know, so the stakes of actions available to an agent has a surprising effect on what they know.

Moral encroachment has been hypothesised to work similarly:[6] by holding certain beliefs one is risking a moral wrong, such as the wrong of falsely thinking less of someone. So, the moral stakes of holding that belief are raised compared to believing an innocuous proposition, so the standards required to hold that belief justifiedly, or for it to count as knowledge, are raised accordingly.

Many opponents of pragmatic encroachment, including Cohen (2019), deny Knowledge-Action Link. According to Cohen, one can know a proposition that is in fact a decisive reason for performing a certain action, without possessing that proposition as a decisive reason for performing that action; Cohen dubs this

---

[4]See, e.g., Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005). However, not all versions of pragmatic encroachment are stakes-based, see Anderson and Hawthorne (2019) for discussion. We note that the "practical adequacy" view they discuss is arguably even more friendly to the Knowledge-Action Link.

[5]Ross and Schroeder (2014) endorse a related view, whereby in high-stakes situations the level of credence required to have a belief at all is raised.

[6]See, e.g., Schroeder (2018) and Basu (2018, pp. 78-89).

phenomenon *reason possession failure*. Therefore, the opponent of pragmatic encroachment can posit knowledge and/or justified belief in cases where acting on the belief in question would be unwise, and thus need not recommend intuitively unwise actions.[7]

The authors are also tempted to resist pragmatic encroachment by denying Knowledge-Action Link. But this strategy has unclear applicability to moral encroachment for two reasons. First, the literature on moral encroachment has tended to focus on justification rather than knowledge.[8] Second, the risky action in alleged cases of moral encroachment is the believing itself, rather than some action has consequences are tied to the content of the belief. Instances of Knowledge-Action Link, where the action is believing, are intuitively harder to deny. To be slightly more precise, if the two actions an agent can perform are either continuing to believing a certain proposition or not, and it is good to believe it if it's true and bad if it's false, then Knowledge-Action Link says in this case the following: if they know the proposition to be true then they should continue to believe it. Denying such instances of Knowledge-Action Link feels particularly surprising. If, as Basu, Schroeder, and other proponents say, you can commit a wrong with your belief even before anything untoward has come of that belief, then choosing not to act on that belief is besides the point. By knowing the proposition, one has already started believing it, so one has already accepted a very risky gamble by that very action. So the gap to be crossed from belief to action is much smaller. Thus, it is less clear that a Cohen-style response to pragmatic encroachment is effective when deployed against moral encroachment. The present arguments against moral encroachment are intended to work independently of what is said about pragmatic encroachment in general.

---

[7]Cohen's basis for denying Knowledge-Action Link turns on reason possession failure (i.e., denying that knowing a proposition suffices for possessing it as a reason for action). There may be other grounds for disputing Knowledge-Action Link that do not turn on positing reason possession failure. For example, whether or not someone possesses the proposition that the bank will be open tomorrow, someone who risks it all on that proposition may manifest problematic dispositions—if one risks it all when one does possess this reason, then, given human architecture, one will almost inevitably be disposed to risk it all when one thinks one possesses the reason but doesn't. Moreover, possessing a proposition as a reason need not suffice for that proposition having evidential probability one (supposing knowledge does not entail evidential probability one). And having less than evidential probability one in a proposition which is in fact a decisive reason for an action may not suffice for reasonably taking that action. For discussion of these issues see Lasonen-Aarnio (forthcoming) and Hawthorne and Stanley (2008).

[8]One might consider a variant of Knowledge-Action Link which substitutes justified belief for knowledge. We will not pursue the credentials of that principle here.

3 Against Harmony

Interpreted literally, the quoted harmony principles are obviously false, so the literal interpretation should be set aside immediately. Imagine a case where an evil psychologist has rigged a nuclear bomb to explode, killing millions, if you come to believe some otherwise innocent proposition, such as that they own an even number of cars. Morality requires you not to believe that proposition, and believing it would morally wrong the victims of the bomb, but it by no means follows that believing it would be epistemically unjustified. You must do what it takes not to believe the proposition, even if that takes disbelieving usually honest friends, doubting your vision, becoming irrational, or knocking yourself unconscious. But if you do end up believing it and killing millions, it could very well be an epistemically impeccable belief.

This case is not what the quoted authors have in mind. Rather, they have in mind cases where an attitude wrongs a person *directly*, rather than in virtue of its untoward consequences. It is thought that an attitude can directly wrong a person by being racist, sexist, uncharitable, disloyal; and the rough-and-ready test for directly wronging is if the subject of the belief would rightfully be offended if they knew it was held (the ideology of direct and indirect wronging is not found in the literature; however in the authors' ordinary usage of 'wronging' it is clear that one often wrongs someone only because of the bad consequences of one's action for that person). Thus, it seems clear that insofar as John Hope Franklin was directly wronged by the clubgoer's belief he was staff, so are the subjects of the four further examples directly wronged respectively by the respective suspicion, lack of belief, or low credence, since they would rightly be offended if they knew of the attitude. Schroeder (2018) tentatively proposes a necessary and sufficient condition for a belief directly wronging a person, namely that a belief wrongs a person directly just in case it falsely diminishes them, where diminishing is glossed as making out their agential contribution to be less than it is. Falsely assuming a person to fit racial stereotypes is presented as a prototypical example.[9]

Understood in this way, the quoted principles are much weaker than they initially seem. We shall argue that they are still implausible, but before this it is also worth noting that it is not entirely clear whether the wrongs involved in the motivating cases are cases of *directly* wronging someone with a belief. In the case of John Hope Franklin, if the belief itself wronged him independently of the offensive behaviour it engendered, then one would expect such a belief to have wronged him even in a case where the behaviour that resulted was benign. To probe this

---

[9]Consequentialists will be suspicious of the whole category of directly wronging, but for the purposes of this paper we are happy to grant that there are cases of direct wronging with attitudes.

hypothesis, let us construct a new, fictional but analogous case, as follows:

Universe Club

Sally is to visit the Universe Club, a similarly exclusive establishment to the Cosmos Club, and one at which the vast majority of persons of a marginalized ethnicity at the club are staff as opposed to members, as a result of systematic discrimination against that ethnicity. Sally is also a member of this ethnicity and a union organizer.

Sally visits the club and notices Jim, a member of the ethnicity that at this club overwhelmingly comprises of staff. She assumes him to be staff and introduces herself, with the intention of promoting the unionization of staff at the club. In fact she is mistaken. Jim is a member of the club and an eminent scholar.

It is far from clear that Sally has wronged Jim in this case.[10] But the belief is formed on an exactly analogous statistical basis as the woman's belief in John Hope Franklin's case, and mischaracterizes Jim in an exactly analogous way. Thus, it seems that the primary way in which John Hope Franklin was wronged was by factors other than the belief itself. Clearly, one major factor is the behaviour that was caused by a false belief, in particular the presumptuous demand. Perhaps Franklin was also harmed independently by other factors surrounding the belief, such as the racist attitudes of the believer. In either case, it seems quite unclear that the belief itself directly wronged. This is to say, it seems that on the guarded interpretation of the quoted harmony principles, the principles are orthogonal to Gendler's primary example of a morally problematic but epistemically impeccable belief.

In the remainder of this section we shall argue that there is no way of interpreting the harmony principles which renders them plausible while also respecting any of the central judgments used to motivate the literature.[11] One straightforward interpretation is Schroeder's, which is easy to refute because he admirably presents a clear necessary and sufficient condition for a belief wronging in the relevant sense, namely that it wrongs just in case it falsely diminishes. After dispensing with the harmony principles interpreted in accordance with Schroeder's theory, we turn to

---

[10]Basu (2019, p. 14) presents a similar case to Universe Club called *Enlightened Northerner*, and judges that a wrong is still committed. However, her case has an additional element of the eponymous 'enlightened' US northerner manifesting a disposition to think of the staff at a club in a problematically objective way. Universe Club avoids this element, so seems to isolate better the belief itself.

[11]Here is one way of giving up standard judgments from the literature concerning which beliefs harm: one might hold that a belief directly harms only if it is unjustified by traditional standards. This would not support moral encroachment. Rather, it would suggest that epistemic factors encroach on the moral.

refuting the harmony principles on the assumption that we are to hold fixed the central judgments of wronging to which they are supposed to apply, such as John Hope Franklin's case. We leave open the possibility of holding onto the harmony principles by dropping both Schroeder's theory of when beliefs wrong and the judgments about particular cases held in the literature on moral encroachment, since this option is of no interest for those attempting to establish moral encroachment in the cases central to the literature. (And, since the harmony principles are obviously false on the literal interpretation, there seems little reason to adhere to them without a particular theoretical need which might be hoped to clarify their best interpretation.)

First is Schroeder's theory, wherein a belief wrongs if and only if it falsely diminishes. This is to be combined with his thesis that 'the beliefs that wrong are never epistemically impeccable'. We take it that by 'epistemically impeccable' Schroeder means epistemically justified or something similar.[12] The problem is that the condition for wronging is not sensitive to how overwhelming the evidence in favour of the belief might be. On this theory there are certain beliefs, namely those which falsely diminish, which are *ipso facto* not epistemically justified. Thus, it is possible to ramp up the traditionally recognized factors in favour of epistemic justification to an arbitrarily high degree while keeping the content of the belief fixed, so Schroeder is absurdly forced to deny epistemic justification is present even in extreme cases. (This argument will generalize to make trouble for harmony principles on any view of doxastic wronging according to which, if some belief wrongs, then so does every belief with the same content.)

A belief diminishes when it makes out an agent's "contribution to be less". A particularly clear case is is one where someone achieves something through their hard work, but is believed to have only achieved it through nepotism. The belief falsely diminishes them because while their agential contribution was in fact great, the belief has it that their agential contribution was small. Thus, according to Schroeder, a false belief that someone achieved a position through nepotism is never epistemically justified. This is absurd. It rules out a historian having a justified but false belief about it being easy for the ancient Chinese monarch Cheng to succeed his father as King of Zhou. The manuscripts and all other evidence might falsely indicate that the succession went smoothly, as a result of an elaborate conspiracy concealing a true history where the he had to fight tooth and nail. Believing the manuscripts would falsely diminish King Cheng's agential contribution. But it is wrong to think that historians never get justified beliefs by believing manuscripts

---

[12]There is some sense in which any false belief is not epistemically impeccable, since by believing a false thing one makes a mistake in some sense, but this cannot be the operative reading, because Schroeder's harmony principle would be trivially true without any moral encroachment, since he thinks that beliefs that wrong are false.

that falsify history in this way.

Schroeder is the main contributer to the literature who adheres to a harmony principle with a specific theory about when beliefs wrong. The rest of the literature relies primarily on paradigm cases, such as that of John Hope Franklin. We now turn to refuting harmony principles on the assumption that whatever 'wronging' amounts to, beliefs in cases like this wrong in that sense. The problem in this case is that harmony principles of this nature are difficult to square with any systematic thought about abstract principles governing epistemic justification.

The first principle posits that a belief is justified when, in normal cases, it would amount to knowledge.

KJ Link If a belief is formed in a way such that it would be knowledge in all normal cases, then that belief is justified.[13]

KJ Link is supported by the thought that someone who believes in a way that would be knowledge in normal circumstances is doing what they are supposed to, epistemically speaking. Someone who forms a belief in such a way is merely unlucky, rather than irresponsible, if it turns out to be false—their belief is still epistemically justified. To illustrate the principle, consider two cases. In the first case, you see Ned steal a book and come to know that he is a thief. In the second, you see someone stealing a book who, by a remarkable coincidence, looks exactly like Ned, and you falsely come to believe that Ned is a thief. The second case is extraordinary, in that you form a belief in such a way that it would be knowledge in all normal cases, but that belief turns out to be false. KJ Link, plausibly says that even in this extraordinary case your belief is justified (indeed, this kind of case is a paradigmatic example of justified but false belief in the literature).

Harmony is in tension with KJ Link for the following reason: in extraordinary cases, one can wrong someone by believing something on a basis that, in ordinary cases, would result in knowledge. By KJ Link, such a belief must nevertheless be epistemically justified. In particular, it seems that in the second Ned case you wrong him by thinking him a thief. But by KJ Link, your belief is justified.

The second structural constraint concerns the transmission of justification by competent deduction.

---

[13]KJ Link bears similarities to various ideas in the literature. For example, Bird (2007) suggests the following criterion for justification:

> If in world w_1 S has mental states M and then forms a judgment, that judgment is justified if and only if there is some world w_2 where, with the same mental states M, S forms a corresponding judgment and that judgment yields knowledge.

Notice that KJ Link is apparently much weaker than Bird's principle, since it is not a biconditional, and the antecdent requires the belief to constitute knowledge in a wide range of cases, namely all the normal ones, rather than just one.

KJ Closure If an agent knows for sure that φ and justifiedly believes that ψ, then they justifiedly believe anything they believe on the basis of competent deduction from φ and ψ.

Abstractly, KJ Closure is in tension with harmony because one can reach a morally problematic belief by competently reasoning on the basis of morally unfortunate knowledge and a morally unproblematic, justified, but false belief.

Let us now consider a concrete example where these principles generate problems for harmony.

Racist Club

A (fictional) club previously adhered to a strict racial segregation policy: only white people could be members, although there are and were many non-white staff. Due to the previously racist owner having an entirely unpredictable moral revelation, this policy has been abolished as of today, and Jacob is the very first non-white member of the club.

Jerry visits the club and falsely believes the segregatory policy is still in effect, having seen it in effect yesterday and presuming that the owner has not had any sudden change of heart. He sees Jacob, assumes him to be staff on the basis of his race, and demands he fetch his coat.

In Racist Club, Jerry wrongs Jacob. The authors' view is that he wrongs Jacob primarily by his presumptuous demand, but if the wrong in John Hope Franklin's case is a case where the belief itself wrongs, then surely also Jerry's belief that Jacob is staff wrongs Jacob in the way that, given harmony, should rule out the belief being epistemically justified.

However, we may argue from either KJ Link or KJ Closure that Jerry's belief is epistemically justified. Using KJ Link, we may point out that if things had been proceeding as usual and there had been no moral revelation that abolished the racist whites-only policy, then Jerry would have known that any non-white person at the club was staff. To deny that Jerry would know in that case is rank skepticism: of course one can know the consequences of a longstanding and stable policy in normal cases, even when that policy is racist. So, by KJ Link, Jerry's false belief is epistemically justified in Racist Club.

Using KJ Closure, it will suffice to establish the following:

(a) Jerry knows for sure that Jacob is not white.

(b) Jerry knows for sure that the policy was in effect yesterday, and is being followed today unless the previously very racist owner had a sudden change of heart.

(c) Jerry is justified in believing that the previously racist owner did not have a sudden change of heart.

Jacob is not wronged by Jerry's knowledge that he is not white. He is not wronged by Jerry's knowledge of the policy being in effect until yesterday. He is also not wronged by Jerry's false belief that the owner did not have a change of heart. Presumably, the standard for justification for this final belief is not unusually high. Surely Jerry can be epistemically justified in believing that the owner did not have a sudden change of heart: the reader can put themselves in Jerry's epistemic position and recognize that it would be bizarre to entertain any doubt that the previously committedly racist owner of a club is still racist. Jerry believes Jacob is staff on the basis of deduction from these premises, so by two applications of KJ Closure, the belief that Jacob is staff is epistemically justified.

Denying KJ Link or KJ Closure is an unpromising way of defending harmony. Of course one might quibble with both principles for various reasons, but it is fairly clear that any systematic thought about epistemic justification will reveal principles which cause analogous problems. It is very difficult to deny that Jerry's belief in Racist Club is epistemically justified, for anyone who wants to go on to use epistemic justification in any serious theoretical capacity. There is also surely some good status a belief can have that obeys these or related principles, whether or not we call it 'justification'. That sort of status is surely what epistemologists care about, and it is what they mean to ascribe when they laud beliefs as 'justified' or 'rational' or 'appropriate'. Moral encroachment on a notion of 'justification' is only epistemologically significant insofar as that notion is one epistemologists care about.

A better option for the defender of harmony is to deny that Jerry's belief wrongs Jacob. The problem, of course, is that the case of Racist Club is so closely related to the original motivating case, of John Hope Franklin at the Cosmos Club. The main difference is the following: in Racist Club, it is more obvious that Jerry's belief is justified by the standards of traditional epistemology. In the original case, the level of justification from the perspective of traditional epistemology is not clear, since it is not clear how strong her statistical evidence was that any African American man at the club would be a staff member, and Franklin was not wearing an employee uniform. Insofar as her belief would be unjustified even from a traditional perspective, the appeal to harmony principles does exactly nothing to militate in favour of moral encroachment. Insofar as it becomes clear that her belief was justified from a traditional perspective, the case starts to look more similar to Racist Club, so it is less clear that the application of the harmony principle is plausible.

4 Moral Encroachment Without Harmony

The foregoing considerations show that the harmony principles used by Bolinger, Schroeder, Basu, King and Babic, and Jackson and Fritz to argue for moral encroachment are implausible, holding fixed either Schroeder's theory of when beliefs wrong, or their shared judgments about which beliefs wrong in particular cases. A force of moral encroachment strong enough to bring about harmony would make a mess of all previous theorizing about epistemic justification, with cases where it is unjustified to derive conclusions from justified premises, or to believe something on a basis that would in all but the most bizarre cases would amount to knowledge. The stated motivation for moral encroachment is in bad standing.

Nevertheless, even with a thoroughgoing harmony off the table, one might still hope for a greater degree of harmony than traditional epistemology suggests. Moral encroachment might be posited to impugn some cases of morally problematic belief as epistemically unjustified, contrary to traditional epistemology, without going so far as to rule out morally problematic but epistemically justified belief in extreme cases like Racist Club.

It becomes less clear that moral encroachment makes a difference in the motivating cases, such as the case of John Hope Franklin at the Cosmos Club. After all, if the woman's degree of justification that Franklin was staff would be very high absent moral encroachment, then her belief could still be epistemically justified, just to a lesser degree, in the actual case. Thus the proponent of moral encroachment would have to agree with the authors and Gendler that the moral and epistemic considerations pull apart in the very cases which motivated them to say otherwise. Nevertheless, it is not a disaster for the proponent of moral encroachment that the belief in the John Hope Franklin case *could* be epistemically justified on some ways of filling in how things played out. It would still be an interesting discovery that any encroachment happens, even if it is hard work to show it makes the difference between justified and unjustified in particular cases. However, the prospects for even a small effect of moral encroachment are dim.

5 Moral Encroachment on Credence

We shall now argue that the effect of moral encroachment must be negligible. The argument will concern the epistemically most appropriate credences for an agent to have. In this section we shall argue, in agreement with Jackson and Fritz (2021), that any significant effect of moral encroachment on any doxastic attitude would have a significant effect on the epistemic appropriateness of morally problematic credences. Then we will show that any effect of moral encroachment on credence

would have to be negligible, given widely accepted constraints.

It is worth noting first of all that although the usual proposed mechanism for moral encroachment discussed in Section 2—that high moral stakes raise the evidential standard required for justified belief—is specific to belief, Gendler's motivating examples and the examples used by proponents of moral encroachment involve attitudes besides belief. In Handbag, Sally certainly does not believe she will be robbed—if she did, she would not merely grip her bag more tightly—she merely has a slight worry or suspicion that she will be robbed. In Assault, the problematic attitude is Mr Y's *lack* of belief in Miss X's testimony, rather than an outright belief that Miss X was not assaulted. And in Workplace, Victor gives little credence to the two women being in management, but need not be supposed to believe that they are not. Whatever link there is between suspecting or doubting something and action is very different from that between (outright) believing that thing and action. Nothing like the Knowledge-Action Link principle comes even close to holding, since it is obvious that merely suspecting something does not license treating that thing as guaranteed in deciding how to act. The sensibility suggesting direct wronging in the John Hope Franklin case surely also suggests direct wronging in these cases too.

Other motivating cases in the literature also do not involve only belief, or even primarily belief. For example:

> The human resources manager who wrongly suspects the young married female applicant of seeking the job for its maternity benefits, the jealous wife who wrongly suspects her husband of having a wandering eye, and the father who doubts his daughter has what it takes to succeed in a career in engineering—all wrong the subjects of their beliefs by what they believe about them. (Schroeder 2018, p. 114)

(Schroeder describes suspicions and doubts as 'beliefs'. In ordinary English this is acceptable—see Hawthorne, Rothschild, and Spectre (2016)—but a suspicion or doubt is not automatically a belief in the philosopher's sense of an *outright* belief, which requires outright commitment to the proposition fully believed. It is plausible that justified outright beliefs can require a very high evidential standard for reasons relating to Knowledge-Action Link, but not at all plausible that suspicions or doubts could.)

Moroever, Basu and Schroeder (2019) are explicitly concerned with belief, but it is evident that in many cases their arguments would work identically if belief were replaced with merely high credence.

> The racist is paradigmatically disposed to be influenced by her perceptions of race in the beliefs that she forms about another person – more

easily persuaded that someone is dangerous, for example, if they are perceived as black. Racist beliefs are naturally taken not just to be morally problematic, but specifically to wrong their subjects. Moreover, an apology for crossing the street to avoid you by someone who still believes that you are dangerous would strike us as insincere, as would an apology for not checking more carefully before forming this belief, by someone who still holds the belief. So that is again at least some prima facie evidence that at least an important part of the wrong lies in the belief, rather than acts leading up to or subsequent to it. (Basu and Schroeder 2019, p. 183)

Being more easily persuaded that someone is dangerous if they are perceived as black is a property one has when one forms a higher credence in someone's being dangerous conditional on their being black. Similarly, being motivated to cross the street is something that would be expected by someone who forms a non-low credence in their being attacked (someone who outright believes they will be attacked will run in the opposite direction, rather than cross the street). And having a high credence that someone is dangerous is generally extremely offensive, whether or not an outright belief is formed. So if the quoted passage tells in favor of moral encroachment on belief at all, then it should also tell in favor of moral encroachment on credence.

Moss (2018a) is more explicit that moral encroachment should apply to "opinions" generally, understood so as to include attitudes besides belief:

[...] opinions are defined to include all doxastic states studied by either traditional or formal epistemologists—full beliefs, for instance, as well as credences and other probabilistic beliefs. As they form their opinions, these subjects all engage in racial profiling, where my use of this term in this paper is restricted to a doxastic practice—namely, forming opinions about a person on the basis of statistics about members of their racial group. As many advocates of moral encroachment have noted, accepting moral encroachment can help us identify one respect in which the practice of racial profiling is problematic. (Moss 2018a, p. 178)

Proponents of moral encroachment are motivated, in opposition to Gendler, to rule out problematic opinions as epistemically inappropriate, despite their apparently being justified according to traditional epistemology. This project requires the generalization of moral encroachment to non-belief attitudes. There is no plausibility at all to the idea that one can avoid forming racist, sexist, or otherwise problematic opinions by meticulously avoiding being fully committed to (i.e., fully

believing) those opinions. Differential suspicions or hypotheses of wrongdoing, doubts or lack-of-belief regarding ability, and other doxastic attitudes, are paradigmatically what a racist, sexist, or otherwise morally dubious person holds. Insofar as moral encroachment is an important factor that renders epistemically unjustified racist, sexist, or otherwise problematic doxastic attitudes where traditional epistemology fails to do so, it must act on these non-full-belief attitudes. (So if an account of doxastic wronging only allows full beliefs to wrong, so much the worse for that account.)

Once it is conceded that moral encroachment acts on a wide range of doxastic attitudes besides full belief, it would be odd to deny that it acts on credences, since credences are just another doxastic attitude. However, we wish to argue also that any effect of moral encroachment on the epistemic appropriateness of doxastic attitudes generally, such as doubting, suspecting, having little credence, should be reflected in an effect on the epistemic appropriateness of credences. Our argument is twofold. First, given that the commonly proposed mechanism for moral encroachment is the same as that for pragmatic encroachment, namely that the evidential standard for belief is raised when the stakes are high, has no bearing on non-belief attitudes, that mechanism cannot be the general mechanism of moral encroachment. It then seems very natural to locate the general mechanism for moral encroachment in its effect on credence, because credence presumably has a tight constitutive connection with doubting, suspecting, and so on (e.g., you can only doubt something if you don't have very high credence in it, you can only suspect something if you have non-negligible credence in it, etc.). It would be very odd to posit a morally-sensitive shifting evidential standard required for doubting, then another one for suspecting, then another one for having little credence, with nothing explaining these all at once. Second, given the tight connection between credence and these other attitudes, it seems impossible for the appropriateness of them to vary without corresponding variation in appropriateness of credences. If the right level of credence for you to have in two propositions is equal, then surely it is equally appropriate for you to doubt/suspect/have little credence in either of them.

6   Against Moral Encroachment on Credence

If there is moral encroachment, then it has a non-negligible effect on which credences are appropriate to have. This section will establish that there is no such effect, so morality does not encroach.

The argument will be made through variations on the Workplace case of Section 1. There, if moral encroachment manifests at all, it should increase Victor's credence that the women outside are in management by at least a little bit. Perhaps

traditional epistemology recommends a credence of 1%, but with moral encroachment taken into account this shoots up to 30%.[14] Crucially, the arguments will not be highly sensitive to how big a difference moral encroachment makes. Even if it only recommends bumping up the Victor's credence from 1% to 1.2%, the following arguments will apply to show that there is no such encroachment.

The argument will proceed using a categorization of cases following Johnson King and Babic (2020). In cases of *direct inference*, the agent forms a credence about a particular member of a population on the basis of known population-wide statistics. To make the Workplace case a case of direct inference, we must suppose that Victor knows exactly what proportion of women in the company are in management. In a *predictive inference* case, the agent forms a credence about a particular member of a population on the basis of known statistics about a *sample* from that population that is not presumed to include the particular member. To make the Workplace case a case of predictive inference, we could suppose Victor has knowledge of a sample of women in the company, and knows how many from that sample are in management.

6.1 Direct Inference

Direct Inference Workplace

Victor is a visitor to a company where he is eager to impress people in management. He also knows that of the 1000 women in the company, 10 (i.e., 1%) are in management, whereas 30% of the 1000 men in the company are in management. Sitting in an office, he hears a woman, Joan, speaking outside. He relaxes, giving little credence to her being in management. In fact, Joan is in management.

From the description of the case, traditional epistemology cannot reach a recommendation about the credence Victor should adopt that Joan is in management. Certainly it is *not* true that, in any case consistent with the description of Direct Inference Workplace, Victor's credence that Joan is in management should be 1%.

_____

[14]Admittedly, there is a much milder version of moral encroachment that does not *require* the credence to be different to what is recommended by traditional epistemology. The idea of this milder version is that moral factors increase the range of permissible credences without rendering the traditionally recommended credences impermissible. This takes away one version of Gendler's tragedy, since the morally problematic credence associated with traditional epistemology is no longer required (though there remains a milder tragedy according on which morally problematic credences are still permitted). This milder version of moral encroachment is suggested by some of the language of Johnson King and Babic (2020). The lessons of our discussion will largely carry over to the milder version, since, for example, our arguments against adopting the grossly partial credences seen later in this section seem to show they are impermissible.

For an obvious counterexample, we might consistently add to the case that, when he overhears Joan, he overhears her telling someone her rank as management. In that case, one would certainly expect a high credence that she is management. But Victor needn't get such straightforward evidence for his credence to diverge significantly from 1%. Maybe he knows that most people with a certain accent are in management, so that by hearing Joan's voice with that accent he gets significant information. Perhaps also Victor antecedantly expects a person in management to be outside, because he is waiting to be interviewed by someone in management. In that case, the fact that Joan is outside may be some evidence that she is in management. Whether this should take Victor to 2% credence or 90% will depend on further details.

For moral encroachment to make a difference, then, it is not sufficient to show that Victor's credence that Joan is in management should be much higher than 1%, since this is entirely compatible with traditional epistemology. What moral encroachment requires is that Victor's credence should, epistemically speaking, be higher than whatever is recommended by traditional epistemology. To fix ideas, then, we will assume that the evidence Victor has does not, from a traditional perspective, particularly favor or disfavor Joan's being in management over any of the other women in the company. In that case, traditional epistemology will recommend a credence somewhere around 1%, and we can suppose moral encroachment recommends something higher, be it 2% or 30%.[15]

The argument against this discrepancy relies on the principle of Probabilism.

Probabilism Any epistemically appropriate credal state is one where your credences are representable by a probability function.

Probabilism is very widely accepted. It is also explicitly endorsed by many proponents of moral encroachment in the literature, including Johnson King and Babic (2020), who argue for moral encroachment on credence.

It also appeals to the following elementary fact.

Fact 1. *If a probability function* Cr *assigns* 1 *to the proposition that exactly n of* $x_1 \ldots x_k$ *(with $x_i \neq x_j$ for $i \neq j$) are F, then the average value of* $\mathrm{Cr}(x_i$ *is* $F)$ *is* n/k.

In light of Probabilism and Fact 1, any epistemically appropriate credence function for Victor gives an average credence of 0.01, or 1%, to any particular woman's

---

[15]It has been suggested to the authors that moral encroachment will be immediately ruled out by the *Principal Principle*, which says that one's credence in a proposition should be equal to the expectation, on one's credences, of the objective chance of that proposition (Lewis 1980). The authors agree that there is *prima facie* tension between this principle and moral encroachment. However, since moral encroachment could also act on one's credences in the objective chance of a given proposition, an effective argument against moral encroachment using the Principal Principle is not totally straightforward to construct.

being in management (assuming, in accordance with the setup of the case, that he rightly has 100% credence in the statistic that 10 of 1000 women are in management, and 100% credence that Joan is a woman in the company).

Any credence above 1% that Victor has in Joan's being in management comes at the cost of some other woman in the company. For the average to stay at 1%, Victor's credence for some other woman's being in management must be lower than 1%, if his credence for Joan is higher.

Given Probabilism, it is simply impossible for the epistemically appropriate credences to be more charitable overall. It is consistent to recommend a high credence in Joan's (or any other particular woman's) being in management, but not enough room is left for moral encroachment to recommend being an overall more charitable person. In the end, on average, women are given 1% credence of being in management, and moral encroachment cannot change this without overthrowing Probabilism. For moral encroachment to recommend a higher-than-traditionally-recommended credence that Joan is in management therefore requires *partiality* to Joan on Victor's behalf. But partiality of this sort is cold comfort for the proponent of moral encroachment, since it is morally objectionable in addition to being epistemologically bizarre. It is objectionable because it is unfair. One is not being an admirable person by arbitrarily being overly charitable to the people one has heard in the corridor at the cost of others.

This argument does not depend on how much Victor's credences are supposed to be modified by moral encroachment. Rather, once one realizes that moral encroachment requires apparently arbitrary partiality, rather than an overall increase in charitability, it seems that any nonzero effect is implausible. Victor's credences should be controlled by traditionally recognized factors, *contra* the thesis of moral encroachment.

The argument also does not depend on Victor being in a situation where traditionally recognized factors point to a credence of 1% in Joan's being in management. Johnson King and Babic (2020) make a great deal of the fact that, in cases of direct inference, one's credences about a particular member of a population need not be controlled only by the overall statistics of that population.

> Since genuine conditions of exchangeability[16] are vanishingly rare in real-life cases, we are not worried about direct inference. The circumstances in which someone could be epistemically required to conform her credences to observed population frequencies are possible, but extremely unlikely to arise. (Johnson King and Babic 2020, p. 101)

---

[16]Conditions of exchangeability hold, in Johnson King's and Babic's sense, when all members of the population are given the same credence of exhibiting the relevant property.

This is a red herring. Conditions of exchangeability are irrelevant to the present argument. The argument is most straightforward under the very natural supposition that traditional epistemology points to approximately equal credence in each woman's being in management, but nothing hangs on this. What is objectionable is the partial treatment of some over others beyond traditionally recognized evidence, and any modification to Victor's credence function suggested by moral encroachment is *ipso facto* an instance of this partiality. Exchangeability is besides the point.

6.2   The Denial of Probabilism

The proponent of moral encroachment might choose to deny Probabilism. Of course Probabilism is not sacrosanct, and it has faced a range of serious objections and promising competitors (for example, imprecise credence models). But it is hard to take seriously the denial of Probabilism for the sake of moral encroachment, a hypothesis which, although interesting, is yet to demonstrate any value in the broader scheme of epistemology.

More importantly, a proponent of moral encroachment who merely denies Probabilism will quickly be forced to confront the more fundamental principles of credence that have traditionally been used in justifying Probabilism. Which axioms of Savage (1956), or Kraft, Pratt, and Seidenberg (1959), for example, would moral encroachment justify rejecting? Note that moral encroachers seem to be saying that even ideally rational agents should depart from the recommendations of traditional epistemology. If they are to do so by violating Probabilism, then the most worrying objections to these axioms, which turn on the limits of ordinary rational agents, will be neither here nor there.

The importance of the fundamental principles underlying Probabilism is made especially clear when credence is connected to action. Thus, let us modify the Direct Inference Workplace case so that Victor must make a series of 1000 choices, one for each woman in the company, that essentially amounts to a bet on whether she is in management. Abstractly, we may say Victor may "accept" or "decline" to perform a certain action for each woman, so that accepting yields a very good payoff if the woman in question is in management, and a neutral payoff otherwise, and so that declining yields a very slightly good payoff no matter what. We assume that the very good, slightly good, and neutral payoffs have the following decision-theoretic structure in a gamble disconnected from moral encroachment (e.g., a series of lotteries known to be fair):

(i) A 1% chance of the very good payoff and a 99% chance of the neutral payoff would be worse than a 100% chance of the slightly good payoff. However it is a close thing, since a slightly higher than 1% chance of the very good payoff

with a slightly less than 99% chance of the neutral payoff is better than 100% chance of the slightly good payoff. This is to say if the neutral outcome has utility 0 and the slightly good payoff has utility 1, the very good payoff is to have utility slightly less than 100.

(ii) (i) holds for each of the 1000 choices independently of which payoffs are ultimately received. This is to say, the utility of receiving the very good payoff m times, the slightly good payoff n times, and the neutral payoff $1000-(m+n)$ times, is $xm + n$, where x is the utility of receiving the very good outcome (x, recall, is presumed to be slightly less than 100).

We submit that Victor would be a fool to accept all 1000 times. It is known that there are 10 women in management of the 1000, so accepting every time would guarantee an outcome worse than declining every time, since the very good outcome is less than 100 times better than the slightly good outcome. So unless Victor is a fool he will decline at least once. However, we also submit that declining in a given case yields a "revealed credence" that is theoretically speaking indistinguishable from a credence that is at most very slightly greater than 1%. If you decline a bet that costs 1 utility to play and has a payoff of either zero or something just under 100, then there is very little to distinguish you from someone who has a credence around 1% or less that accepting would yield the very good payoff. Indeed, no distinction is usually made between an agent's credence and their revealed credence, if a revealed credence can be found,[17] and endorsing moral encroachment on credence is surely not a good reason to make such a distinction. Insofar as a credence of 1% wrongs Joan, surely also a revealed credence of 1% wrongs Joan for the exact same reason, since it amounts to Victor having a doxastic state which inclines him to accept bets that effectively treat Joan as being very probably not in management, on the basis of her gender. Moreover, by distinguishing credence from revealed credence, the importance of credence is drastically reduced: it no longer is held to connect to action in the traditional Bayesian way; rather revealed credence bears that connection. So surely revealed credence is the much more important doxastic state of a person. To admit moral encroachment acts on credence but not revealed credence would be to correspondingly curtail its importance.

Thus, we may concede to moral encroachment the denial of Probabilism for the sake of the argument. Perhaps Victor should have 10% credence in each woman's being in management, being vastly more charitable in every case than what traditional epistemology would recommend, at the cost of a massively non-probabilistic credence function. Or perhaps he should refrain from forming credences about any of the women, thus avoiding forming a *low* credence about any (e.g., by forming

---

[17]Ramsey 1926.

imprecise credences instead). All of this doesn't help avoid the problem that in cases of direct inference, charity to some comes at the cost of others, since no matter his credences, if Victor is not a fool then his revealed credences will not be much greater than 1% for at least some women in the company.

6.3   Predictive Inference

Predictive inference cases are those where a credence is formed on the basis of known statistics about an incomplete sample of the population. Unlike cases of direct inference, it is very common to find oneself in a predictive inference case, by learning some statistic about a sample from a population (e.g., by accessing survey data) and then considering a particular member of that population. As such, it is important for moral encroachment that it makes a difference in predictive inference cases, even once it is recognized that it makes no difference in direct inference cases.

> **Predictive Inference Workplace**
>
> Victor is a visitor to a company where he is eager to impress people in management. He knows that of the 1000 men at the company, 30% are in management. He also has a roster of the 1000 women at the company (in fact, 10 are in management, but he does not know this). He selects 300 women seemingly at random from the roster, and queries someone in the know if they are in management, and learns that none of those 300 are. Later, he hears Joan outside, and relaxes, since he can tell she is a woman and on this basis has little credence that she is in management. In fact, Joan is in management.

In Predictive Inference Workplace it would be inappropriate for Victor to have credence 1 in a particular number of women in management, so Fact 1 does not apply. Instead, the case against moral encroachment will be made by appeal to a slightly more general fact about probability functions:

Fact 2. *For any probability function* Cr, *the average value of* $\mathrm{Cr}(x_i\ \text{is}\ F)$, *for* $i = 1$ *through* k, *is equal to the expectation, on* Cr, *of the number of* x_1 *through* x_k *that are* F.

Thus, if in light of his apparently random sampling the expected number of women in management is, on Victor's credences, less than 10, Fact 2 then implies that Victor's average credence in a given woman's being in management must be less than 1%. If so, then given Probabilism, moral encroachment is prevented from recommending more charitable credences for Victor, exactly as it is in the direct

23

Number of red balls of 1000	10	20	30	40	50	100	500
Prob. of 0 red of 300 draws	0.028	10^{-4}	10^{-5}	10^{-7}	10^{-8}	10^{-17}	10^{-119}
Prob. of 0 red of 200 draws	0.11	0.01	10^{-3}	10^{-4}	10^{-5}	10^{-11}	10^{-71}
Prob. of 0 red of 100 draws	0.35	0.12	0.04	0.01	10^{-3}	10^{-5}	10^{-33}

Table 1. Approximate probabilities of finding 0 red balls from 300, 200, or 100 draws from an urn with 1000 balls, of which 10, 20, 30, 40, 50, 100, or 500 are red. (The probability distribution is called the hypergeometric distribution.)

inference case. Therefore, if moral encroachment makes a difference in predictive inference cases, then it must act by recommending credences for Victor with a significantly higher expected number of women in management. We shall now argue that there is no such effect, since it would be manifestly epistemically irrational, as well as morally irresponsible, for Victor to have a high expectation of the number of women in management.

Some basic statistics will put the example in context. If one is selecting fairly at random, without replacement, from an urn of 1000 balls, the probability of not finding a single red ball after 300 selections depends on the number of red balls as given in the first row of Table 1. By inspection one can see that with random selection, finding 0 red balls from 300 draws overwhelmingly favors a low number of red balls among the 1000 balls in the urn, for example it favours there being 10 or fewer about a million billion times more than there being 100 or more. So unless one starts out overwhelmingly confident in a high number of balls in the urn, drawing 300 times and finding no reds will lead to drastically lower expectations for the number of red balls in the urn.

A similar point holds in Victor's situation. If the selection of women is random from the perspective of Victor's credences, which is to say that each woman is equally likely to be selected on his credences, and the pattern of selection is independent of which women and how many happen to be in management, then conditionalizing on none of the 300 picked being in management would crater the expectation of the number of women in management, unless for some reason he starts overwhelmingly confident that there are many women in management (such as being a million billion times more confident that there are 100 or more woman than that there are less than 10). Such overwhelming confidence would be manifestly irrational if Victor doesn't have any particular indication of the number of women in management to begin with. It would also be morally problematic: someone who is overwhelmingly confident in approximately even gender ratios in socially desirable positions will be oblivious to the systematic factors that prevent women from reaching such positions.

Therefore, if Victor's credences are probabilistic before he makes the seem-

ingly random selection, and the selection is random from the perspective of his credences, and he then updates by conditionalizing on the realization that none of the 300 women are in management, the expected number of women in management will be very low, so the arguments used in the direct inference case against moral encroachment will apply.

The proponent of moral encroachment can again deny one of the general structural principles: either Probabilism or the thesis that Victor should update by conditionalization. For the reasons already mentioned, this sort of move is extremely unpromising (no argument specific to conditionalization will be supplied here; it is just another principle that one could deny if one did not care about doing catastrophic damage to epistemology to save moral encroachment). However, there is a much more plausible way out, which is to deny that Victor's credences ought to treat the selection of women as random. In fact, it is clear that barring exceptional circumstances, one's credences ought to treat an apparently random sampling as not perfectly random. Even in the case of the red balls in the urn, picking 300 non-red in a row speaks very strongly against there being 500 red balls of 1000 in the urn, but still not nearly as strongly as the 10^{-119} in Table 1 suggests. After all, in normal circumstances one should apportion some credence much greater than 10^{-119} that the urn has not been thoroughly mixed, and that you don't pick any red balls because they are all at the bottom of the urn. This region of credence-space can only be favoured by learning that none of the 300 balls drawn were red, so there remains hope for you to retain a high credence that many red balls are in the urn.

If moral encroachment prevents Victor from having a low average credence in each particular woman's being in management, it must do so by preventing him from having a low expectation in the number of women in management. Since the regions of Victor's credence-space where the sampling of women is random and where there are a high number of women in management are massively disfavored by none of the 300 women turning out to be in management, the majority of the region of credence-space where there are a high number of women in management must be a region that, before learning that none of the 300 women in management, did not treat the sampling of women as random. That is, moral encroachment, if it exists, weights Victor's credences in favor of hypotheses to the effect that his selection process was biased in some way to select women who are not in management

6.4   *Conspiratorial Thinking for Moral Encroachers*

We shall now argue that this weighting in favour of hypotheses of biased sampling amounts to the claim that we should assign alarmingly high credence to outlandish conspiracy theories. The arguments for moral encroachment are not strong enough to make this plausible, so moral encroachment should be rejected.

Certainly, one should not unthinkingly assume one's sampling method to be random.[18] This point is central to statistics and traditional epistemology; anyone sampling anything must be aware of it. However, what is also clear is that it is not usually extremely difficult to have a high credence that a given sampling process is fair; i.e., to have high credence in some proposition conditional on which your credence treats the sampling as random. In the case of drawing balls from the urn, one might have various worries about the randomness of the selection. But it is not hard to take steps to ensure the selection is very probably random. For example, one can ensure that the urn is well-mixed before sampling, that the red and non-red balls are indistinguishable except for their color, and so on. If reasonable precautions are taken, it just would be irrational to draw 300 non-red balls and not end up with a very low expectation of the number of red balls in the urn.

The same goes for taking an apparently random sample of women. Is it really so hard to be approximately random when selecting names off a roster? Hypotheses that support bias seem outlandish. Should Victor think he can somehow tell by looking at someone's name whether she is in management? Or should he think that there is a shadowy conspiracy subtly guiding his hand to choose women who happen to not be in management? Hypotheses such as these must be highly magnified in Victor's final credences, if he retains a high credence in a high number of women in management after selecting 300 not in management. But having high credence in such conspiracy theories seems outright bizarre, and not the sort of thing that should be recommended by a sound epistemological theory.

The same goes in cases like Predictive Inference Workplace that occur in real life. In real life it is fairly easy to get statistical knowledge about a sample from officially recognized sources, such as the US Bureau of Labor Statistics, that may or may not support the idea that women in the US are less likely to be in management in certain sectors of the economy (overall, women made up just over 45% of those in "management occupations" in the US in 2021 (U.S. Bureau of Labor Statistics 2023, Table 11), so the company in the Workplace case does not reflect the overall economy of the US). Insofar as such a source does find that the sampled

---

[18]Thus, our argument will not appeal to the following incontrovertibly false principle that Johnson King and Babic correctly deride:

Frequency-Credence Connection If (a) I know that a is an F and (b) know that x% of previously sampled Fs are G, and (c) I have no further evidence bearing on whether a is G, then my credence that a is G should be x%. (Johnson King and Babic 2020, p. 91)

Counterexamples to this principle are easy to generate. A straightforward one is where a is Charles III, F is the property of being English, G is the property of being identical to Charles, and less than 100% of English people sampled were identical to Charles (i.e., someone other than Charles was in the sample). Where Johnson King and Babic go wrong is in thinking that denying the Frequency-Credence Connection one makes room for moral encroachment.

women were mostly not in management, then by analogy with Predictive Inference Workplace, moral encroachment ought to increase our credence that the sampling was not done fairly. Thus, if moral encroachment is right we should write to the US Bureau of Labor Statistics and other government agencies to inform them about the previously unrecognized probability that their sampling process was biased in some way. Of course, such a letter would rightly be dismissed by anyone serious.

In addition to being epistemically bizarre, the recommendations of moral encroachment in the Predictive Inference Workplace case seem deeply morally problematic. They are morally problematic both because they blatantly disrespect the 300 women sampled, and also because they are overly charitable to the sexist company in the case.

Suppose that Victor makes his seemingly random selection of 300 women, but is yet to discover that none of them are in management. Suppose also that moral encroachment is correct, and learning these 300 not to be in management will leave him still with a fairly high expectation of how many of the 1000 women are in management. We judge this to be extremely disrespectful of the 300 women selected in the absence of traditionally recognized evidence that would support this disposition. Victor would, in effect, be thinking "these 300 women are probably not representative of the women in the company; given that I selected them, finding that *they* weren't in management does little damage to my overall expectation of how many women are in management in this company". For Victor to behave in this way despite no traditionally recognized evidence indicating he should do so would be deeply offensive to these women, insofar as they care about Victor's opinion.

In addition to disrespecting those 300 women, the credences recommended by moral encroachment are also problematic on account of being vastly too charitable to the sexist company in Predictive Inference Workplace. Suppose that Victor knows that it is the CEO, Jim, that chooses who gets to be in management, and suppose that in fact the reason why there are vastly more men than women in management in this company is that Jim is extremely sexist, but Victor has no reason to suspect this before visiting the company. Now, if after selecting the 300 women and finding none are in management, traditional epistemology would suggest, correctly, that this is very good evidence that Jim is sexist. The opposite holds if we adhere to moral encroachment. Since Victor's credence in a high number of woman in management is to stay high, he will *not* receive very strong evidence that Jim is sexist. Rather, he will think "those women I selected were probably duds anyway. I don't think Jim is sexist because I'm still optimistic that there are many women in management among those I didn't select." To recommend such reasoning is reprehensible, so moral encroachment should be rejected on moral grounds as well as epistemological ones.

6.5 The Effect of Learning About Oppresssion

Proponents of moral encroachment generally take a traditional approach to episte-mology on questions which are not morally fraught. They identify a great epistemic difference between drawing conclusions about the color of the next ball drawn from the urn, where traditionally recognized factors dominate, and drawing conclusions about whether a given woman is in management, where moral encroachment plays an important role. Therefore, a well-adjusted agent can be expected, given moral encroachment, to have a different doxastic response to statistical information de-pending on whether they understand the topic under consideration to be morally important or not.

Imagine that in addition to Victor, there are some aliens observing Victor and also forming opinions about how many women in the company are in management. These aliens, we suppose, know almost nothing of Earth—Victor just happens to be the first thing they can see looking through their telescope. Later, they will learn basic facts about humanity, including that womanhood is a socially important category, that being in management is a socially desirable position, and that sexism and other factors have lead to systematic barriers for women to reach such positions. But as they look through the telescope at Victor, they can't tell that one's gender matters more than what side of the street one lives on, or that being in management is any more socially desirable than being on probation for poor behaviour.

We submit that after they aliens learn these basic facts, they will be in a sim-ilar epistemic situation as Victor. Thus, if Victor should have a high expectation of the number of women in management, then so should the aliens. However, we also submit that before the aliens learn these basic facts, if they are well-adjusted, traditionally recognized factors will dominate and they will have a very low ex-pectation of the number of women in management, after seeing Victor select 300 of the 1000 seemingly at random and finding none in management. Therefore, if the aliens are well-adjusted, they will initially have a very low expectation of the number of women in management, but after learning about sexism and systematic barriers to women being in management, their expectation of the number of women in management shoots up.

This is exactly the wrong result. Learning about systematic barriers to women being in management should reinforce one's credence that there are very few in management. So the recommendation of moral encroachment, that the aliens should end with a high expected number of women in management, is incorrect.

7 Conclusion

Since there is no moral encroachment, Gendler's tragedy is forced upon us. There will be situations where the epistemically most appropriate doxastic state is not one that a well-adjusted person would like to have.

There arises no difficult conflict between the demands of epistemology and those of morality. Moral considerations trump epistemic ones. If morality requires you to do something, then do it! It is not the slightest bit surprising that morality and epistemology would pull apart in some cases, since their goals are different. Epistemology evaluates beliefs in terms of parameters like accuaracy and knowledge, whereas the moral evaluation of a belief should, at the very least, take into account broader consequences of having that belief, such as whether having it will detonate a nuclear bomb, killing millions.

Insofar as adopting a certain doxastic state is morally wrong, we recommend against adopting it. It does not follow, however, that in the motivating cases morality requires the agents of the cases to adopt different doxastic states than they do. Gendler's tragedy is that epistemic demands pull in favour of adopting beliefs or other attitudes that are in some respect unfortunate or regrettable. Just because an action is unfortunate or regrettable in some respect does not make it bad overall, and in many cases it is highly plausible that overall morality weighs in favour of adopting the doxastic state. This may include the motivating cases in the introduction, depending on how the cases are fleshed out. After all, there is something extremely morally important about being an epistemically well-adjusted agent, because an epistemically ill-adjusted agent will tend to make unwise decisions even if their priorities are impeccable. Moreover, we have repeatedly seen that insofar as an agent adheres to probabilistic credences, any attempt at charity turns into morally problematic partiality. Thus, a moral advance on the epistemically most appropriate credences will require non-probabilistic credences. Of course it is possible for the adoption of non-probabilistic credences to be morally correct: if the evil psychologist threatens to detonate the bomb unless you make your credences non-probabilistic, then you had better do it. However, it is usually very costly to deliberately make your credences non-probabilistic, since doing so would usually require becoming a thoroughly irrational person who is entirely unsuited to achieving their goals, no matter how laudable they may be. It is in general difficult to know what to do, epistemically or morally. Running these distinct modes of evaluation together does not help.[19]

References

Anderson, Charity and John Hawthorne (2019). "Knowledge, Practical Adequacy, and Stakes". In: *Oxford Studies in Epistemology* 6.

Basu, Rima (2018). "Beliefs That Wrong". PhD thesis. University of Southern California.

— (2019). "Radical Moral Encroachment: The Moral Stakes of Racist Beliefs". In: *Philosophical Issues* 29.1, pp. 9–23.

Basu, Rima and Mark Schroeder (2019). "Doxastic Wronging". In: *Pragmatic Encroachment in Epistemology*. Ed. by Brian Kim and Matthew McGrath. Routledge, pp. 181–205.

Bird, Alexander (2007). "Justified Judging". In: *Philosophy and Phenomenological Research* 74.1, pp. 81–110.

Bolinger, Renée Jorgensen (2020). "Varieties of Moral Encroachment". In: *Philosophical Perspectives* 34.1, pp. 5–26.

Cohen, Stewart (2019). "Pragmatic Encroachment and Having Reasons". In: *Pragmatic Encroachment in Epistemology*. Ed. by Brian Kim and Matthew McGrath. Routledge.

Fantl, Jeremy and Matthew McGrath (2002). "Evidence, Pragmatics, and Justification". In: *Philosophical Review* 111.1, pp. 67–94.

Gardiner, Georgi (2018). "Evidentialism and Moral Encroachment". In: *Believing in Accordance with the Evidence*. Ed. by McCain Kevin. Springer Verlag.

Gendler, Tamar Szabó (2011). "On the Epistemic Costs of Implicit Bias". In: *Philosophical Studies* 156.1, pp. 33–63.

Hawthorne, John (2004). *Knowledge and Lotteries*. New York: Oxford University Press.

Hawthorne, John, Daniel Rothschild, and Levi Spectre (2016). "Belief is Weak". In: *Philosophical Studies* 173.5, pp. 1393–1404.

Hawthorne, John and Jason Stanley (2008). "Knowledge and Action". In: *Journal of Philosophy* 105.10, pp. 571–590.

Jackson, Elizabeth and James Fritz (2021). "Belief, Credence, and Moral Encroachment". In: *Synthese* 199.1-2, pp. 1387–1408.

Johnson King, Zoe and Boris Babic (2020). "Moral Obligation and Epistemic Risk". In: *Oxford Studies in Normative Ethics* 10, pp. 81–105.

Kraft, Charles H., John W. Pratt, and A. Seidenberg (1959). "Intuitive Probability on Finite Sets". In: *The Annals of Mathematical Statistics* 30.2, pp. 408–419.

Lasonen-Aarnio, Maria (forthcoming). "Perspectives and Good Dispositions". In: *Philosophy and Phenomenological Research*.

Lewis, David K. (1980). "A Subjectivist's Guide to Objective Chance". In: *Studies in Inductive Logic and Probability, Volume II*. Ed. by Richard C. Jeffrey. University of California Press, pp. 263–293.

Moss, Sarah (2018a). "Moral Encroachment". In: *Proceedings of the Aristotelian Society* 118.2, pp. 177–205.

— (2018b). *Probabilistic Knowledge*. Oxford University Press.

Ramsey, F. P. (1926). "Truth and Probability". In: *The Foundations of Mathematics and other Logical Essays*. Ed. by R. B. Braithwaite. McMaster University Archive for the History of Economic Thought. Chap. 7, pp. 156–198.

Ross, Jacob and Mark Schroeder (2014). "Belief, Credence, and Pragmatic Encroachment". In: *Philosophy and Phenomenological Research* 88.2, pp. 259–288.

Savage, Leonard J. (1956). "The Foundations of Statistics". In: *Philosophy of Science* 23.2, pp. 166–166.

Schroeder, Mark (2018). "When Beliefs Wrong". In: *Philosophical Topics* 46.1, pp. 115–127.

Stanley, Jason (2005). *Knowledge and Practical Interests*. New York: Oxford University Press.

U.S. Bureau of Labor Statistics (2023). *Women in the labor force: a databook*.

Williamson, Timothy (forthcoming). "Justifications, Excuses, and Sceptical Scenarios". In: *The New Evil Demon*. Ed. by Fabian Dorsch and Julien Dutant. Oxford University Press.