

Prior Adjustments and Moral Encroachment on Credences

Yumin Liu
Department of Philosophy
The University of Hong Kong

Abstract

This paper develops a more fine-grained account of how moral considerations can shape credences. Rather than merely defending existing formulations, the paper offers a framework in which moral influences on credence operate through priors and in which prior adoption is constrained by both epistemic and moral considerations. Against the recent objection that moral encroachment conflicts with evidentialism and probabilism, I argue that agents can and should adjust their priors in relevant contexts so that their credences can jointly satisfy both moral and epistemic requirements.

Keywords: Moral encroachment on credences; prior adjustments; moral norms; context-sensitivity; evidentialism; probabilism

1 Introduction

Moral encroachment is the thesis that the epistemic status of a doxastic state can depend on its moral features. Philosophical discussions in this area have largely focused on outright belief. In particular, many philosophers advocate a radical version of moral encroachment on beliefs,¹ holding that a belief can itself be morally wrong — and thereby epistemically unjustified — rather than merely wrong in virtue of its causes or effects (Basu 2018, 2019; Marušić and White 2018; Schroeder 2018; Basu and Schroeder 2019).

If this idea is taken seriously, it is natural to extend it to credences (see Jackson and Fritz 2021; Goodsell and Hawthorne 2025).² However, this extension is widely thought to conflict with central epistemic norms — namely, evidentialism and probabilism, and thus seems hard to sustain (e.g. Pace 2010; Ross and Schroeder 2014;

¹Hereafter, I use “radical moral encroachment” and “moral encroachment” interchangeably to refer to the general radical view.

²Hereafter, I use “moral encroachment on credences” to refer to this view.

Schroeder 2018; Bolinger 2020; Gardiner 2018; Munton 2019; Fritz 2020; Jackson and Fritz 2021; Goodsell and Hawthorne 2025). While existing accounts identify important moral concerns, they leave open how such considerations should be incorporated into credence formation in a principled and systematic way.

This paper addresses that gap. Rather than focusing on whether moral encroachment holds in its existing formulations, it develops a more fine-grained account of how moral considerations can shape credences. This approach preserves the core idea that moral factors can affect epistemic rationality, while offering a framework in which credences can remain compatible with evidentialism and probabilism. In particular, it provides a descriptive account of how boundedly rational agents, operating with coarse evidence and limited reasoning resources, can form credences that are both morally and epistemically permissible.

The account focuses on a fundamental component of credence formation — priors. Specifically, I propose a *prior adjustment approach*, according to which agents can rationally adjust their priors to satisfy both epistemic and moral norms, thereby ensuring that their resulting credences are epistemically and morally permissible.

The key point is that our priors can — and sometimes should — be rationally adjusted to meet context-sensitive requirements on credences. As Ye (2024) argues, priors can be rationally adjusted without evidential change if doing so promotes long-term accuracy. Woodard (2022, p. 326, 332) also holds that we are epistemically entitled to change our minds whenever doing so is epistemically improving. This paper develops a distinct account, extending these ideas to show that priors can also be rationally adjusted without changes in evidence when doing so helps to uphold both epistemic and moral integrity.

Section 2 examines evidentialist and probabilist objections, arguing that two assumptions — the epistemic requirement that credences be proportionate to statistical base rates and the moral requirement of impartiality — can be replaced by more fundamental norms: (1) updating credences through conditionalization and (2) avoiding harming others.

The central contribution of this paper is showing how moral constraints can operate through priors without violating fundamental epistemic norms. A key insight is the distinction between two moral norms governing credences — Impartiality and Avoid Harming — and the argument that the latter can be reconciled with epistemic norms by identifying the right fundamental constraints and adjusting priors accordingly.

Section 3 develops the prior adjustment approach within the hypothetical prior framework. I argue that by integrating epistemic and moral norms into agents' evidential standards, agents can adopt priors that yield both morally and epistemically permissible credences in morally sensitive contexts. I also address potential objections and compare this approach with Johnson King and Babic (2020), arguing that it has distinct advantages. Section 4 addresses evidentialist and probabilist objections. Section 5 concludes.

2 Evidentialist and probabilist objections

2.1 Evidentialism

Evidentialist objections to moral encroachment on credences hold that rational credences must fit the agent’s evidence:

Evidentialism (Credence) An agent’s credence is epistemically rational if and only if it *fits* her evidence (e.g. Pace 2010; Jackson and Fritz 2021; cf. Conee and Feldman 2004).

How a credence “fits” evidence is an underspecified requirement. Many would agree, however, that when an agent receives information her priors treat as probabilistically relevant, her credence should move accordingly. But how to move requires further specification. This basic idea can be combined with different assumptions, for example, that priors should remain stable, that credences should be updated by conditionalization, or that only a single set of priors can be rational.

Many such assumptions, especially those concerning priors, are not directly grounded in the evidence and therefore seem not to be genuine requirements of evidentialism. Nevertheless, it is widely accepted that rational credences should be formed by conditionalization, since this directly captures how credences ought to respond to evidence³:

Conditionalization A rational credence should be formed via conditionalization:

For any time t_i and later time t_j , suppose an agent receives evidence E between t_i and t_j , and her credence at t_i in E satisfies $cr_i(E) > 0$, then for any proposition P ,

$$cr_j(P) = cr_i(P|E)$$

Critics argue that moral encroachment violates evidentialism by requiring departures from evidence. In particular, they appeal to a simple principle linking credences to frequencies:

Credence-Statistics Link (CS Link) If an agent knows that σ belongs to a reference class G in which $x\%$ have feature F , then she should assign credence $x\%$ to the proposition that σ is F .⁴

Granting this principle, moral encroachment appears to conflict with evidentialism, since it sometimes requires credences to diverge from statistical frequencies (Fritz and Jackson 2021). Consider the following example (adapted from Goodsell and Hawthorne 2025):

Workplace Victor learns that in a company, only 1% of female employees (10/1,000) are in management. While at the company, he overhears a woman named Joan speaking and assigns a 0.01 credence to her being in management.

According to the CS Link, this credence is rational. However, moral encroachment theorists argue that credences based solely on statistical profiling can be morally

³Which further assumptions need to be combined with conditionalization will be discussed in the remainder of the paper.

⁴CS Link is often understood as a form of direct inference and is widely used in objections to moral encroachment (e.g. Moss 2018; Munton 2019; Johnson King and Babic 2020). More sophisticated statistical approaches have been proposed (e.g. Babic et al. 2024), but this paper focuses on the simple form commonly assumed in the literature.

problematic; if too low, they may contribute to harm and warrant upward revision (e.g. Basu 2018, 2019; Marušić and White 2018; Schroeder 2018). Avoiding this may therefore require assigning a higher credence, departing from statistical evidence.

This objection is often motivated by the CS Link, which reflects the intuitive idea that rational credence should track statistical evidence. However, I argue in Section 3.2 that this assumption is not required.

2.2 Probabilism

Probabilist objections focus on the probabilistic coherence requirements of rational credences, which are typically based on the following norm of probabilism and a statistical fact directly derived from it:

Probabilism An epistemically rational credence can be represented by a probability function that satisfies the probability axioms.

Fact If a credence function cr assigns 1 to the proposition that exactly n of $x_1 \dots x_k$ (with $x_i \neq x_j$ for $i \neq j$) are F , then *the average value* of $cr(x_i \text{ is } F) = \frac{n}{k}$ (Goodsell and Hawthorne 2025, Sect. 6.1).

Critics argue that moral encroachment on credences violates probabilism by allowing rational credences to deviate from those justified by probabilism and basic statistics. For instance, Goodsell and Hawthorne (2025) argue that such deviations are impermissible, as they have to either conflict with probabilism or violate the moral requirement of impartiality, which they characterize — summarized in my own words — as follows:

Impartiality A moral agent should not show excessive charity to an individual when she has no special evidence about that individual, especially when doing so comes at the partial expense of other individuals for whom she has exactly the same evidence.

As they explain, ‘Given Probabilism, it is simply impossible for epistemically appropriate credences to be more charitable overall’; otherwise, it would require partiality — that is, assigning a higher- or lower-than-traditionally-recommended credence while maintaining probabilism requires adjusting other credences in the opposite direction. Yet, ‘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’ (2025, Sect. 6.1).⁵

This can be illustrated in the Workplace example. Consider again Victor’s 0.01 credence. Suppose this credence is morally wrong and therefore irrational. In that case, moral encroachment on credences would require him to raise it — say, to 0.5 (or to some value above 0.01).

Since Victor has no special evidence about any individual woman and relies solely on the statistical evidence — namely, that 1% of women in the company are managers, then, according to Impartiality, he should assign the same, impartial credence to each woman. However, if Victor wants to meet both the requirement of moral encroachment and Impartiality by assigning a uniform, higher-than-the-base-rate credence — such

⁵Indeed, Probabilism, the Fact, along with Impartiality together imply the CS Link, according to which credences in direct-inference propositions are rationally required to equal the statistical base rate. Moral encroachment on credences deems such credences irrational when they are morally wrong, thereby violating both evidentialism and probabilism.

as 0.5 — to each woman, the resulting credence distribution will be probabilistically incoherent — the average credence would be 0.5, vastly exceeding 0.01.

By contrast, if Victor wants to satisfy the requirement of moral encroachment while preserving probabilism, increasing his credence about Joan would require lowering the credences assigned to the remaining 999 women in order to keep the overall average at 0.01. This adjustment would lower the average credence for the remaining women to approximately 0.0095. If he also assigns a credence of 0.5 to another woman employee being in management, the average for the remaining 998 women drops further to around 0.009.

This pattern continues: each time Victor raises the credence for one individual, the average for the rest decreases. As a result, no matter how he adjusts his credences, at least one will always fall below the original average of 0.01, which creates a distribution partial to some individuals. Therefore, it seems that if Victor adheres strictly to probabilism, his credences will violate Impartiality and thus can never be fully morally acceptable.

However compelling, this objection is not decisive either. It seems that if we can rationally reject Impartiality but uphold probabilism, Victor may escape the charge. However, this move introduces a new problem: at least one of Victor's credences will still fall below the original average. This seems morally problematic — if the credence of 0.01 is already deemed morally wrong (and thus irrational) for being too low, assigning even lower credences to others only amplifies the moral wrongness. How, then, could such credences be morally permissible?

The answer lies in identifying which moral norm we should adopt and how it can impose context-sensitive requirements on credence formation. In the next section, I argue that Impartiality should be set aside in favor of the norm of Avoid Harming, which is implicitly endorsed by moral encroachers and allows Victor's credences to be morally permissible without violating probabilism.

3 The prior adjustment approach

The aim of this section is to develop a novel account of moral encroachment — the prior adjustment approach. Rather than treating priors as fixed, the core idea is that priors can directly encode an agent's epistemic and moral values, shaping how evidence is interpreted and credences are formed. On this view, agents can, and in some cases should, adopt different priors across contexts, even in the absence of new evidence.

Prior adjustment is not unconstrained. Any adjustment must remain responsive to the agent's total evidence. Properly understood, prior adjustment provides a systematic way of forming credences that satisfy both moral and epistemic requirements. This section clarifies how these constraints operate.

3.1 Context-sensitive priors within conditionalization

Priors are a central and contested notion in Bayesian epistemology. Some take them to be any pre-update credences, while others understand them as credences prior to

all evidence (so-called *ur-priors* or *hypothetical priors*). This paper adopts the latter interpretation.⁶

On this view, credence formation is governed by conditionalization (Williamson 2000; Hedden 2015; Meacham 2016; Titelbaum 2022):

Hypothetical Prior Conditionalization At every time t_j , given an agent’s total evidence E_j at t_j , her credence cr_j in a proposition P at t_j should be determined by a rational prior pr satisfying probability axioms and the regularity principle such that

$$cr_j(P) = pr(P|E_j)$$

This is a synchronic constraint⁷: it determines rational credences at a time but does not fix priors uniquely or impose diachronic restrictions. As a result, it leaves room for different rational priors, and thus different credences, at different times or contexts, given the same evidence.⁸

The key question, then, is which priors are permissible. At a minimum, priors should satisfy standard probabilistic constraints (e.g., probability axioms, regularity, and related principles).⁹ These constraints are synchronic: they govern which priors are admissible at a given time. They therefore do not rule out context-sensitive priors.

I argue that priors can also encode an agent’s epistemic and moral commitments, insofar as their adoption is governed by the epistemic and moral norms the agent endorses. Since these commitments may vary across contexts, priors — and hence credences — may also rationally vary even when the agent’s evidence remains fixed. This provides a basis for context-sensitive credence formation in morally sensitive cases.¹⁰

But why should moral considerations matter in credence formation? And how can priors encode epistemic and moral values?

The first question is not the primary focus of this paper, though I briefly address it in the next subsection. The central idea is that *the risk of harm* matters: such risk can itself be morally wrong and may generate a requirement to avoid it.

⁶Throughout, when I refer to “priors”, I mean hypothetical priors.

⁷On other interpretations of priors, conditionalization may take a different form. For example, credences may be required to be updated by conditionalizing previously held credences on new evidence, rather than on the agent’s total evidence. This is a diachronic requirement.

⁸This paper assumes precise credences. Although multiple credence functions may be rationally permissible for an agent at a given time, she should adopt only one at any given time. Developing an imprecise-credence approach to moral encroachment is an important and promising direction for future work, but it lies beyond the scope of this paper.

⁹Beyond these requirements, philosophers have proposed many further constraints. Some hold that rational priors must determine a unique credence function for any given body of evidence (e.g., White 2005), while others allow that priors may be permissive and subjective, so that more than one credence function can be rationally supported by the same evidence (e.g., Kelly 2013).

¹⁰This view resonates with intrapersonal permissivism, a theory that an agent can rationally hold different doxastic attitudes based on the same evidence (e.g. White 2005, Kelly 2013, Schoenfeld 2014). Intrapersonal permissivism is controversial. Critics, such as White (2005), argue that it entails epistemic arbitrariness or inconsistency, but Kelly (2013) suggests a conceptual bridge: if interpersonal permissivism holds, intrapersonal permissivism may follow. Some time-slice epistemologists (e.g. Moss 2014, 2015; Hedden 2015) also claim that an agent can rationally hold different doxastic attitudes based on the same evidence at different times.

Like moral encroachment, discussions of permissivism on credences can face evidentialist and probabilist challenges. However, within the hypothetical prior framework, permissivism becomes more natural: both intra- and interpersonal differences in credences based on the same evidence can be rational, provided rational underlying priors.

Many philosophers have argued that moral considerations can affect the rationality of beliefs, and potentially credences. A common motivation is that allowing rational beliefs to be morally wrong creates tension between epistemic rationality and moral permissibility (e.g., Gendler 2011; Basu 2018).¹¹

My primary focus is on the second question. The idea that priors can encode agents' values is not unusual. For example, Robert (2007) highlights a duality between priors and loss functions, suggesting that value judgments may be incorporated into the prior.

However, I do not pursue this duality or appeal to sophisticated loss functions here. Instead, I argue that priors themselves can be shaped by agents' values in a more direct and intuitive way: by adjusting them in response to normative requirements on rational credence.

This proposal is supported by interpreting priors as encoding an agent's constant *evidential standards* throughout her life: abstract, relatively stable dispositions for interpreting and responding to evidence (Meacham 2016; Titelbaum 2022). Evidential standards include the agent's deep-seated *normative commitments* in interpreting evidence and forming credences. Such commitments may incorporate 'pragmatic influences, a predilection for hypotheses with certain features, a tendency toward mistrust or skepticism' (Titelbaum 2022: 130; Lewis 1980).¹²

For example, a risk-averse agent may be more hesitant to assign a high credence that it won't rain tomorrow than a risk-seeker, even given the same evidence. A skeptic might be more inclined to believe we are all Brains in Vats than an ordinary person given their same everyday experience.

Similarly, a morally conscientious agent may be less willing to form a high credence that a woman is not in management based solely on demographic statistics, compared to a biased agent. These differences in credences need not arise solely from differences in evidence, but can also reflect differences in how agents interpret evidence and form credences, as governed by the commitments encoded in their priors.

Building on this idea, I argue that moral considerations can influence credences by shaping which priors are permissible. The central claim is that priors can be rationally adjusted in response to the relevant normative requirements even when the evidence does not change, since these requirements on credences can differ across contexts. For example, when a credence risks being morally wrong, avoiding moral wrongdoing may require assigning a lower credence to certain propositions than would be required in contexts where no such moral risk is present. This can be explained by the context-sensitive requirements on priors: in morally sensitive contexts, priors are governed by

¹¹Many philosophers have argued that this conflict can be avoided. For example, Johnson King and Babic (2020) and Babic et al. (2024) argue that a sufficiently sophisticated Bayesian can resolve the tension by assessing epistemic risk, specifying an appropriate probability space, and analyzing data. A full defense of this objection lies beyond the scope of this paper. However, as I suggest later, these arguments may be open to challenge.

¹²Philosophers disagree about the nature of hypothetical priors (ur-priors), offering several interpretations: as agents' initial credences (Skyrms 1983); as appropriate probability functions (Meacham 2007; Titelbaum 2013); or as evidential or ultimate epistemic standards, which I talk about here. Each interpretation has its trade-offs, but Meacham (2016) and Titelbaum (2022) argue that the evidential-standards interpretation is the most promising overall. Titelbaum (2022) explicitly uses the term "hypothetical priors" to refer to this interpretation. While my argument assumes the evidential-standards view, it might also be adapted to other interpretations, as long as the hypothetical prior conditionalization holds — for example, under the functional view, moral considerations may influence prior functions as weighting coefficients.

both moral and epistemic norms, whereas in morally insensitive contexts they need only satisfy epistemic norms. In this way, moral encroachment on credences can be sustained.¹³

The next subsection develops this proposal by explaining how normative constraints determine permissible priors and shape the resulting credences. I argue that objections to moral encroachment often rely on norms — such as the CS Link and Impartiality — that should not be treated as fundamental requirements on credence formation. Instead, I clarify which norms should govern priors.

3.2 Priors and norms

The aim of the prior adjustment approach is to provide a more fine-grained and plausible account of moral encroachment on credences that remains compatible with evidentialism and probabilism. The tensions between moral and epistemic norms may suggest a deeper incompatibility. For those committed to maintaining both moral and epistemic integrity, such a conflict is unacceptable. This section argues that moral encroachment in fact calls for a *morally cautious attitude*, which is compatible with fundamental epistemic norms.

To begin, we must determine which norms a moral encroacher should adopt.

Philosophers disagree on many norms, but fortunately, we have identified several plausible candidates. For example, many would agree that Evidentialism (Credences), Conditionalization, and Probabilism should be accepted.¹⁴ However, the CS Link and Impartiality should not. Instead, we should adopt a moral norm focusing on avoiding harming others. In fact, without further constraints, the CS Link can lead to epistemic irrationality, while Impartiality can lead to harm when uniform credences are morally wrong. Let's explain these points one by one.

First, the CS Link, at least in its simple form, should not be applied unreflectively. It assumes that statistical data can be used directly without resolving familiar issues such as reference class selection or data quality (e.g., Buchak 2014; Babic et al. 2024), and thus need not be accepted as a fundamental norm even by orthodox epistemologists. What matters instead is whether morally constrained credences can satisfy more basic norms such as evidentialism and conditionalization.

Still, the CS Link highlights a genuine problem: agents often rely on coarse and underspecified statistical information, and unreflective use of such data can lead to epistemic and moral error. One response to this problem is to refine statistical reasoning. Another, developed here, is to explain how agents can form morally and epistemically permissible credences even under imperfect evidential conditions

¹³Some might worry that allowing prior adjustment conflicts with accuracy-first defenses of conditionalization (e.g., Greaves and Wallace 2006). However, the proposal here is not meant to reject conditionalization as a rational updating rule. The idea is instead that moral considerations may constrain which priors are admissible prior to updating before updating occurs. Once a prior is fixed, credences can still be updated through standard hypothetical prior conditionalization. In this respect, prior adjustment is intended to supplement, rather than replace, conditionalization-based approaches. A fuller treatment of its relation to accuracy-theoretic frameworks lies beyond the scope of this paper.

¹⁴There are many other epistemic norms — such as the Principal Principle, reliability, sharp credences, and accuracy maximization — that may also bear on the prior adjustment approach. This paper focuses only on evidentialism and probabilism, though some of these further norms may ultimately prove compatible with the view. A fuller treatment lies beyond the scope of this paper.

Second, I argue that radical moral encroachment should reject Impartiality, as it can conflict with a more fundamental moral norm: to avoid harming others.

One might worry that rejecting Impartiality would lead to injustice or unfairness. This need not be the case, however, if we instead adopt a more general norm that constrains credence formation when injustice or unfairness would risk harming others. While the prior adjustment approach rejects Impartiality, it endorses a more fundamental norm — the moral obligation to avoid harming others — which still requires treating others justly and fairly insofar as such treatment is necessary to prevent harm.

This norm serves as a safeguard, ensuring that agents who follow it do not form morally wrong credences:

Avoid Harming A moral agent should avoid harming others, so long as this does not conflict with other moral obligations.

Several clarifications are in order regarding this norm. This norm may be contentious, as it touches on broader debates about moral responsibility — for example, in cases involving self-defense or the trolley problem. Such discussions often draw a further distinction between actively bringing about harm that should not have occurred and passively failing to prevent harm (e.g. Thomson 1976; Harman 2009).

Both kinds of harm may arise in moral encroachment cases. When an agent forms a new credence in a morally sensitive context, she faces the risk of actively bringing about harm, in the sense that she should not actively form morally wrong credences in such contexts. By contrast, when an agent already holds a credence but moves from a morally insensitive context to a morally sensitive one, she faces a problem of preventing harm: her existing credence may cause harm unless she adjusts her priors accordingly. The Avoid Harming norm can therefore apply to both kinds of harm, so long as avoiding them does not conflict with other moral obligations.

A useful characterization of the relevant harm, developed by philosophers such as Basu and Schroeder (2018), Basu (2019), Marušić and White (2018), and Schroeder (2018), is that it involves four features: (1) it is directed at specific individuals (including collective entities, such as social groups);¹⁵ (2) it expresses disrespect for an agent's personal agency and contributions; (3) were the agent to become aware of the credence held toward her, she would reasonably feel harmed by it; and importantly (4) such harm can arise from holding a belief or credence itself, rather than from its upstream causes or downstream effects.

Thus, the harm at issue is not only harm that is deemed to occur if a certain credence is formed, but also the risk of such harm. When an agent is uncertain whether such harm will in fact occur, but knows that the credence would be harmful if it were formed, she should also avoid taking on that risk.

But there are two worries about these features. First, one might worry that this account trivializes harm and implies that many ordinarily permissible activities should be avoided merely because they carry some risk of harm. This need not be the case, since the requirement to avoid harm applies only when doing so does not conflict with

¹⁵Moral encroachers often do not specify whether the individuals who may be harmed include collective entities, such as social groups. Here, I take such groups to be included, since attitudes or credences directed at a group can plausibly harm that group as a whole. For example, forming credences that diminish the perceived working capacity of women can be harmful to women as a group, even when no particular individual is singled out.

other moral obligations. In cases such as credence formation, where no such conflict arises, it is often not problematic to take a better available option that avoids the risk of harm.

Moreover, there is rarely any “deemed” or deterministic harm whose occurrence can be predicted with certainty. Even in cases where harm is widely thought to be morally salient and ought to be avoided, there is typically a nontrivial probability that the harm will not in fact occur.

For example, when driving near an elementary school, a driver should slow down even though no child may actually step into the road. The harm may never occur, but slowing down is still morally appropriate, and often required, because it reduces the risk of serious harm at little cost.

Second, a natural concern is that merely holding a credence without any action that reflects a biased attitude does not harm others. If an agent holds a biased credence about someone but nonetheless acts fairly, or even favorably, toward that person, many would question whether the credence itself has moral significance. How, in such cases, could the credence be harmful?

This question can be explained in several ways, since the considerations that make belief itself morally wrong can also explain why credences may be morally wrong and harmful. For example, Basu (2019) argues that belief can itself demean, socially subordinate, or fail to respect individuals appropriately. The same point applies to credences. Credences formed on the basis of coarse-grained statistical generalizations can inherit the social biases embedded in those statistics, thereby socially subordinating individuals and failing to respect their agency. Such credences also risk treating individuals as mere interchangeable objects of assessment rather than as particular persons. Thus, holding such credences can itself already be morally objectionable and harmful.

Moreover, credences can be morally wrong because of the risk of harm they impose. Bolinger (2020: 2423–24) argues that imposing a risk of harm can itself be morally wrong: “So long as A has a moral claim against S imposing H, S wrongs A if she imposes an unjustified risk of H on A.” Although this discussion concerns harms arising from actions based on belief, rather than harms caused merely by holding the belief, the same reasoning extends naturally once we acknowledge that beliefs and credences can themselves impose risks of harm. In such cases, imposing a risk of harm through one’s credences may itself constitute a moral wrong.

Indeed, credences can carry a risk of harm, in the sense that, were the person concerned to become aware of them, she could reasonably be harmed. Even if an agent acts fairly toward someone while privately holding a biased credence about her, the credence itself can pose a morally salient risk — indeed, in some cases a more serious one than overtly biased action. One may reasonably assume that fair treatment reflects an unbiased attitude, so discovering the underlying credence can therefore feel like a betrayal, causing harm by undermining trust, inducing self-doubt, and damaging future interpersonal relations.

Sometimes such credences may not directly cause harm, and sometimes they will. Some individuals may not feel harmed by discovering them, while others reasonably will. Given this uncertainty and the morally salient risk involved, it is reasonable, and

in morally sensitive contexts necessary, to avoid the risk of harm — that is, avoid forming such credences in the first place. Thus, even merely holding a credence can be subject to the Avoid Harming norm. This norm provides additional support for moral encroachment — one that Impartiality alone cannot capture.

Indeed, I suggest that *avoiding risk of harm* exemplifies the kind of *moral caution* radical moral encroachers implicitly endorsed and seek to promote in social life. As they emphasize, morality can be demanding: it requires agents to exercise heightened caution in upholding their moral responsibilities, even when full compliance is difficult or rarely achievable (e.g., Basu 2019).

However, one might still worry that reducing moral encroachment to a moral norm such as Avoid Harming is overly demanding. If the norm is objective moral obligation, then it could seem to imply that many agents are constantly morally blameworthy, as many people are not in a position to satisfy this norm. For example, individuals living in environments shaped by entrenched social and cultural biases may lack the means to identify which credences will harm others and which contexts are morally sensitive. Such agents may genuinely strive to avoid harming others but, due to these constraints, may nonetheless fail to do so.

This worry can be addressed by distinguishing moral obligation from moral blameworthiness. Even if the Avoid Harming norm is violated in such cases, agents need not be morally blameworthy, since their failures may be excused given their epistemic and social limitations. Moral encroachment is therefore not overly demanding: it imposes substantive moral constraints on credence formation while allowing for excuse in cases where agents are not reasonably positioned to comply.¹⁶

Moreover, even in cases where agents are not excused yet are not well positioned to satisfy the norm, this does not undermine the norm itself. Rather, it reflects a central idea behind radical moral encroachment: morality is demanding. It requires agents to exercise heightened moral caution in discharging their responsibilities, even when full compliance is difficult or rarely achievable (Basu 2019).

With these clarifications in place, we can now see how this norm bears on credence formation. In this context, it is natural to hold that if forming a credence would harm others, the Avoid Harming norm requires agents to refrain from forming it.

This aligns with the main defense of radical moral encroachment: credences can harm others by making them feel harmed when discovered, diminishing their personal agency and contributions, or treating them as objects rather than moral agents. Such credences are morally impermissible and, on the radical encroachment view, epistemically irrational. In this sense, rational credences are constrained by the requirement to avoid harming others.

By contrast, credences formed in strict accordance with Impartiality are not necessarily rational in the moral-encroachment sense, since they can sometimes harm others, especially when they are based solely on coarse-grained statistical data. In the Workplace case, for example, Victor assigns a credence of 0.01 to Joan's being in management based solely on gender statistics. Even if this credence is impartial to every

¹⁶Evidentialist and probabilist norms face a parallel worry. Agents who lack training in epistemology or probability theory may likewise be poorly positioned to fully satisfy these norms. But this does not undermine their normative authority: they remain standards that agents should aspire to meet and by which their epistemic states can be appropriately evaluated.

woman in the company given his total evidence, radical moral encroachment theorists hold that it still wrongs Joan and is therefore morally impermissible.

On the other hand, partial credences do not automatically give rise to morally salient harm. Whether a credence is morally wrong depends not only on its distribution, but also on the context in which it is formed and on which propositions are morally salient — a point I will explain shortly.

Therefore, Impartiality alone fails to fully capture the core commitment of moral encroachment — refraining from forming credences that harm others. Instead, the norm Avoid Harming more accurately captures this commitment.

In sum, by integrating Avoid Harming — rather than Impartiality — into agents' evidential standards, their priors can be constrained to yield credences that do not risk harming others. Agents who fail to adopt such priors thereby fall short of their own commitments regarding how to process evidence and form credences, rendering their resulting credences morally impermissible and, on the moral-encroachment view, epistemically irrational.

As mentioned earlier, which credences risk harming others may vary across contexts. Accordingly, the priors that are permissible under epistemic and moral constraints are also context-sensitive. In this sense, moral encroachment is context-dependent. The next section explains how and why this is so.

3.3 Context-dependent prior adjustments

The core idea of the prior adjustment approach is that moral considerations influence credences by requiring agents to adopt only priors that, when conditionalized on the total evidence, avoid generating credences that risk harming others in a relevant context. How priors should be adjusted varies across contexts because moral requirements on credences vary by context.

In some situations, credences can cause no harm and no prior adjustment is required. While in others, the very same credences, formed by conditionalizing the same priors on the same total evidence, can involve a serious risk of harm, and thus call for prior adjustments. A credence cannot be evaluated as morally wrong in isolation; it must be assessed within its specific context.

This is because the moral status of a credence depends not merely on its value, but in fact on the moral stance it expresses. In different contexts, the same credence can carry different moral implications. For example, a credence equal to base rates based solely on coarse-grained statistical data can risk harming the subject when it is formed in a context that expresses a moral stance diminishing the subject's personal agency and individual contributions. However, in other contexts, such a credence may be morally defensible or even favorable, especially when it is used to reveal the structural biases and motivate social reform.

In the workplace example, a credence of 0.01 in the proposition that Joan is in management can be morally wrong when it reflects an assessment of Joan based solely on social profiling. However, the very same credence in “Joan (or any woman in the company) is in management” can be morally permissible if it is intended to highlight the underrepresentation of women in leadership roles within the company.

Thus, to avoid credences that risk harm requires avoiding credences that express morally harmful stances in relevant contexts. The problem can be solved if agents can rationally form credences that reflect morally favorable stances.

What counts as a morally favorable stance? Some might argue that it requires strategic caution, even at the cost of departing from what otherwise seems epistemically optimal. For example, agents may be inclined to assign relatively high credence to a person's being a manager, since mistaking a janitor for a manager causes little harm, whereas mistaking a manager for a janitor can seriously disrespect their agency.

Such a stance may be morally favorable, but it is not epistemically favorable and is therefore not permitted by the prior adjustment approach. Since an agent is committed to both moral and epistemic norms, her credences cannot satisfy only one of them when both are relevant and must be taken into account.

A genuinely morally favorable stance that can also be rational, and, when appropriate, required by the prior adjustment approach, is a morally *neutral* stance. To be morally neutral is to refrain from adopting a clear positive or negative assessment, especially when such assessments would risk harm. For example, when asked whether Joan is a manager in a context where the evidence is ambiguous and the risk of harm is salient, a morally neutral response would be: "I do not know," "I am not sure," or "I have no opinion." In the credal context, moral neutrality can likewise be represented by holding credence within a certain range that appropriately expresses such neutrality.¹⁷

Which credences count as morally neutral or favorable may also vary across contexts. For example, in the *Workplace* case, a morally neutral credence that Joan is in management needs to be higher than the base rate, whereas a morally neutral credence that Joan is *not* in management needs to be correspondingly lower than the base rate. Nevertheless, I argue that there can be a range of credences that represent a morally neutral stance in a given context. For simplicity and clarity, I assume a morally neutral credal range centered near 0.5 to illustrate how priors should be adjusted in this paper. Yet the approach can apply equally to other ranges of morally permissible credences.

That is to say, to avoid harming others by credences, agents should adjust their priors so that, when updated on their total evidence, they yield credences that can express a morally neutral stance while being rational in the given morally sensitive context.

I now offer a clarification of what counts as a morally sensitive context, a notion that has so far been left implicit in the paper despite being extensively relied upon. The prior adjustment approach applies only in such contexts and is not intended to govern all cases of credence formation. More specifically, morally sensitive contexts are those in which forming a credence involves a risk of harming others and is directed at particular individuals or identifiable groups.

On standard moral-encroachment views, morally sensitive contexts typically involve forming credences about particular human individuals (Basu and Schroeder 2018; Schroeder 2018; Basu 2019). This helps explain why the prior adjustment

¹⁷Some may also take such neutrality to amount to suspension of judgment (e.g., Friedman 2017). On this view, morally and epistemically rational credence in morally sensitive contexts can be understood as a credence that represents suspension.

approach is not meant to apply universally. In morally insensitive contexts — such as assigning a high credence to the proposition that your neighbor’s beagle barks equal the statistical base rate (e.g. 99% beagles bark) — forming such credences poses no plausible risk of harm. Moral encroachment therefore treats these credences as both epistemically and morally permissible.

By contrast, in morally sensitive contexts — such as assigning a high credence that a man is a criminal solely because statistical data suggest a higher crime rate in his community — moral encroachment holds that such credences risk harming others and are both morally wrong and epistemically irrational.

Moreover, the morally sensitive context also depends on the inquiry the agent is pursuing. By inquiry, I mean the agent’s *epistemic focus*: what she directs her attention toward and aims to find out. This can be represented either in terms of unobservable internal epistemic states, such as the agent’s mental activities,¹⁸ or in terms of outward behavioral patterns, such as asking questions or searching for evidence, through which her inquiry is manifested.

A further feature of morally sensitive contexts concerns what is morally salient within them. By morally salient, I mean an agent’s epistemic focus, whether understood as an internal state or as manifested in action, can make her doxastic attitude toward focus morally evaluable and responsible in the given context. The object of the epistemic focus determines which propositions and which individuals are *morally salient* in the given context. Here, I understand moral salience as context-relative moral evaluability.¹⁹ Different morally sensitive contexts can involve different morally salient issues.

It is important to distinguish moral salience from moral importance. By moral importance, I mean the agent’s standing responsibilities grounded in relationships, vulnerability, or structural injustice (e.g., responsibilities toward one’s children, dependents, or marginalized individuals). By contrast, moral salience concerns what is contextually prominent or calls for attention within a particular inquiry or decision.

As it stands, the view may suggest that mere salience triggers prior adjustment, which would be problematic. I agree that salience itself should not justify boosting credences. The normative basis of prior adjustment is moral importance, which is grounded in standing reasons and does not depend on what is currently salient. Rather, salience plays a role in determining relevance: it determines which morally important individuals become operative within a particular context of inquiry and decision-making.

Moral salience is especially important for boundedly rational agents. Salience helps structure attention and makes credence formation tractable. Without some mechanism

¹⁸One might further wonder whether such mental states must be under the agent’s control, or consciously accessible to her, in order for her doxastic states in the relevant context to be subject to moral evaluation. It is widely accepted that agents are morally responsible only for what lies within their control, so some form of control is required. What counts as control, however, admits of different interpretations. While conscious awareness plausibly suffices for moral evaluation, it is not necessary. There are cases in which an agent is unaware of her epistemic state yet still bears moral responsibility (for example, by owing an apology) although such unawareness may excuse or mitigate blame. Owing to space limitations, I do not pursue this issue further here.

¹⁹One might worry that this assumes epistemic focus determines moral focus, even though the two can in principle diverge. Here, my discussion is limited to moral-encroachment cases where moral focus should be guided by and rely on epistemic inquiry, and within this scope the assumption is unproblematic.

of salience, boundedly rational agents could not deliberate or form credences reasonably, since they cannot attend to all morally important individuals simultaneously. In this sense, moral importance tracks which individuals morally matter, while salience tracks which individuals currently call for attention within a given inquiry, credence formation, or decision-making.

Accordingly, although many individuals may be morally important, prior adjustment applies only when such individuals are salient to the inquiry. Salience does not determine who matters, but only when morally grounded weights become operative in reasoning.

This also explains why moral importance and salience can diverge. Even when I am not attending to my child, my child remains morally important, and it may still be appropriate for me to remain partial to them. However, this standing partiality need not affect my credences unless my child becomes salient to the inquiry at hand.

In sum, once a morally important individual is salient, priors may be adjusted; where no such moral importance is present, salience alone does not justify prior adjustment.

These ideas can be captured by allowing priors to reflect the moral importance of individuals. One way to represent this is by adjusting a baseline prior $P(H_i)$ such that it satisfies a morally weighted score. This score determines whether the prior is morally permissible and may be represented as $w(i)h_i(pr)$, where $w(i)$ denotes the moral weight of i and $h_i(pr)$ captures the risk of harm the prior poses to i in a given context. Both the weights and the associated harms may vary across agents and contexts, depending on their moral statuses and circumstances.²⁰

Against this background, we can now see how prior adjustment operates in morally sensitive inquiry. When an agent is inquiring whether a specific individual (or subgroup) has a certain trait within a broader population, moral salience fixes the relevant domain of evaluation — namely, the propositions that the individual has the trait or does not — while moral importance determines whether and how those propositions are weighted. In such cases, moral considerations influence credence formation by requiring priors that, when updated on the total evidence, yield morally neutral credences about the specific individual in question.

By contrast, when an agent inquires about a random or generic individual, the inquiry does not focus on any particular person. Credences formed in such inquiries therefore do not give rise to morally salient harm. For this reason, such contexts are not morally sensitive, since no individual is morally salient within them and thus do not typically require any prior adjustment.

Before turning to how this approach can be implemented in formal models, some clarifications are needed. First, what doxastic states can the agent hold toward individuals outside her epistemic focus? Second, should unfocused states also be subject to moral evaluation?

Consider the first question. When an agent forms a (precise) credence about a focused individual, she may simultaneously, and quite rationally, lack any determinate credence, or have only imprecise credences, or even hold precise credences with

²⁰Given space constraints, I do not develop the model in full here, but provide a simple illustration in Appendix B. Many thanks to an anonymous reviewer for suggesting that I consider this.

respect to other, unfocused propositions. In virtue of probabilism, her current credence function may implicitly determine credences concerning unfocused individuals, and she may recognize that increasing a focused credence entails lowering some unfocused ones. Nevertheless, these unfocused credences are not consciously assigned: the agent may be unaware of their precise values and need not track them explicitly. For this reason, she can reasonably decline to specify which unfocused credences are low, precisely because they lie outside her epistemic focus.

Requiring an agent to explicitly clarify her doxastic states — whether precise credences, imprecise credences, or the absence of any credence — toward unfocused propositions would be implausibly demanding. As boundedly rational agents, we have limited computational and deliberative capacities. If such a requirement applied to one unfocused proposition, it would equally apply to many others, quickly exceeding any plausible cognitive limit. While there may be some upper bound on what a human mind can manage, that bound is neither fixed nor easily specified.

Thus, in response to the second problem, credences concerning morally non-salient individuals in the given context may neither be actively deliberated nor explicitly specified, and for that reason they are not subject to moral evaluation. If an agent forms a moral and rational credence concerning a focused individual at the cost of lowering credences concerning some unfocused ones, she does not thereby wrong or harm anyone, since those unfocused credences are not directed at any particular individual and are therefore not morally salient.

Treating unfocused others as morally non-salient is not morally problematic, since the agent has no reason to actively consider and specify them in a context where her epistemic, pragmatic, and moral concerns are focused on a particular other individual. Rather, by focusing on the morally salient individual at a given time, the agent can deliberate more carefully and show genuine respect to that individual. In this way, respect can be extended to all individuals, since any individual may become morally salient as the agent’s focus shifts.

When she turns her focus to another proposition or individual, what is morally salient also shifts — now concerning whether that new individual has the trait. As a result, the agent should adjust her priors accordingly so that the credences about the new individual in this context remain morally permissible given the same total evidence.

For propositions that are not morally salient in a given context, agents can rationally adopt any prior distribution over them without being morally wrong (remember this paper endorses precise credences), as long as it yields morally neutral credences for the morally salient propositions — even if the same prior distribution might yield credences that, in other contexts, could seem morally wrong. The key point is that an agent’s priors should ensure that the resulting credences do not lead to *morally salient* harm in the given *morally sensitive* context; while credences about morally non-salient propositions do not cause such harm and thus fall outside this moral constraint.

As Basu and Schroeder (2018) argue, the formation of doxastic states can be shaped by agents’ (epistemic, moral, or pragmatic) needs and interests, as reflected in the focus of their inquiries. When agents are not inquiring about something, they

can rationally refrain from forming or expressing explicit views about it. For example, someone uninterested in a celebrity’s personal life may rationally refrain from expressing an opinion about what the name of the celebrity’s dog is, even when told directly (Basu and Schroeder 2018:192-3). Likewise, agents can refrain from expressing any opinion to morally non-salient propositions. Even if there are credences in those propositions, these credences can rationally express no morally significant judgment or stance, so long as the agent refrains from treating them as such.

One might object that, even when a person is not morally salient, the agent’s unfocused credence assigning a low probability to her being a manager could still harm her, who might reasonably demand an apology. However, this mischaracterizes the role of unfocused credences. An unfocused credence about someone is not an answer to the question of whether this person is a manager. Answering that question requires the agent to shift her epistemic focus and thereby change the context of inquiry. In a context in which a person is morally non-salient, the agent is not deliberating about how to assign a credence to the person at all. It seems that holding that the agent is morally at fault in such cases would therefore be overly demanding. By contrast, once the person becomes morally salient — once the agent’s focus shifts to her — assigning a low credence would be morally evaluable and morally wrong.

We are now in a position to illustrate this approach more clearly. Consider the following simplified *Workplace* example:

Small Workplace Victoria learns that in a small company with three female employees, Joanna, Anna, and Bella, only one is in management. She overhears a woman named Joanna and assigns a $\frac{1}{3}$ credence to her being in management based solely on statistical data.

Assume that Victoria’s credence in the proposition “Joanna is in management” (denoted as $cr(J) = \frac{1}{3}$) is low enough to harm Joanna and therefore morally wrong. Since Victoria is inquiring specifically about Joanna, the morally salient propositions are J and $\neg J$. Given that her total evidence is morally inadequate, she should adjust her priors and adopt a morally permissible credence which is higher than $\frac{1}{3}$, say, 0.5, for J . Then, if her credences for the other two women are greater than 0 and sum to 0.5, the average credence across all three can equal the base rate, thus preserving probabilism. Multiple possible assignments can satisfy these requirements, as shown in Table 1 (where $cr(A)$ and $cr(B)$ denote Victoria’s credences that Anna and Bella are in management, respectively).²¹ Which of these assignments should Victoria adopt? At first glance, none seem to guarantee that her credence for each woman is morally permissible. However, since only propositions concerning Joanna are morally salient, Victoria is required to ensure only her credence for J is morally permissible in this context. Therefore, all of these assignments are permissible — Victoria can rationally and morally adopt any of them when her inquiry is focused only on Joanna.

This is because, when Victoria’s inquiry concerns only whether Joanna is in management, she is required to form a judgment only about that question and need not

²¹Notice that the permissible credence assignments should not include a credence of 0, because (1) no proposition is a logical contradiction, (2) the hypothetical priors should be regular, and (3) there is no decisive evidence proving any proposition to be false.

Table 1

	Assignment 1	Assignment 2	Assignment 3	...
$cr(J)$	0.5	0.5	0.5	...
$cr(A)$	0.2	0.1	0.49	...
$cr(B)$	0.3	0.4	0.01	...

take any stance toward Anna or Bella. Her credences about the other two women therefore carry no morally salient judgment toward them.²²

The key idea is that the prior adjustment approach requires agents to assess individuals on the basis of evidence connected to their particular traits, rather than treating them merely as interchangeable members of a demographic group. This reflects the core motivation of moral encroachment: respecting individuals as distinct moral agents and preventing credences from being unduly shaped by potentially harmful demographic statistics.^{23 24}

²²Some might think the credence that “Joanna is not in management” implicitly carries attitudes toward Anna and Bella. However, since neither is morally salient in this context, the credence assignment is not directed at them specifically and need not generate morally salient harm.

²³It is important to clarify how the relevant individual is fixed in context, since this determines the appropriate credence assignment. Which credence should be assigned depends on who is the focus of inquiry. The target of a credence is the individual on whom the agent’s inquiry is focused, identified through available contextual evidence such as name, appearance, voice, behavior, or social interaction. One might worry that misidentification could lead to harmful prior adjustment by assigning the relevant credence to the wrong person. But a mistake about one identifying feature, such as a name, need not change the target of the credence. When Victoria asks whether “Joanna” is a manager, her credence concerns the woman she is interacting with under that description, even if the woman’s actual name is Anna. Since the inquiry remains focused on the same individual, the appropriate credence assignment should still concern Anna rather than Joanna, and no violation of Avoid Harming follows. The mistake concerns only the label, not the target of the credence.

By contrast, if the agent substantially misidentifies the individual under consideration, then the target of inquiry changes, and the appropriate credence assignment should change accordingly. In such cases, the agent may violate Avoid Harming. However, this is not a distinctive problem for the prior adjustment approach. Instead, errors of identification are a general limitation of belief and credence formation under imperfect information (see, e.g., Chalmers 2011). Like other forms of rational inquiry, the framework may still permit occasional mistakes under imperfect epistemic conditions, including morally suboptimal outcomes. The central idea of the prior adjustment approach is that, given the evidence available, agents should adopt priors that minimize the risk of moral harm while remaining responsive to the evidence. Therefore, full information is not required for appropriate prior adjustment. What matters is how the target of the credence is fixed in context and what evidence is accessible to the agent, and prior adjustment should track the individual whom the available evidence and inquiry concern. I thank an anonymous reviewer for raising this issue.

²⁴One might worry that shifting credences across contexts makes Victoria seem inconsistent or untrustworthy. However, context-sensitive credence reflects rational responsiveness to changing evidential and moral features rather than instability. In the misidentification case, once Victoria realizes that she is speaking to Anna rather than Joanna, her evidential situation changes, and the resulting adjustment in credence is therefore best understood as a rational response to corrected information about the identity of the person she is evaluating, rather than a violation of epistemic stability. In this respect, the change resembles ordinary rational updating under improved information. The governing norms and updating procedures remain stable even if particular credences vary in response to evidence and contexts.

Moreover, trustworthiness does not require maintaining identical credences across all contexts, but responding to evidence in a principled and explainable way. In prior adjustment cases, stability lies in the governing norms and updating rules rather than fixed credences. Victoria is not changing her view about a single person, but applying the same morally cautious norm across different contexts.

The worry about instability and loss of trustworthiness become more compelling if the agent strategically misrepresents or conceals her credences. Absent such deception or manipulation, however, context-sensitive credence adjustment does not by itself seem sufficient to undermine trustworthiness.

Table 2

	Joanna	Anna	Bella
$cr(J)$	0.5	0.3	0.2
$cr(A)$	0.2	0.5	0.3
$cr(B)$	0.3	0.2	0.5

This pattern can extend to inquiries about each woman in the company. When we ask whether each specific woman in the company is in management, in order to ensure our credences cause no harm, we should conduct separate inquiries, each focused on one individual at a time: for example, we should first ask, “Is Anna in management?” and form credences about Anna, then proceed to Bella, and so on. These separate inquiries confine the focus to a specific individual at a time, ensuring only one specific woman is morally salient in each case. Therefore, in each inquiry, the agent can adopt any credence assignment that both avoids harming the inquired individual and meets epistemic requirements.

As shown in Table 2, in the Small Workplace example, Victoria can rationally adopt different credence assignments for each separate inquiry about each of the three women. The key requirement is that, in each context, she should ensure that the credence assigned to the woman being inquired about is morally permissible, while also preserving evidentialism and probabilism. To achieve this, the prior adjustment approach requires her to adjust her priors for each inquiry so that her resulting credences are both morally and epistemically permissible.

For instance, when inquiring about Joanna, she should adjust priors to ensure $cr(J)$ higher than $\frac{1}{3}$ — for example, setting it at 0.5 — while also satisfying evidentialism and probabilism. When inquiring about Anna, she should adjust her priors again to ensure that $cr(A) = 0.5$ while meeting the same epistemic requirements; and likewise for Bella.²⁵

This framework can also be extended to specific groups.²⁶

Consider the following group-based case.

Small Workplace, Water-cooler Conversation Victoria visits a small company with four women, only one of whom is a manager, and she does not know which one. She happens upon Joanna, Anna, and Bella chatting together at the water-cooler.

In this case, three individuals become salient. How should Victoria assign credences about which of them is the manager? The individual version of the framework struggles with this case. Assigning each salient individual credence 0.5 violates probabilism, since the probabilities sum to 1.5. But preserving probabilism by assigning only one person credence 0.5 would require arbitrarily lowering the credences of the others. Victoria therefore faces a dilemma: either violate probabilism or treat some participants unfairly.

²⁵This suggests that Victoria may be required to adopt different priors for the same proposition in different contexts, even given the same total evidence. As previously argued, this is permissible within the hypothetical prior framework, provided that these adjustments are guided by constant evidential standards and updated via hypothetical prior conditionalization.

²⁶I am grateful to an anonymous reviewer for encouraging me to explore this question.

Cases of this kind are plausible everyday occurrences. They can be handled in a group version of the framework. In the individual version, we distinguish the individual under consideration from other non-salient individuals. In the group version, we can similarly distinguish the salient group from the non-salient group, which I call the in-group and the out-group.²⁷ Here, the in-group consists of Anna, Bella, and Joanna, while the out-group is the fourth woman in the company.

Similarly, there can be only one salient in-group, while there may be multiple out-groups. This captures ordinary cases in which attention is focused on one group while many others remain in the background. Mirroring the individual case, to maintain morally and epistemically permissible credences, we can assign higher, morally neutral credences to the in-group and lower credences to the out-group.

In the Water-cooler Conversation case, we can assign credence 0.7 to the in-group and 0.3 to the out-group. Then, we can treat each in-group member generically and equivalently by assigning credence to the group-level proposition, that is, to the disjunction of all in-group members. Specifically, we assign credence 0.7 to the group-level proposition that one of Anna, Bella, or Joanna is a manager, and credence 0.3 to the proposition that the fourth woman is a manager.

Probabilities are assigned to exhaustive group partitions, so individual credences need not be specified. This avoids imposing arbitrary differences among in-group members while preserving probabilistic coherence.

This group-level framework naturally subsumes the individual version, since the in-group and out-group can each consist of a single agent. Thus, when we instead assign credences to each individual, the framework reduces to the individual version of prior adjustment.

This extension captures the idea that moral judgment should be cautious: even if a group's statistics include negative cases, any individual — or any proper subgroup — may still be good. Coarse-grained statistics can misrepresent a demographic by overemphasizing negative instances and obscuring positive ones. By focusing on more refined and salient subgroups, the framework aims to minimize moral harm under evidential constraints and avoid unfairly disadvantaging members of the salient group. Even if some moral tension cannot be fully eliminated, the approach seeks to reduce unnecessary harm and avoid treating individuals merely as interchangeable members of a demographic category.

This also enables the framework to apply more naturally to ordinary conversational settings in which multiple individuals become simultaneously salient. In this way, the framework allows us to justify a more favorable credence toward the wider group, even in the presence of potentially misleading statistical data.

3.4 The appeal of the prior adjustment approach

With the framework in place, I now turn to the advantages of the prior adjustment approach as a defense of moral encroachment on credences.

²⁷One might worry about cases in which the in-group includes all individuals in the relevant population. I address this issue in Section 4.2.

This approach explains how moral considerations can influence rational credence through priors. Priors encode both epistemic and moral commitments — such as Evidentialism, Conditionalization, Probabilism, and Avoid Harming — and should be adjusted so that, when updated on the total evidence, they yield credences that are both morally and epistemically permissible by satisfying these norms. The prior adjustment approach therefore grounds moral encroachment in more fundamental epistemic and moral norms.

By allowing context-dependent adjustment of priors, the approach captures the context-sensitive nature of moral encroachment, restricts its effects to morally sensitive contexts and morally salient individuals, and avoids overgeneralization.

This approach also provides a further defense to radical moral encroachment, explains why only morally permissible credences are epistemically rational by emphasizing the moral significance of the risk of harming.

However, because it allows, and in some cases even requires, priors to change over time without any change in evidence, this approach may strike some as uncomfortable. First, it can seem fickle and raise concerns about making credences less accurate. However, the approach is not fickle at all: prior adjustment is constrained by epistemic and moral norms, context, and evidence. Within a given context, absent evidential change, and the norms impose no further requirements on credences, the agent should not adjust her priors.

Such prior adjustments need not undermine the pursuit of accuracy. By requiring resulting credences to remain epistemically and morally neutral in contexts involving ambiguous evidence, this approach can help avoid confirmation bias and support long-run expected accuracy, even for agents who lack sophisticated mathematical tools and capacities. This is because neutral credences are generated from priors that balance the influence of ambiguous or socially profiling evidence, while still appropriately accommodating evidence that directly concerns the personal traits of the focused individual.

When an agent’s total evidence is ambiguous or socially biased, uncritical reliance on it can entrench bias. Such reliance can impair the agent’s ability to form credences in similar cases without the influence of earlier biased reasoning, thereby perpetuating systematically distorted judgments. These tendencies can persist even when the resulting credences are demonstrably inaccurate, as biases replicate across analogous contexts (e.g., Festinger et al. 2017).

Adopting neutral credences in such cases can prevent the reinforcement of biased epistemic tendencies, preserve critical engagement with future evidence, and thereby support more reliable and accurate belief formation over time.²⁸

Second, one might object that this approach renders agents vulnerable to Dutch Book arguments. When an agent is offered a bet on whether Anna is a manager based solely on coarse-grained statistical evidence, she will not want to bet in accordance with a morally charitable credence.

²⁸These explanations are at least *pro tanto* compelling. For an illustrative example showing how a morally neutral credence—assumed here to be 0.5—need not impair long-run accuracy, see the Appendix. This paper does not aim to provide a full justification of the approach, since its primary goal is to articulate a plausible account of moral encroachment on credences. I pursue this further justification in a companion paper.

However, this may not pose a serious problem for this approach. First, many philosophers have expressed skepticism about the diachronic Dutch Book argument, arguing that vulnerability to a diachronic Dutch Book does not necessarily reveal probabilistic incoherence (Christensen 2000; Williamson 2000; Briggs 2009; Hedden 2015; Moss 2014, 2015). Critics of this argument often maintain that requirements on rational credences fundamentally govern the credence distribution at a given time (i.e. time-slice rationality), rather than to an agent’s entire course of credences across time. The prior adjustment approach aligns well with this view: by imposing moral and epistemic requirements only within a given context, it can effectively encode the idea of time-slice rationality.

A stronger worry, however, is that the prior adjustment approach generates aggregate inconsistency. If Victoria is willing to bet at favorable odds that Anna is in management, and later does the same for Bella, then across cases she appears to be betting as though more women are in management than she knows to be true.

This worry is important, but it also does not undermine the prior adjustment approach. The crucial point is that the prior adjustment approach does not license unconstrained aggregation of such adjusted credences across contexts. Adjusted credences operate locally within a particular inquiry while remaining constrained by the agent’s background knowledge about aggregate frequencies. It applies only when the inquiry concerns a particular individual, not when evaluating multiple individuals simultaneously.

When Anna is salient, Victoria’s credence is adjusted only within that local context and remains compatible with her background knowledge about how many women are managers. Elevated credences for different individuals across contexts do not jointly imply that more individuals possess the trait than the agent takes to be possible. This does not commit her to incoherent aggregate beliefs or betting behavior; rather, it reflects a context-sensitive judgment about which individual might have the trait.²⁹

Moreover, the betting interpretation itself may be inappropriate in morally sensitive contexts. A moral encroacher can recognize that these bets exploit her moral commitments and reasonably reject bets that treat individuals merely as objects of profitable evaluation or social profiling. The relevant question is therefore not whether every locally adjusted credence can be aggregated into a single betting strategy, but whether each credence is rationally and morally permissible within its context of inquiry.

Thus, moral considerations may justify locally adjusting credences toward particular individuals without abandoning probabilistic coherence or the agent’s knowledge of overall frequencies.³⁰

Third, one might worry that this approach, insofar as it supports the moral encroachment thesis, inherits its familiar objections. von Klemperer (2023) argues that if believing that P , suspending on P , and setting the question aside would all wrong someone, belief in not- P may become the only morally permissible option, even when

²⁹This can be captured by allowing moral weighting to redistribute probability mass within a fixed hypothesis space $H_1, \dots, H_n$ with a known number k of true propositions — rather than introducing independent probabilistic commitments across individuals.

³⁰I do not claim that this fully resolves all diachronic Dutch-book worries. Such arguments remain controversial, but residual difficulty may reflect a broader debate within epistemology and decision theory rather than a distinctive problem for the present account.

unsupported by the evidence. She suggests that a parallel problem arises for credences: “if Esther must lower her credence that Franklin is a staff member, she may thereby have to raise her credence that he is not” (2023: 2187).

This worry, however, is less forceful in the credence case, which differs importantly from the belief case. The objection assumes that the evidence independently supports the proposition. Within a Bayesian framework, however, evidential confirmation is always relative to a prior: rational credence is not determined by evidence alone, since what the evidence confirms also depends on the agent’s prior. An agent should believe a proposition only if it is confirmed by the evidence, but whether evidence confirms a proposition depends on how that evidence is interpreted through a prior distribution.

The core idea of the prior adjustment approach is that agents may rationally adopt different priors according to moral requirements and then update those priors on their total evidence through conditionalization. Evidence such as coarse-grained statistical data may seem confirmatory under one prior while failing to confirm the proposition once morally constrained priors are adopted. Prior adjustment is therefore not a departure from proper evidential response, but part of how evidence is rationally interpreted in morally sensitive contexts.

More specifically, in a Bayesian framework, evidence E confirms a proposition P iff $cr(P|E) > cr(P)$, where $cr(P)$ is the prior and $cr(P|E)$ the posterior credence conditional on E . On the prior adjustment approach, morally problematic or inadmissibly coarse-grained statistical evidence requires agents to adopt a higher prior before conditionalizing on that evidence, in order to yield morally permissible credences. In such cases, the posterior credence may fail to exceed the adjusted prior, and the evidence therefore does not confirm the proposition or support belief in it. Once more individualized and admissible evidence becomes available, no further prior adjustment is required, and the evidence may then support a higher credence or belief in the usual way.

Consider the Small Workplace example. Some might argue that the statistical evidence does not support the claim that Joanna is a manager, whereas moral encroachment would require Victoria to believe against the evidence. However, on the present approach, the evidence may support the proposition once interpreted through a different prior. Victoria may rationally adopt an intermediate prior, say 0.5, that Joanna is a manager. Suppose conditionalizing on the statistical evidence yields a posterior credence 0.6.³¹ Since the posterior credence is higher than the prior, the evidence confirms the proposition that Joanna is a manager. Thus, the view can justify Victoria’s belief that Joanna is a manager without requiring belief against the evidence; rather, different priors can alter what the evidence confirms.

The possibility that the same evidence rationally supports different credences is not unique to the present proposal. Since the prior adjustment approach allows priors to rationally vary across contexts, the same evidence can rationally support different credences, provided that agents’ priors differ in rationally and morally permissible

³¹This credence can be derived in many ways. For example, by Bayes’ theorem:

$$cr(P|E) = \frac{cr(E \cap P)}{cr(E)}$$

many probability assignments can yield a posterior credence 0.6.

ways. This kind of variation is familiar from permissivist theories in epistemology, which hold that different agents, and even the same agent over time, may rationally hold different credences based on the same evidence (e.g., Kelly 2013).³²

Thus, in cases such as Cosmos Club, assigning a lower credence to the proposition that Franklin is a staff member does not require believing against the evidence. Rather, the claim is that coarse-grained demographic evidence, once filtered through morally permissible priors, may fail to confirm the proposition in the first place. Prior adjustment applies only when the available evidence is inadmissible and does not bear sufficiently on the particular individual under consideration. Once individualized evidence concerning Franklin becomes available, credences should update accordingly. As shown in the Appendices, doing so need not compromise epistemic accuracy.

A further worry is that any credence insufficiently favorable to Franklin may itself be harmful, just as suspension can also be harmful. But this seems not obvious in the credence framework. It is unclear why Franklin could reasonably claim harm merely because someone assigns equal or higher credence to his not being a staff member than to his being one. What would be more plausibly harmful is a comparative pattern — for example, assigning a nearby white man a significantly higher credence that he is not staff on the same evidence. This is because given only statistical evidence, such differential treatment would be neither morally nor epistemically justified.

A stronger objection may arise if we accept theories on which credences map onto belief, suspension, and disbelief. Even in that case, however, the objection seems limited. Suppose, in some cases, suspension itself is harmful and only one doxastic attitude remains morally permissible. The agent may still rationally adopt a credence within the range corresponding to that attitude without thereby believing against the evidence, since the prior adjustment approach allows priors to be adjusted in an evidence-sensitive way before conditionalization, and morally constrained priors can still determine what the evidence confirms.

The core idea is that priors determine how evidence is interpreted, and prior adjustment must therefore remain evidence-sensitive. When only inadmissible evidence is available, agents need to adjust their priors to avoid moral harm. When individualized and admissible evidence becomes available, credence should shift accordingly. In both cases, the resulting credences remain evidence-based; what changes is the prior through which the evidence is interpreted. The view therefore does not require believing against the evidence.

At the same time, I suggest that the prior adjustment approach provides a way of defending moral encroachment on credences that differs from other prior based accounts which also emphasize priors but deny moral encroachment by reducing moral considerations to epistemic risk — most notably, Johnson King and Babic’s epistemic risk account (2020).

Johnson King and Babic also argue that moral considerations affect the adoption of priors, but they explain this through the norm of minimizing epistemic risk (2020: 90). Epistemic risk concerns the costs of assigning too high a credence to a falsehood or too low a credence to a truth. Such risk influences the agent’s adoption of priors. They

³²This may also be explained by epistemic fragmentation, according to which epistemic status, and even the application of evidence, can rationally shift across contexts (e.g., Greco 2014, 2019).

argue that agents should adopt priors whose resulting credences minimize expected epistemic risk.

However, this account seems to face three problems. First, it seems merely to delay, rather than prevent, the formation of morally wrong credences. For example, if the agent receives strong statistical data indicating an extremely low base rate of women in management, then — unless the mean of her prior distribution is extremely high — she will ultimately adopt a correspondingly low, and thus morally wrong, credence about the specific woman. This result is precisely what moral encroachment theorists seek to avoid.

Second, the evaluation of epistemic risk seems to depend on expected (in)accuracy, which requires assigning probabilities to possible outcomes — that is, estimating how likely the proposition is true or false — and then calculating the cost of being wrong. But such probabilities are difficult to determine, since assessing epistemic risk already requires estimating how likely the proposition is to be true or false. When the agent has only statistical evidence, those estimates will likely track the underlying statistics. If so, the resulting risk assessment may reproduce the same problematic statistical assumptions and moral mistakes the framework aims to avoid. Moreover, it seems that one must already rely on certain credences or statistical data to assess which credences are more or less risky. Since epistemic risk is then used to guide the choice of priors, the account threatens to be worryingly circular.

Third, it seems to offer no method for addressing the probabilist objection — namely, that increasing credences about some individuals may violate probabilism or come at the cost of lowering credences about others, and vice versa.

By contrast, the prior adjustment approach addresses these problems by clarifying that it is moral norms that influence credence formation, rather than reducing moral influence to purely epistemic considerations. By directly integrating moral and epistemic norms into agents' evidential standards, it allows agents to adopt priors that, when conditionalized on their total evidence, yield credences that are both morally and epistemically permissible in a given morally sensitive context.

Importantly, this approach allows priors to be adjusted in response to the requirements of moral and epistemic norms across different contexts. By requiring credences to remain morally neutral and epistemically rational in morally sensitive cases, it avoids the first problem faced by Johnson King and Babic's account. No matter how much statistical evidence an agent has, so long as that evidence is coarse-grained, biased, or underspecified, she may adjust her priors to counterbalance its effects and maintain morally neutral credences.

Another key appeal of this approach is that it helps to reconcile moral and epistemic requirements. The next section explains how this reconciliation is achieved by showing how the approach responds to objections from evidentialism and probabilism.

4 Responses to objections

4.1 Evidentialism

I have argued that the *genuine* evidentialist objection to moral encroachment rests on two principles: Evidentialism (Credences) and Conditionalization.

The prior adjustment approach satisfies both principles, since credences are formed by conditionalizing priors on the agent's total evidence. Evidential constraints therefore apply through the priors from which credences are derived.

Priors may vary across contexts, but their adjustment must still fit the agent's total evidence. If contextual factors count as part of the evidence, then this approach also allows credences to vary across contexts in proper response to total evidence.

Thus, on this account, morally constrained credences can remain evidentially grounded. Without evidence, an agent cannot legitimately move beyond her priors to form any further credence, let alone moral and rational ones. In this sense, moral encroachment does not replace evidential constraints, but operates through them.

4.2 Probabilism

I now turn to objections concerning probabilistic coherence. Critics argue that moral encroachment may conflict with probabilism, especially when all individuals are simultaneously salient.

The prior adjustment approach preserves probabilistic coherence by treating inquiry as context-sensitive. By endorsing Avoid Harming rather than Impartiality, it allows agents to assign credences relative to morally salient individuals or subgroups. However, both the individual- and subgroup-level frameworks still depend on moral salience, and thus do not by themselves resolve the full-population case.

Goodsell and Hawthorne's argument seems to suggest a stronger objection: if all individuals in a sample are simultaneously morally salient, how can moral encroachment be satisfied without violating probabilism?

However, in practice, agents need not, and should not, evaluate all individuals at once. Moral encroachment, by its nature, rejects treating all individuals in this way solely based on general demographic statistics, since this amounts to social profiling — which treats specific individuals as interchangeable group members and undermines respect for their distinct subjectivity.

Instead, the agent can evaluate each inquiry separately — whether concerning an individual or a salient subgroup — and assign a credence above the base rate within that context. Since each judgment concerns a distinct morally salient target in a given context, probabilistic coherence is also preserved.

If one insists that all individuals are simultaneously morally salient, the outcome depends on which moral norm is prioritized. Given only coarse demographic evidence, the agent faces two options: assign base-rate credences to all (impartiality), resulting in equal harm, or minimize overall harm by protecting as many as possible at the cost of harming a few. Both preserve probabilism but carry costs: the first weakens moral encroachment, while the second allows that it may sometimes require harm.

This suggests that the difficulty may arise from limitations in the available evidence rather than from the framework itself.

Indeed, the full-population dilemma arises primarily when the available evidence is too coarse-grained to permit meaningful differentiation among individuals or subgroups, for example when agents rely solely on demographic statistics. In such cases, the account may not fully eliminate moral tension. The central claim is that moral salience should track evidential resolution: when sufficiently individualized evidence

is available, agents should avoid relying solely on coarse demographic statistics in ways that risk harming individuals by treating individuals merely as interchangeable members of a group.

As evidence increases, moral salience can become correspondingly more fine-grained. More individualized evidence allows agents to partition populations into more specific and evidentially relevant subgroups, thereby enabling prior adjustment to operate through those morally salient distinctions.

By contrast, when only highly coarse-grained evidence is available, agents should not treat all individuals as simultaneously morally salient. The apparent difficulty at the full-population level is therefore evidential rather than substantive: it reflects the general limitations of credence formation under highly limited information, rather than a defect specific to the framework itself. The full-population case should thus be understood as a limiting case rather than a stable counterexample.

For those who reject restricting simultaneous moral salience, moral encroachment may be less compelling as a general account. However, the aim here has been to offer compelling support for moral encroachment. Given the pressure of this objection, there may be limits to what can be achieved. There is no simple or definitive solution, but this may be as far as the argument can reasonably go.

5 Conclusion

This paper argues that the prospects for moral encroachment on credences are more promising than some critics suggest. Accepting the prior adjustment approach makes it possible to account for moral encroachment on credences without undermining fundamental epistemic requirements or presupposing sophisticated evidential processing capacities.

Moreover, this approach is flexible enough to incorporate additional moral³³ and even pragmatic norms. It can extend beyond moral encroachment on credences to belief, given views that connect credence with belief, suspension, and disbelief. It may also accommodate both radical accounts and more moderate views that locate harm in actions or evidential practices.

As such, the framework may provide a promising basis for pragmatic encroachment. For instance, pragmatic norms (such as exercising heightened caution when the cost of acting wrong is high) could justify context-sensitive adjustments of priors, though any such extension would require showing that the additional norms are compatible with the existing epistemic constraints.

I do not claim that this approach should be universally accepted. Its aim is to show how moral considerations can shape the credences of boundedly rational agents under coarse evidence and limited reasoning. Those who reject moral encroachment or favor more technical approaches may find it less appealing.

Further work is needed to fully develop this account, including assessing the hypothetical prior framework and clarifying morally neutral and favorable credences.

³³In this paper, this approach is primarily used to explain how being morally wrong affects the rationality of credences, but it can also account for other moral features — such as moral favorability — by integrating the relevant norms into evidential standards.

Still, I hope this paper offers a plausible support and advances discussion of moral encroachment.

Appendix A An Illustrative case of morally neutral credence and accuracy

Proof Let $\{A, B\}$ be an agent's total evidence about proposition P about an individual at a given time t , where A is ambiguous and biased social statistical evidence and B is evidence about the individual's personal traits. According to the hypothetical prior conditionalization and the prior adjustment approach (assume the neutral credence is $pr(P|A) = 0.5$), the agent's credence in P at t should be:

$$\begin{aligned} cr_t(P) &= pr(P|A, B) \\ &= \frac{pr(B|P, A) \cdot pr(P|A)}{pr(B|A)} \\ &= \frac{pr(B|P, A) \cdot 0.5}{pr(B|A)} \\ &= \frac{pr(B|P, A) \cdot 0.5}{pr(B|P, A) \cdot pr(P|A) + pr(B|\neg P, A) \cdot pr(\neg P|A)} \\ &= \frac{pr(B|P, A) \cdot 0.5}{pr(B|P, A) \cdot 0.5 + pr(B|\neg P, A) \cdot 0.5} \end{aligned}$$

Define the likelihood ratio for B :

$$R := \frac{pr(B|P, A)}{pr(B|\neg P, A)}$$

We can infer that $pr(B|P, A) = pr(B|P)$ and $pr(B|\neg P, A) = pr(B|\neg P)$, provided that evidence A (e.g., statistical data) and evidence B (e.g., more decisive evidence about the individual) are conditionally independent given P . For example, once the truth value of P is known (e.g., Joan is in management), statistical evidence A (e.g., "1% of women are managers") no longer affects evidence B (e.g., "Joan is working in manager's office").

Suppose evidence B confirms P , then we have:

$$pr(P|B) > pr(P) > pr(\neg P|B)$$

Thus, $R > 1$.

Express $cr_t(P)$ in terms of R :

$$cr_t(P) = \frac{R}{R+1} > \frac{1}{2}$$

Similarly, when evidence B disconfirms P , then

$$cr_t(P) = \frac{R}{R+1} < \frac{1}{2}$$

□

Therefore, the prior adjustment approach can maintain the accurate long-term credence formation as evidence about the individual's personal traits unfolds.

Appendix B An Illustrative Model of Prior Adjustment

In this appendix, I present a simple model illustrating the prior adjustment approach and how priors may be selected by integrating epistemic and moral considerations. Conditionalizing these priors on the agent’s total evidence yields credences that are both epistemically and morally permissible. The aim is not to provide a complete theory, but to illustrate a plausible framework in which moral considerations can affect which credence function is operative.

The central idea is that prior selection is subject to both epistemic and moral norms. Building on accuracy-first approaches — typically taken to support evidentialism and probabilism — we can model this by assigning each candidate prior both an accuracy score and a moral score. A prior is then permissible if it performs sufficiently well on both dimensions, or if it maximizes an appropriate weighted combination of epistemic accuracy and moral considerations.

Let $\Pi = \{pr_1, pr_2, \dots, pr_n\}$ be the set of available priors. Each prior pr receives an accuracy score, $A_D(pr)$, and a moral score, $M_c(pr)$.

The moral score measures the moral importance of expected risk of harm to a salient individual, which can be defined as:

$$M_c(pr, \omega) = - \sum_{i \in I_c} w(i) h_i(pr)$$

where I_c is the set of individuals relevant in context c ; $w(i)$ measures the moral importance of individual i ,³⁴ and $h_i(pr)$ measures the expected risk of harm to i induced by pr .

The accuracy score measures the expected diachronic accuracy³⁵ of a prior assignment, which can be defined as:

$$A_D(pr) = \mathbb{E} \left[\sum_{t=1}^T a(pr_t, s_t) \right]$$

where $a(pr_t, s_t)$ is the accuracy score of the credence function generated by prior pr at time t in state s_t , evaluated over a time horizon T that runs from the stage at which the agent has no admissible evidence to the stage at which she has acquired all admissible evidence.

³⁴In morally insensitive contexts, or when no morally important individuals are at stake, the moral weight w_i is negligible or zero, so prior selection is determined primarily by considerations of accuracy. In morally sensitive contexts involving morally important individuals, $w(i)$ is assigned in proportion to the moral importance of individual i , such that individuals with greater moral importance receive higher weights.

³⁵Here, “expected diachronic accuracy” refers to the sum of the expected epistemic accuracy of a credence function at each time-slice as it is updated on new evidence over time. It measures the total epistemic accuracy an agent is expected to achieve in the long run. A credence function may therefore have higher diachronic accuracy even if it is initially less accurate, provided that it becomes more accurate as further admissible evidence is acquired. This captures the prior adjustment approach: morally adjusted credences may sacrifice some short-term accuracy while achieving greater long-run accuracy by correcting for the influence of biased statistical evidence.

Table B1

Prior	$A_D(pr)$	$M_c(pr, w_A)$
pr_1 : coarse statistical prior	0.7	0.3
pr_2 : moderately increased prior	0.8	0.85
pr_3 : strongly increased prior	0.6	0.95

Then the morally admissible set of priors is:

$$\Pi_c^M = \{pr \in \Pi : M_c(pr, w) \geq \mu_c\},$$

or equivalently:

$$\Pi_c^M = \{pr \in \Pi : \sum_{i \in I_c} w(i)h_i(pr) \leq \rho_c\},$$

where ρ_c is the maximum morally permissible level of weighted risk.

The selected prior is then the maximally diachronically accurate one:

$$pr^* = \arg \max_{pr \in \Pi_c^M} A_D(pr).$$

where pr^* is the admissible prior.

This says: moral importance and risk of harm constrain which priors are admissible; among those, the agent chooses the prior with maximal diachronic accuracy.³⁶

Thus, prior adjustment is not arbitrary. The agent does not simply “boost” credences whenever someone is in focus. Rather, she selects a prior from a set of available priors that satisfies both epistemic and moral constraints, and among these chooses the one that maximizes expected accuracy.

Consider the Small Workplace example. Suppose Victoria is evaluating whether Anna is a manager, and considers three candidate priors (assume, for simplicity, that all women have equal moral importance) as listed in Table B1.

In a morally sensitive context c , suppose the moral admissibility threshold is $\mu_c = 0.7$. Then pr_2 and pr_3 are morally admissible, since $M_c(pr_2, w_A) = 0.85 > \mu_c$ and $M_c(pr_3, w_A) = 0.95 > \mu_c$. By contrast, pr_1 is excluded. Among the admissible priors, pr_2 has the highest diachronic accuracy, since $A_D(pr_2) > A_D(pr_3)$. Victoria should therefore adopt pr_2 .

References

- [1] Babic, B., Gaba, A., Tsetlin, I., and Winkler, R. L. (2024). Normativity, epistemic rationality, and noisy statistical evidence. *The British Journal for the Philosophy*

³⁶One might worry that the prior that maximizes expected accuracy lies outside the morally admissible set. I do not attempt a general proof here. However, as suggested in Appendix A, morally permissible prior adjustments — those that yield appropriately neutral credences — typically fall within a moderate range that avoids extreme or biased assignments. Priors in this range are often competitive in, and may even improve, expected accuracy relative to more biased alternatives. A fuller treatment of this issue is beyond the scope of the present paper. A more detailed treatment of this issue is left for future work.

of Science, 75 (1), 153-176.

- [2] Basu, Rima. 2018. The wrongs of racist beliefs. *Philosophical Studies*, 176 (9):2497-2515.
- [3] Basu, Rima. 2019. Radical moral encroachment: The moral stakes of racist beliefs. *Philosophical Issues*, 29 (1):9-23.
- [4] Basu, Rima. and Mark Schroeder. 2018. Doxastic Wronging. In Kim, Brian. and Matthew McGrath. *Pragmatic Encroachment in Epistemology*. New York: Routledge. pp.181-205.
- [5] Bolinger, Renée Jorgensen. 2020. The rational impermissibility of accepting (some) racial generalizations. *Synthese*, 197 (6):2415-31.
- [6] Briggs, R. 2009. Distorted reflection. *Philosophical Review* 118 (1):59-85.
- [7] Buchak, Lara. 2014. Belief, credence, and norms. *Philosophical studies*, 169: 285-311.
- [8] Chalmers, David. 2011. Frege's Puzzle and the Objects of Credence. *Mind*, 120: 587-635.
- [9] Christensen, David. 2000. Diachronic Coherence Versus Epistemic Impartiality. *Philosophical Review*, 109: 349-71.
- [10] Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in epistemology*. Clarendon Press.
- [11] Festinger, L., Riecken, H. and Schachter, S. 2017. *When prophecy fails: A social and psychological study of a modern group that predicted the destruction of the world*. Lulu Press, Inc.
- [12] Friedman, J. 2017. Why suspend judging? *Noûs*, 51 (2), 302-326.
- [13] Fritz, James. 2020. Moral encroachment and reasons of the wrong kind. *Philosophical Studies*, 177 (10):3051-3070.
- [14] Fritz, James, and Elizabeth Jackson. 2021. Belief, credence, and moral encroachment. *Synthese*, 199 (1): 1387-1408.
- [15] Gardiner, Georgi. 2018. Evidentialism and Moral Encroachment. In McCain Kevin. *Believing in Accordance with the Evidence: New Essays on Evidentialism*. Cham: Springer Verlag.
- [16] Goodsell, Zachary, and John Hawthorne. 2025. Morality Does Not Encroach. In Comesaña, Juan, and Matthew McGrath. *Knowledge and rationality: essays in honor of Stewart Cohen*. Abingdon, Oxon: Routledge.

- [17] Greaves, H. and D. Wallace. 2006. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind* 115: 607–32.
- [18] Greco, Daniel. 2014. Iteration and Fragmentation. *Philosophy and Phenomenological Research*, 88.
- [19] Greco, Daniel (2019). Fragmentation and Higher-Order Evidence. In Matthias Skipper & Asbjørn Steglich-Petersen, *Higher-Order Evidence: New Essays*. Oxford, United Kingdom: Oxford University Press. pp. 84-104.
- [20] Harman, Elizabeth. 2009. Harming as causing harm. In Wasserman, David, and Melinda, Roberts. *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*. Springer. pp. 137-54.
- [21] Hedden, Brian. 2015. Time-Slice Rationality. *Mind*, 124 (494): 449–491.
- [22] Kelly, Thomas. 2013. Evidence can be permissive. *Contemporary debates in epistemology*, 298-311.
- [23] King, Zoe Johnson, and Boris, Babic. 2020. Moral obligation and epistemic risk. *Oxford studies in normative ethics*, 10: 81-105.
- [24] Lewis, David. 1980. A Subjectivist’s Guide to Objective Chance. In Richard C. J. (Ed.), *Studies in Inductive Logic and Probability 2*: 83–132. University of California Press.
- [25] Marušić, Berislav. and White, Stephen. 2018. How Can Beliefs Wrong?: A Strawsonian Epistemology. *Philosophical Topics*, 46 (1):97-114.
- [26] Meacham, Christopher JG. 2007. Chance and the Dynamics of de Se Beliefs (Doctoral dissertation). Retrieved from <http://dx.doi.org/doi:10.7282/T30Z73N5>
- [27] Meacham, Christopher JG. 2016. Ur-Priors, Conditionalization, and Ur-Prior Conditionalization. *Ergo: An Open Access Journal of Philosophy* 3: 444-92.
- [28] Moss, Sarah. 2014. Credal Dilemmas. *Noûs*, 48 (3): 665–683.
- [29] Moss, Sarah. 2015. Time-slice epistemology and action under indeterminacy. In T. Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology*, 5: 172-193.
- [30] Moss, Sarah. 2018. “IX — Moral encroachment.” In *Proceedings of the aristotelian society*, 118(2):177-205. Oxford University Press.
- [31] Munton, Jessie. 2019. Beyond accuracy: Epistemic flaws with statistical generalizations. *Philosophical Issues*, 29 (1):228-240.

- [32] Pace, Michael. 2010. The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe'. *Noûs*, 45 (2):239-268.
- [33] Robert, C. P. 2007. *The Bayesian choice: from decision-theoretic foundations to computational implementation*. New York, NY: Springer New York.
- [34] Ross, Jacob, and Mark Schroeder. 2014. Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88 (2):259-288.
- [35] Schoenfield, Mark. 2014. Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, 48 (2): 193-218.
- [36] Schroeder, Mark. 2018. When Beliefs Wrong. *Philosophical Topics*, 46(1):115-127.
- [37] Skyrms, Brian. 1983. Three Ways to Give a Probability Function a Memory. In Earman John. (Ed.), *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science*, 10: 157-161. Minnesota University Press.
- [38] Thomson, Judith Jarvis. 1976. Killing, Letting Die, and the Trolley Problem. *The Monist* 59 (2):204–17.
- [39] Titelbaum, Michael G. 2013. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press.
- [40] Titelbaum, Michael G. 2022. *Fundamentals of Bayesian epistemology 1: Introducing credences*. Oxford University Press.
- [41] von Klemperer, Caroline. 2023. Moral Encroachment, Symmetry, and Believing Against the Evidence. *Philosophical Studies*, 180 (7): 2175-2190.
- [42] White, Roger. 2005. Epistemic permissiveness. *Philosophical perspectives*, 19, 445-59.
- [43] Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- [44] Woodard, Elise. 2022. A puzzle about fickleness. *Noûs*, 56 (2): 323–342.
- [45] Ye, Ru. 2024. It can be rational to change priors. *Analysis* 84:822-30.