

A Bayesian explanation of the irrationality of sexist and racist beliefs involving generic content

Paul Silva


Synthese

An International Journal for
Epistemology, Methodology and
Philosophy of Science

ISSN 0039-7857

Synthese

DOI 10.1007/s11229-018-1813-9



A Bayesian explanation of the irrationality of sexist and racist beliefs involving generic content

Paul Silva Jr.¹

Received: 15 December 2017 / Accepted: 10 May 2018 / Published online: 26 May 2018
© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract Various sexist and racist beliefs ascribe certain negative qualities to people of a given sex or race. Epistemic allies are people who think that in normal circumstances rationality requires the rejection of such sexist and racist beliefs upon learning of many counter-instances. This is a common view among philosophers and non-philosophers. But epistemic allies face three problems. First, sexist and racist beliefs often involve generic propositions. These sorts of propositions are notoriously resilient in the face of counter-instances. Second, background beliefs can enable one to explain away counter-instances to one's beliefs, thus making it rational to retain one's beliefs in generics in the face of many counter-instances. The final problem is that the kinds of judgements epistemic allies want to make about the irrationality of sexist and racist beliefs upon encountering many counter-instances is at odds with the judgements that we are inclined to make in seemingly parallel cases about the rationality of non-sexist and non-racist generic beliefs. Thus epistemic allies may end up having to give up on plausible normative supervenience principles. In what follows I explain how a Bayesian approach to the relation between evidence and belief can neatly untie these knots. The basic story is one of defeat: Bayesianism explains when one is required to become increasingly confident in chance propositions, and confidence in chance propositions can make belief in corresponding generics irrational.

Keywords Sexism · Racism · Generics · Epistemology · Bayesianism

✉ Paul Silva Jr.
psilvajr@gmail.com

¹ University of Pennsylvania, Philadelphia, USA

1 Introduction

Various sexist and racist beliefs ascribe certain negative qualities to people of a given sex or race. *Epistemic allies* are people who think that in normal circumstances rationality requires the rejection of such sexist and racist beliefs upon learning of many counter-instances, i.e. members of these groups who lack the target negative quality. Accordingly, epistemic allies think that those who give up their sexist or racist beliefs in such circumstances are rationally responding to their evidence, while those who do not are irrational in failing to respond to their evidence by giving up their belief. This is a common view among philosophers and non-philosophers.¹

But epistemic allies face three problems. First, sexist and racist beliefs often involve generic propositions. These sorts of propositions are notoriously resilient in the face of counter-instances since the truth of generic propositions is typically compatible with the existence of many counter-instances. Second, background beliefs can enable one to explain away counter-instances to one's beliefs. So even when counter-instances might otherwise constitute strong evidence against the truth of the generic, the ability to explain the counter-instances away with relevant background beliefs can make it rational to retain one's belief in the generic despite the existence of many counter-instances. The final problem is that the kinds of judgements epistemic allies want to make about the irrationality of sexist and racist beliefs upon encountering many counter-instances is at odds with the judgements that we are inclined to make in seemingly parallel cases about the rationality of non-sexist and non-racist generic beliefs. Thus epistemic allies may end up having to give up on plausible normative supervenience principles. All together, these problems pose a significant *prima facie* challenge to epistemic allies. In what follows I explain how a Bayesian approach to the relation between evidence and belief can neatly untie these knots. The basic story is one of defeat: Bayesianism explains when one is required to become increasingly confident in chance propositions, and confidence in chance propositions can make belief in corresponding generics irrational.

2 Some quick clarifications

In speaking of 'sexist' and 'racist' beliefs I have in mind, quite roughly, beliefs like this:

- The belief that women are less intelligent than men.
- The belief that black people are not as diligent as white people.
- The belief that Latinos are more violent than white people.

I take these sorts of beliefs to be paradigmatic instances of sexist and racist beliefs that epistemic allies are keen to oppose.

It will help us in what follows to fix our attention on a question addressed by Arpaly (2003), Fricker (2007), and Begby (2013):

¹ Arpaly (2003) and Fricker (2007). Anecdotally, I've heard this position espoused by numerous colleagues, university staff, family, and friends.

Q1. In normal circumstances is it irrational to maintain one's belief that women are less intelligent than men *upon learning of many intelligent women*?

In limiting our attention to 'normal circumstances' we draw attention to the sorts of circumstances people actually tend to find themselves in; the sorts of circumstances you, your neighbors, your friends, and your family members tend to find themselves. This is just to set aside some of the more remote and unusual scenarios that philosophers might come up with.²

There are a few things to say about the content of the belief *that women are less intelligent than men* and the idea of learning of *many intelligent women*. First, the term 'intelligence' can be used in different senses insofar as it has become customary to speak of different kinds of intelligence (emotional intelligence, kinesthetic intelligence, etc.). The target notion of intelligence in this paper is whatever notion involves or is most closely related to a thinker's abstract reasoning abilities—since that's the notion Arpaly (2003), Fricker (2007), and Begby (2013) have in mind. Being more specific than that won't make a difference to the general lessons in what follows; what matters is that our judgements concerning this sort of intelligence are meaningful and that our evidence can support such judgements.

Second, the terms 'men' and 'women' are kind terms and to say that men or women are intelligent is just to ascribe a property of intelligence *to these kinds*.³ Sometimes in discussions of the comparative intelligence of men and women what people mean when they assert *that women are less intelligent than men* is a claim about the *most intelligent* of these kinds: *the most intelligent women are less intelligent than the most intelligent men*. I'll return to this at the end of the paper.

Third, in speaking of learning of *intelligent women* in a sense that would count as a counter-instance to the claim that women are less intelligent than men, we have to be referring to learning of individual women whose degree of intelligence is *at or above* the threshold implicitly ascribed to men when making the comparison that women are less intelligent than men.

Finally, and perhaps most importantly for present purposes, the proposition *that women are less intelligent than men* is a generic proposition. I'll say more about what this means and its relevance in the next section.

² Q1 is an epistemically important question in thinking about the epistemology of sexist and racist beliefs. But it's not the only important question. For example, consider the following questions:

Q2. In normal circumstances is it irrational to maintain one's belief that women are less intelligent than men given the fact that *one ought to be aware of many intelligent women*?

Q3. In normal circumstances is it irrational to maintain one's belief that women are less intelligent than men *given all the reliable testimony one has, or could easily have, to the effect that it's false*.

Space unfortunately prohibits me from discussing these questions in further detail. For more on Q2 see Goldberg (2016) and Benton (2016).

³ We regularly ascribe properties to kinds, e.g. the dodo bird is extinct, a duck is a biological organism, coffee beans are small, etc. Notice that the truth of such statements doesn't depend on any member of the kind existing. 'a duck is a biological organism' is true even if ducks ceased to exist, and 'Dodo birds are extinct' is true in part *because* no dodo exists.

3 Three problems for epistemic allies

When it comes to sexist beliefs people either find themselves in an evidential position where they *lack sufficient reason* to think these sexist beliefs are true; or people find themselves in an evidential position where they *have sufficient reason* to think these sexist beliefs are true. Questions about the epistemic rationality of sexist beliefs in the former position are uninteresting; almost by definition, sexist beliefs are irrational when one has insufficient reason to believe them. Perhaps most people are in such a position: their sexist beliefs are just irrational from the get-go. That said, it's fairly easy to imagine people who have misleading evidence that supports the sexist belief that women are less intelligent than men. And we should hope that an adequate epistemology could explain how learning of many intelligent women could undermine that misleading evidence in the way epistemic allies think it does, and without running afoul of other serious problems.

So the interesting question I want to address is what we should think of Q1 against the backdrop of someone who starts out in an evidential position where it is rational to have a sexist belief and then later encounters many counter-instances to that belief. So consider a slight variation on the toy case that Arpaly (2003, p. 103ff), Fricker (2007, p. 33ff), and Begby (2013, p. 91ff) work with:

Solomon grew up in an isolated town where he was taught that:

(X) Women are not as intelligent as men.

His testimonial evidence for X is robust: all of the most informed and thoughtful people in his town taught him that their experiences with women have confirmed that X is true. Moreover, while living in his home town, Solomon failed to acquire any evidence for thinking X is false. Indeed, [let us implausibly suppose] that *in Solomon's town* it's actually true that women are rarely as intelligent as men, and Solomon gets strong evidence for thinking this from classroom discussions with women, women's test scores, and everyday interactions.

Assuming that the rationality of Solomon's believing X in these circumstances is determined by what his evidence supports, with this kind of background there's little room for doubt that it is epistemically rational for Solomon to believe X. After all, his testimonial evidence as well as his observations of low women's intelligence strongly support X. Arpaly (2003) and Fricker (2007, p. 34) agree that someone like Solomon "could not be accused of any marked irrationality" in maintaining their belief about the intelligence of women.

However, Arpaly and Fricker invite us to consider an evolution in Solomon's circumstances. Suppose he moves to another city to attend university and his story continues respectively:

... Solomon has just moved to New City to attend university, and while living in New City Solomon frequently encounters intelligent women. This is especially (but not exclusively) the case at his university where he studies alongside many intelligent women.

At this point, epistemic allies like Arpaly (2003, p. 103ff) and Fricker (2007, p. 33ff) claim that *rationality requires* Solomon to reject X. Summarizing and affirming Arpaly's view, Fricker (2007, p. 34) writes that if Solomon's belief about the intelligence of women doesn't change after having studied alongside so many intelligent women at university then:

...his belief *is revealed as irrational*, and moreover a prejudice: the stubbornness of Solomon's belief *in the face of manifest counter-evidence* would at once reveal him as both epistemically and ethically flawed.

The exact reasoning towards this conclusion is not explicitly stated, but I suspect it's something like this: as Solomon encounters more and more intelligent women in New City he ought to become less and less confident in X. For he ought to treat each instance of an intelligent woman as evidence against X, and slowly his credence in X should erode to the point where he no longer believes X and he becomes quite confident that men are not more intelligent than women.⁴

But epistemic allies face various problems at this point. The first problem concerns the nature of the content of X. For X is a *generic proposition*—like 'birds fly' or 'ducks lay eggs' or 'trees grow leaves'. Generic statements like these are generalizations about kinds.⁵ But generics are not like other generalizations we are familiar with. To say that ducks lay eggs is not to say that *all* or *most* or even that some *exact proportion* of ducks lay eggs; it's also not to say that *the average* duck lays eggs;⁶ it's not even to say that *in normal circumstances* all or most or some exact proportion of ducks lay eggs; nor is it to say that in normal circumstances the average duck lays eggs (Cohen 1999; Leslie 2007, 2008; Leslie and Lerner 2016). For notice that *most ducks don't lay eggs* (half are male and some females are infertile) and this is true in perfectly normal circumstances. For this reason generics can be true in the face of many counterinstances. After all, discovering that more than half of ducks don't lay eggs doesn't falsify the claim that ducks lay eggs; and discovering many flightless birds doesn't falsify the claim that birds fly; and discovering many leafless trees doesn't falsify the claim that trees grow leaves; etc. The question this raises for epistemic allies is this: since the truth of a generic like X is consistent with there being many intelligent women, how could it be the case that *rationality requires* someone like Solomon to reject X upon encountering many intelligent women?

⁴ This is not the only explanation for why encountering so many intelligent women ought to lead Solomon to reject X. Perhaps, as Pace (2011) and Basu (2018) argue, moral considerations can encroach on rational belief even as more familiar pragmatic considerations can. But this does not seem to be the kind of explanation Arpaly and Fricker invite us to consider here, and those resistant to pragmatic encroachment will obviously take issue. It would also be surprising if there were no purely evidential justification for epistemic allies. So my hope is that some kind of purely evidential story can be told even if epistemic allies can help themselves to further resources.

⁵ Cf. footnote 3.

⁶ To say that *the average* F is G is not to say Fs are G. Suppose every blonde-haired human died. Then it would be true that *the average human is not blonde-haired*; but it doesn't follow from this that *humans are not blonde-haired*. The reason for this is that statements about averages are mathematical functions *over a population*, but generics are about *kinds*. Begby (2013) sometimes wrongly treats generic claims as claims about averages.

The second problem epistemic allies face concerns the role background beliefs can play in helping explain away counter-instances to one's sexist and racist beliefs. It's long been observed that people regularly rationalize the existence of exceptions to their beliefs, and that this is no less true of sexist and racist beliefs (Allport 1979, pp. 23–24; Weber and Crocker 1983; Kunda and Oleson 1995; Schneider 2004, p. 209ff). Still you might think that it's possible for Solomon to encounter a large enough number of intelligent women that he would have to give up his belief in X. But as Begby (2013, p. 96) nicely explains, Solomon seems able to rationally explain away his encounters with many intelligent women, and hence continue to rationally believe X:

It is by now a familiar insight from confirmation theory that people with different background beliefs (“theories”) may reasonably differ in their assessment of the significance of new evidence. The point can be conveniently stated in terms of Bayesian theory... Solomon might assign a very low prior probability to the hypothesis that women are as intelligent as men at the same time as he assigns a relatively high prior probability to finding evidence of female intelligence in a university setting. In light of these prior probabilities, his encounters with intelligent university women demands little if any adjustment to his background beliefs. But with different background beliefs, the significance of the evidence might be very different.

This is not incorrect; and I will further explain why below. But the basic idea is that epistemologists have long observed that the rational response to one's new evidence is not simply determined by one's new evidence. It's also in part determined by one's background information. So different agents with different background information can rationally respond *in different ways* to the same new evidence. Accordingly, it really is a possibility that Solomon rationally responds to his repeated encounters with many intelligent women without ever decreasing, or at most only slightly decreasing, his credence in X. But epistemic allies seem to assume that this is not a possibility. So epistemic allies seem to have another problem.

The third problem concerns the extent to which epistemic allies are able to evaluate like cases alike—as normative supervenience principles advise. Consider the fact that we rationally believe a wide range of generic propositions: we believe that trees grow leaves, that mammals walk on land, that dogs have longer tails than chipmunks, and so forth. Moreover, we would likely continue to rationally believe these things even if we were to encounter many counter-instances to them. For example, we'd continue to rationally believe that dogs have longer tails than chipmunks even if we encountered many tailless dogs; it's not as if we believed that *all* dogs have longer tails than chipmunks. So were we to encounter many tailless dogs it would be rational for us to simply assume that our circumstances are for some unknown reason presently favorable to encountering tailless dogs, even though dogs have longer tails than chipmunks. But it's not clear how epistemic allies could coherently maintain this position in Solomon's case. For Solomon's belief in X has strong rational support just as our belief that dogs have longer tails than chipmunks has strong rational support. So if epistemic allies think that Solomon is rationally required to reject X upon encountering many intelligent women, then it seems like they should say that we would be rationally required to reject the claim that dogs have longer tails than chipmunks upon encountering many

tailless dogs. But, arguably (and I will argue for this below) we would not ordinarily be rationally required to give up our belief that dogs have longer tails than chipmunks upon encountering many tailless dogs. If epistemic allies accept this, they seem to be violating normative supervenience principles that tell us to treat like cases alike.⁷

In summary, epistemic allies face three problems. They seem insensitive to the fact that generics can be true in the presence of many counter-instances; they seem insensitive to the fact that one's background beliefs can enable one to rationally explain away counter-evidence; and they seem committed to evaluating like cases differently and so seem to run afoul of normative supervenience principles.

4 Bayesianism for epistemic allies

The aim of this section is to explain how a lite form of Bayesianism can resolve each of the three problems above on behalf of epistemic allies. I first explain the conditions under which encountering intelligent women should not impact one's confidence in X, and then I explain why those conditions don't plausibly obtain in the range of cases epistemic allies care about. The argument has a few steps. I first explain the range of circumstances where encountering many intelligent women would not rationally require Solomon to give up his belief in X (Sect. 4.1). I then explain the range of circumstances where encountering many intelligent women would rationally require Solomon to give up his belief in X on the assumption that X is to be understood as a disjunctive statement about chances (Sect. 4.2). I next explain why X needn't be understood as a disjunctive statement about chances for one's belief in X to be undermined by encounters with intelligent women (Sect. 4.3). I then explain why epistemic allies are not left with a problem involving normative supervenience principles (Sect. 4.4).

The lite form of Bayesianism to follow consists of the following assumptions. First, agents are required to have probabilistically coherent credences and update by conditionalization just upon learning new information.⁸ Second, if one has rational, probabilistically coherent credences and just upon learning new information one updates on that new information by conditionalization, then one's resulting credences are rational. Third, agents should, other things being equal, calibrate their credences to the chances (discussed below). These three broadly Bayesian assumptions are consistent with the idea that rationality is not merely a matter of probabilistic coherence. I'm deeply suspicious, along with many other epistemologists, that an agent could simply have a probabilistically coherent set of attitudes and thereby count as fully rational in the usual sense that epistemologists have cared about. Those who share this suspicion will be sympathetic with my arguments below that there are some probabilistically

⁷ Since I will resolve this problem we can set aside questions about just how the intuitive advice to "evaluate like cases alike" should be precisely formulated. I'm grateful to Begby (2013) for drawing my attention to the need to treat sexist/racist generics and non-sexist/non-racist generics alike, and thus the puzzle allies face in this regard.

⁸ Under the assumption that *learning that P* is distinct from having an *experience as of P*, this rule for updating is consistent with allowing agents to update in other ways (Miller 2016). The 'just' in 'just upon learning new information' is meant to exclude complications where an agent might simultaneously learn that P and have an experience that would somehow conflict with updating on P.

coherent credences that are not rational to hold. While this is very much not in the spirit of extreme subjective Bayesianism, even extreme subjective Bayesians will be able to draw helpful lessons from this project.

There is one last thing worth noting before diving into the formal epistemology. It may seem strange to some that I introduce a formal theory of rationality to help address questions about the (ir)rationality of sexist and racist beliefs. After all, it's well known that cognitively limited agents such as ourselves are incapable of living up to the ideals of formal theories of rationality, and yet cognitively limited agents can be rational in some sense that is compatible with falling short of the ideals of formal theories. So what's the point in highlighting the fact that sexists and racists also fall short of these ideals if *everyone* falls short of the ideals anyway?

Here's one kind of answer to this question. Even if ideal rationality is unattainable for actual agents, rationality comes in degrees. That is, there are different degrees to which one can be less than ideally rational, and some of these degrees fall further short of ideal rationality than others. So even though ordinary thinkers are typically less than ideally rational, they can also be *sufficiently rational* for the purposes of our ordinary evaluations of each other. But ordinary thinkers can also be *so far* from ideal rationality in what they believe that they are *not* sufficiently rational for the purposes of our ordinary evaluations of each other. As it will turn out, under the conditions to be discussed, people who retain their sexist and racist beliefs in the face of many counterinstances will not only be irrational by the ideal standards, but they will also fall so far from these ideal standards that epistemic allies are right to think that these people are nowhere near sufficiently rational for the purposes of our ordinary evaluations of each other.⁹

4.1 Conditions of probabilistic irrelevance

There are two familiar Bayesian conditions that explain how it's possible for Solomon's background beliefs (priors) to be such that learning that there are many intelligent women does not rationally require him to reduce his credence in X. The first irrelevance condition concerns *certainty*.¹⁰ Suppose Solomon is certain that X is true, i.e. $p(X) = 1$. If so, then *no* future encounter with an intelligent woman will, by conditionalization, reduce his credence in X. For it's a theorem of Bayesianism that no future evidence can undermine a certainty if one is just updating on coherent credences by conditionalization. Thus, if $p(X) = 1$ at any time, then $p(X) = 1$ *at all future times no matter what new evidence arrives*.¹¹

⁹ For further discussion of closely related responses to this issue see Christensen (2007), Wedgwood (2014), Titelbaum (forthcoming), and Staffel (MS).

¹⁰ I'm grateful to Alan Hinkle for drawing my attention to some of the implications of this condition in the present discussion.

¹¹ If an agent can rationally update in ways other than conditionalization (Miller 2016), then being certain that P can be undermined by one's future experience. But I don't want to defend epistemic allies on these grounds here.

To appreciate the intuitiveness of this condition suppose you knew *with certainty* that you had a fair coin. Suppose after 100 tosses it landed heads 90 times. This is an extremely improbable thing for a fair coin to do. So what should you think about the fairness of the coin? Well, if you were truly certain that it was a fair coin to begin with, you should just think that the improbable has happened—improbable things do happen after all. The only way it makes sense for you to revise your belief in the fairness of the coin after so many tosses that land heads is if you had some pre-existing doubt about the coin's fairness. Similarly for Solomon: if he's truly rationally certain of X, he shouldn't revise his credence in X upon learning of many intelligent women. Rather, he should just regard it as an improbable occurrence.

What might epistemic allies think of this first condition of irrelevance? Should epistemic allies be worried that Solomon might be rationally certain of X? Notice that it is not clear that Solomon's epistemic position is *so* good that he can be *rationally* certain of X and thus rationally have absolutely no doubt at all about X's truth. For we should assume that Solomon is like us in having a somewhat vague idea of what intelligence is—at least if we want to keep Solomon's case somewhat realistic. But having a vague notion of the trait to be identified will surely generate some false positives and false negatives. So some unintelligent men will be classified as intelligent, and some intelligent women will be classified as unintelligent. Moreover, the kind of mental averaging one implicitly engages in when seeking to compare men's intelligence to women's intelligence is surely very imprecise. So while Solomon's evidence in favor of X is strong, it's clearly fallible in a way that warrants an attitude that falls short of certainty.¹² The lesson is that effectively resisting the judgment of epistemic allies that Solomon is irrational in continuing to believe X by appealing to the idea that he is (or should be) certain or near-certain of X will require (i) an explanation for how Solomon could be *rationally* certain or near-certain of X in his circumstances despite the facts just noted, and (ii) that the circumstances in which Solomon can be rationally certain or near-certain of X are quite normal—for recall that epistemic allies make a claim about normal circumstances. While it's implausible that Solomon could be rationally certain of X, there's no parallel problem in thinking he could be rationally *highly confident* in X (e.g. having a credence around .95). But as we will see below high confidence in X that falls short of certainty (or near-certainty) will not enable Solomon to rationally retain his belief in X upon encountering many intelligent women in a very wide range of cases.

The second irrelevance condition that could prevent Solomon from having to reject his belief in X upon encountering many intelligent women is this. Suppose Solomon's priors are such that X is probabilistically independent of:

(E) The next intelligent person encountered is a woman.

That is, $p(X|E) = p(X)$. If this condition holds, then learning E will not impact Solomon's credence in X. For upon learning E conditionalization tells us to set our

¹² This point is incompatible with extreme subjective Bayesianism. But even if extreme subjective Bayesianism were the correct account of the kind of rationality epistemologists tend to care about, I suspect that in most practical cases people who believe X will have *some* degree of doubt in X. And that's enough to enable extreme subjective Bayesians to make use of the remarks to follow.

updated credence in X , $p^+(X)$, to our prior conditional credence in X given E . That is, $p^+(X) = p(X|E)$. Thus, if $p(X|E) = p(X)$, it will follow that $p^+(X) = p(X)$. From this it's easy to see that learning E will not affect one's updated credence in X . So if X is probabilistically independent of E conditionalizing upon learning E will not affect Solomon's credence in X *even if* Solomon is not certain of X .

But there is a concern here stemming from Bayesian confirmation theory. Since we can very easily revise this second case, I will just quickly state the concern. Bayesian confirmation theory says that whenever you have two hypotheses that predict an outcome, O , with different probabilities, learning O will more strongly support the hypothesis that more strongly predicted O .¹³ Notice that $\neg X$ seems to more strongly predict E than X does: for it seems more likely that the next intelligent person encountered will be a woman if it's false that women are less intelligent than men. Accordingly, it seems as if Solomon's priors should make X probabilistically dependent on E in the following way: $p(X|E) < p(X)$. So if Solomon were to learn E and nothing else relevant, he should decrease his credence in X .

It is plausible that this last thought is what underlies Arpaly and Fricker's view about the marked irrationality of Solomon above. For, arguably, if E is negatively probabilistically relevant to X —i.e. $p(X|E) < p(X)$ —then encountering intelligent women should drive one's credence in X down. And, arguably, repeatedly encountering intelligent women should drive one's credence in X down dramatically.

The trouble is that there is a way of recovering the idea that encountering intelligent women is probabilistically independent of X that doesn't seem to face immediate problems. For while X may not be probabilistically independent of E , it may be probabilistically independent of E *and further information*. For example, Begby (2013, p. 94) implicitly makes this point when he writes that:

[Suppose Solomon encounters many intelligent women at university.] Does this last batch of evidence require him to reassess his prejudice against women? We can imagine Solomon reasoning as follows: while women are on average less intelligent than men, some women are more intelligent than others, and some are indeed so intelligent as to be the functional peers of all but the most brilliant men. Where might one hope to meet some of these women? At the university, of course!

The idea, then, is that Solomon needn't revise his credence in X in the envisioned case. For even though generally (E) encountering intelligent women is evidence against X , Solomon's credence in X is probabilistically independent of:

($E \& U$) The next intelligent person encountered is a woman, *and that person is encountered at a university*.

That is, while $p(X|E) < p(X)$, Solomon's credences are such that $p(X|E \& U) = p(X)$. Put differently, U screens off X from E . This point should be unobjectionable, at least if we add the intended assumption that only intelligent people are admitted to Solomon's university and that his university admits a roughly equal ratio of men to women.

¹³ Hawthorne (1994) and Earman (1992). For introductory discussions see Skyrms (1975) and Strevens (MS).

4.2 The irrelevance of the second irrelevance condition

Recall what epistemic allies think: they think that in normal circumstances sexist beliefs are undermined by counter-instances, and that enough counter-instances will require the rejection of the sexist belief. The second irrelevance condition just discussed explains why there is a range of normal circumstances where learning E will not impact one's credence in X . But whether or not this observation threatens epistemic allies depends on whether or not sexists who believe X will *always* (or at least *very often*) have some further bit of information that screens off X from E . This is crucial. For epistemic allies are not, nor should they be, of the view that learning E is *always* evidence against X . Indeed, they should agree with the fact that Solomon's credences are such that $p(X|E \& U) = p(X)$. But just as it's possible to learn E in connection with further information that makes E probabilistically irrelevant to X , *it is possible to learn E in the absence of such information*. In other words, $p(X|E \& U) = p(X)$ is compatible with $p(X|E \& \neg U_i) < p(X)$.

To illustrate how things could play out in a way that undermines the rationality of Solomon's belief in X take the following case:

Solomon regularly uses public transportation to get around New City and he has no reason to think that intelligent people are for some reason more likely to use public transit than non-intelligent people—everyone needs to get around town after all. One day while Solomon is using public transit he overhears someone behind him having a very intelligent conversation on the phone (i.e. the kind of conversation only very intelligent people could have). After a bit of eavesdropping he looks behind and sees that it's a woman talking.

What should happen to Solomon's credence in X at this point? The lite Bayesian answer is that it should decrease. That is, Solomon should become less confident that women are less intelligent than men, and if he encounters enough intelligent women in the right circumstances he ought to reject X entirely.

To help see why this is take the following theses relating the property of being a man to intelligence:

- (H9) The chances that the next intelligent person encountered is a man are .9.
- (H8) The chances that the next intelligent person encountered is a man are .8.
- (H7) The chances that the next intelligent person encountered is a man are .7.
- (H6) The chances that the next intelligent person encountered is a man are .6.
- (H5) The chances that the next intelligent person encountered is a man are .5.

Treat these as specifying the objective chances that a randomly selected intelligent person is a man. It's very widely thought that calibrating one's credences to the chances is, other things being equal, *rationally required*—though there's some dispute about just when “other things are equal” (Lewis 1994; Pettigrew 2012). I obviously cannot lay out the case for requirements along these lines here. Since there's nothing exceptional about Solomon's case as envisaged by Arpaly, Fricker, or Begby we'll assume that Solomon is a rational agent who calibrates his credences to the chances in the present

situation.¹⁴ That is, $p(\text{the next intelligent person randomly encountered is a man} | H_n)$ is equal to the chances specified by H_n . So, for example, $p(\text{the next intelligent person randomly encountered is a man} | H_9) = .9$. For convenience let's also assume that $p(\text{the next intelligent person randomly encountered is } woman | H_n) = 1 - p(\text{the next intelligent person randomly encountered is man} | H_n)$. Thus, $p(\text{the next intelligent person randomly encountered is } woman | H_9) = .1$. This is a substantive assumption that can impact the rate of convergence on H_5 , but the convergence on H_5 under the specified assumptions will still occur under different ways of specifying these conditional probabilities.¹⁵ So the philosophical points I aim to make in this paper will remain.

There is now a puzzle concerning how we should think about X in this formal framework, and what might constitute evidence against it. For, as noted above, the generic proposition *that women are less intelligent than men* is not a proposition that entails that *all*, or *most*, or some exact *proportion* of women are less intelligent than men; nor does it entail these claims when limited to *normal circumstances*; nor does it entail a related claim about *averages*. Fortunately, we needn't get bogged down in the vexing issues surrounding the truth conditions for generics. Notice that H_9 – H_6 are each a way of holding that men are *likely* to be more intelligent than women. So let's *temporarily treat* X as the disjunction of these claims, i.e. $X = H_9 \vee H_8 \vee H_7 \vee H_6$. The advantage of this is that—as I'll explain—it's clear in the Bayesian framework outlined what constitutes counterevidence to $H_9 \vee H_8 \vee H_7 \vee H_6$ and it's also clear when that counterevidence can rationally require one to reject $H_9 \vee H_8 \vee H_7 \vee H_6$ and accept H_5 . In the subsequent section I'll explain why being rationally required to accept H_5 would rationally require one to reject the generic claim that women are less intelligent than men.

Since Solomon is to have a high prior credence in the generic claim that women are less intelligent than men, let's set his credence in $X (= H_9 \vee H_8 \vee H_7 \vee H_6)$ to .95. Let's also suppose that Solomon has a meager .05 credence in H_5 and thus has a credence of 0 in all other possible hypotheses concerning the chances of encountering an intelligent man. The fact that we're modeling Solomon's credences in such a way that he assigns a 0 to so many alternative hypotheses is a realistic over-simplification: it's realistic because actual agents don't have credences for the continuum many hypotheses about the chances that exist; it's an over-simplification because agents will tend to have some non-zero credence in many more hypotheses than the five under consideration here. This doesn't impact the general lessons I wish to draw since the kind of convergence on H_5 that I'm illustrating would remain even if we were to consider an agent with non-zero credences in many more chance hypotheses.¹⁶

¹⁴ Extreme subjective Bayesians reject this, requiring only that agents have coherent credences and conditionalize. But even so, agents have a tendency to be sensitive to their views about the chances even if they're not required to be. So extreme subjective Bayesians can still make some use of the justification of epistemic allies to follow. Though it will be somewhat more limited. For they will have to hold that if Solomon chooses not to calibrate his credences to the chances, then he *might* well be rational in retaining his belief in X -depending on his priors.

¹⁵ Hawthorne (1994) and Earman (1992). For introductory discussions see Skyrms (1975) and Strevens (MS).

¹⁶ Again, see Hawthorne (1994) and Earman (1992). For introductory discussions see Skyrms (1975) and Strevens (MS).

Next, to see how randomly encountering intelligent women should impact Solomon's credence in X we have to specify credences for H9, H8, H7, and H6 that sum to .95. So suppose that prior to encountering any intelligent woman in New City, Solomon's evidence makes it rational for him to put the following credence in H9-H5:

$$\begin{aligned} p(H9) &= .35 \\ p(H8) &= .3 \\ p(H7) &= .2 \\ p(H6) &= .1 \\ p(H5) &= .05 \end{aligned}$$

The graph below illustrates a number of alternative prior distributions that strongly favor X, but it will help to focus on just this one. All the main points I make with this distribution could be made with a very wide range of other distributions, as the chart below illustrates.

Let E1 = the first random recognizably intelligent person encountered is a woman.¹⁷ Given Solomon's initial distribution and what each of H9-H5 imply about the chances of randomly encountering an intelligent woman, Solomon's credence in E1 will be .22, i.e. $p(E1) = .22$.¹⁸

Now, recall that while Solomon was traveling on public transportation the first random intelligent person he encountered was a woman. If Solomon conditionalizes on this information (and gains no further relevant information), then his new credence distribution, $p^+(\cdot)$, in H9-H5 will be roughly:¹⁹

$$\begin{aligned} p^+(H9) &\approx .16 \\ p^+(H8) &\approx .27 \\ p^+(H7) &\approx .27 \\ p^+(H6) &\approx .18 \\ p^+(H5) &\approx .11 \end{aligned}$$

There are a few things to observe at this point. First, Solomon's credence in H5 has more than doubled from .05 to .11. Second, even though Solomon's credence in X is high he has become noticeably less confident in X since his credence in it has

¹⁷ ...where 'random' is to imply not at the university or any other circumstance that screens for intelligence and an equal ratio of men to women.

¹⁸ That is, we reach $p(E1) = .22$ with an application of the law of total probability:

$$p(E1) = p(E1|H9)p(H9) + p(E1|H8)p(H8) + p(E1|H7)p(H7) + p(E1|H6)p(H6) + p(E1|H5)p(H5).$$

Because H9-H5 specify the chances that a randomly encountered intelligent person is a man, it follows that the chances for encountering an intelligent woman are 1- the chances of randomly encountering a man as specified by H9-H5. Since Solomon has calibrated his credences to the chances, $p(E1) = .1 \times .35 + .2 \times .3 + .3 \times .2 + .4 \times .1 + .5 \times .05 = .22$.

¹⁹ These are easy to calculate. For by conditionalization $p^+(Hn) = p(Hn|E1) = p(E1|Hn)p(Hn)/p(E1)$, and since Solomon has set his conditional credences in E1 given Hn to the chances, $p(E1|Hn) =$ the chances that E1 obtains given Hn. Solomon's priors for H9-H5 are given above, and we've already determined that his credence in E1 is .22.

decreased from .95 to roughly .89.²⁰ So Solomon has become noticeably less confident in X than he was before. And assuming that updating by conditionalization on rational credences results in rational credences, it follows that Solomon's reduced credence in X is rational.

One might be tempted to think that this near-trivial reduction of credence in X justifies the idea that it is rational for Solomon to continue to have a high credence in X upon encountering many more intelligent women. But caution must be urged here. *For randomly encountering even a fairly small number of intelligent women can, under certain conditions, be enough to undermine the rationality of Solomon's continued high credence in X.* For example, suppose that from this point forward as Solomon goes about fairly common activities (reading at coffee shops, dancing, exercising, attending church, shopping, reading comment sections on-line, etc.) and when he encounters recognizably intelligent people he encounters just as many intelligent women as men—as one would expect given that H5 is true. If Solomon encounters only 42 recognizably intelligent people half of which are women, the rational credence he should have in H5 after conditionalizing will be roughly .512 and thus his rational credence in X will be roughly .488.²¹ Accordingly, if Solomon conditionalizes on his evidence as Bayesians advise, his credence in X will be so low that he no longer counts as believing X. And if Solomon randomly encounters just 152 recognizably intelligent people half of which are women, his credence in H5 after conditionalizing will be greater than .9 and his credence in X will be less than .1.²² In other words, his credence in H5 and X are virtually the opposite of what they were when he arrived in New City.

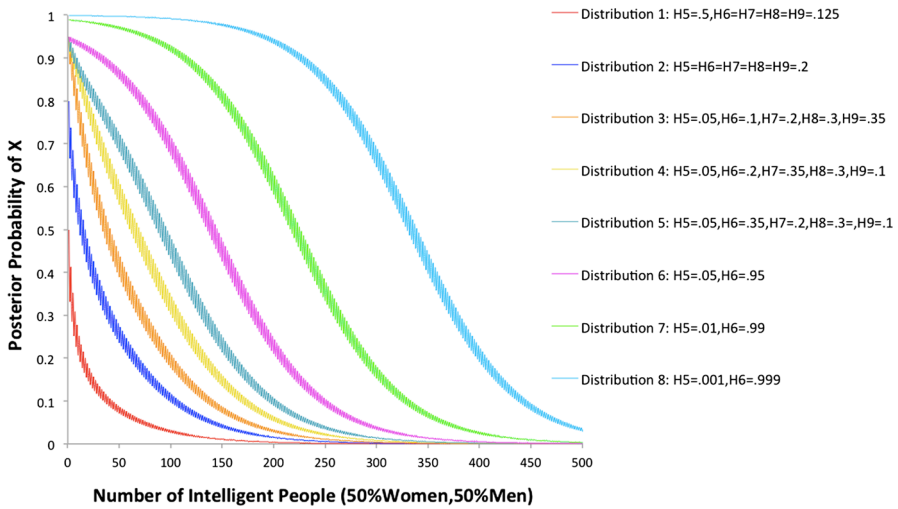
Just how many random encounters with intelligent people it takes for Solomon to reach this position depends heavily on his priors. Here's a graph that illustrates how a few additional prior distributions will converge on H5:

²⁰ Since H9–H6 are mutually exclusive and X is the disjunction of them, $p(X) = p(H9) + p(H8) + p(H7) + p(H6)$.

²¹ For if Solomon has randomly encountered 42 recognizably intelligent individuals—half of which are men and half of which are women—and if he conditionalizes on each encounter, then his credence in H5 after learning E_{42} will be $p(H5|E_{42}) = p(E_{42}|H5)p(H5)/p(E_{42}) \approx .512$. And his credence in X will be $p(X|E_{42}) = 1 - p(H5|E_{42}) \approx .488$.

These values are easy to calculate. Since H5 implies that the chances that a randomly selected intelligent person is a woman or a man is .5, it follows that $p(E_{42}|H5) = .5^{42}$. It's already part of Solomon's credence distribution that $p(H5) = .05$. And $p(E_{42}) = p(E_{42}|H5)p(H5) + p(E_{42}|H9)p(H9) + p(E_{42}|H8)p(H8) + p(E_{42}|H7)p(H7) + p(E_{42}|H6)p(H6)$. Which is: $(.5^{42} \times .5^{21} \times .05) + (.9^{21} \times .1^{21} \times .35) + (.8^{21} \times .2^{21} \times .3) + (.7^{21} \times .3^{21} \times .2) + (.6^{21} \times .4^{21} \times .1)$. Thus, $p(H5|E_{42}) \approx .512$. Thus, $p(X|E_{42}) = 1 - p(H5|E_{42}) \approx .488$.

²² To calculate this just take the formulas from the previous footnote and replace the relevant values.



Each prior distribution’s number corresponds to its numerical order from left to right. So Distribution 1 is represented by the first line on the left, Distribution 2 is represented by the second line from the left, and so forth. For each even number n along the horizontal axis, the number of women encountered is half of n . So each line plots the course of an agent whose encounters with intelligent people always alternate between men and women. While this assumption impacts the exact path of each line, it doesn’t impact the relevant outcome because evidence commutes for agents who update by conditionalization. Thus when assessing the credence of agents after having met, say, 100 intelligent people, the exact order in which they encountered intelligent men and women doesn’t matter so long as about 50% of them are women.

There are many things to observe about this graph.²³ But the main point for present purposes is that each distribution converges on H5 as one encounters more and more intelligent people where about 50% of them are women. What differs between these distributions is just how quickly this convergence occurs. So while convergence can occur quite quickly (see Distribution 1), it can also take quite a bit longer (see Distribution 8). But so long as Solomon’s credence in X is not 1, he will eventually end up assigning a trivial credence to X and assigning a very high credence to H5 after enough random encounters with intelligent individuals where about 50% are women. How many is “enough” will depend on the priors.²⁴

²³ For example, it makes it seem as if Solomon, if he’s rational, is bound to converge on H5 and relatively quickly. That’s a function of our model’s setup. Were Solomon to assign a non-trivial credence to more chance hypotheses, he’d converge so quickly only on a clump of hypotheses surrounding H5 upon encountering many intelligent women about half of which are women—e.g. hypotheses that say the chances that the next encountered intelligent person is a man are either, say, .48 or .49 or .5 or .51 or .52. But which clump of hypotheses near .5 and how quickly is all a function of the priors. I’ll say a bit more about why this doesn’t compromise some of the epistemic lessons I draw in footnote 29.

²⁴ Notice also that Solomon’s sexism could be *more extreme* in the sense that he might be very confident that women are objectively far less likely to be as intelligent as men (e.g. he might assign a much higher credence to H7–H9 than to H6). On the other hand, Solomon’s sexism can be *less extreme* in the sense that he is very confident that women are just somewhat less likely to be as intelligent as men (e.g. he might

At this point some readers might have a practical worry. For it can look like rationally overcoming sexist beliefs can take a long time and a lot of effort; it really would take a bit of time to randomly encounter, say, 500 recognizably intelligent people. But talk of “encountering” people here has just been an expressive device stemming from reflection on Solomon’s case. The general point is that one only needs to *learn* of such people, and one can learn of such people quite quickly through testimony from others. For example, Solomon’s new friends who have grown up in New City could just tell Solomon that their encounters with men and women have shown that men are no more intelligent than women. Or perhaps in university Solomon may end up reading about the national studies that show that for decades women’s school grades have been increasing, and have recently *surpassed* men in overall school performance.²⁵ There is much more to be said about the ways in which Solomon could gain relevant testimonial evidence. But that drags us into Q3 territory, and there’s no use addressing that question unless we can be confident that Q1 has a positive answer.²⁶

Footnote 24 continued

assign a low to credence to H7-H9 and a very high credence in H6). This is approximately the difference between Distribution 6 and Distribution 3 on the graph above, where Distribution 3 is the original, more extreme form of sexism we assigned Solomon at the beginning. But from the graph it’s easy to see that being less of a sexist in this sense entails that it will take *many more* encounters with intelligent people for one to assign a credence of less than .5 to H9vH8vH7vH6– it will take approximately 150 encounters where about 50% are women. This is about 3.5 times more encounters than the 42 it would have taken Solomon to reach the same credence in H9vH8vH7vH6 on his original, more sexist distribution. So, surprisingly, being more sexist (i.e. assigning a higher credence to H7-H9) entails that one’s sexism can be *more easily corrected* than if one is less of a sexist. So the less of sexist one is the less lucky one is in the sense that one’s sexism will be harder to correct. This fact is also perspicuous in the case of Distributions 7 and 8.

²⁵ This is also true of women at the high school level. That is particularly of note for someone like Solomon since the vast majority of high schools, unlike most universities, do not have admission requirements that exclude unintelligent people. So were someone like Solomon to learn of these studies he’d not be able to appeal to the screening off condition that he did above. For some discussion of the closing of the performance gap in math and science between men and women in high school and college see Goldin et al. (2006).

²⁶ One referee pointed out that the prejudice and stereotyping literature (Fiske 1998) strongly suggests that stereotypes and prejudices persist because these attitudes are highly resistant to counter-evidence. Plausibly, this is because subjects have something like further information in a variety of these circumstances which screens-off the relevance of the counter-evidence.

Whether or not this observation about human psychology impacts the epistemological points I’m making depends on whether the actual resistance we witness to revising beliefs in claims like X in light of counter-evidence E is *rational*. This in turn depends on (i) whether or not one’s belief in X was rational to *begin with*, and on (ii) whether or not one *has*–in an epistemically relevant sense–further information that screens-off E from X. But regarding (i), it’s not clear that people tend to have nearly the sort of strong evidential support for high credence in X to begin with (recall how carefully constructed Solomon’s case was). And regarding (ii), it is not obvious that people ‘have’–in a relevant sense–further information that allows for screening-off. For in order to *have* further information in a sense that’s relevant to preserving the *epistemic rationality* of one’s belief in the face of counter-evidence, it is not enough that one *merely believe* further propositions. They have to *rationally* believe them, or *be in a position rationally* believe them, or stand in some other positive epistemic relation to that further information. So the ‘have’ in “having further information” that’s relevant in thinking about the epistemology of sexist/tracist beliefs is itself hiding a normative element. And observing that as a matter of fact people *act as if they have* further information in an epistemically relevant sense is not enough to show that they *do* have such information.

4.3 Does it matter that $X \neq H9vH8vH7vH6$?

Some readers might have the following theoretical worry. For it can seem that in treating the generic statement expressed by X as the disjunction of $H9-H6$ I've falsely assumed that the generic proposition *that women are less intelligent than men* expresses the same thought as $H9vH8vH7vH6$. But recall that I treated X as $H9vH8vH7vH6$ only for convenience. Still some may worry that the fact that $X \neq H9vH8vH7vH6$ must limit the epistemic significance of the formal results from the previous section.²⁷

But my main point concerns the *irrationality* of failing to significantly reduce one's credence in X upon randomly encountering many intelligent women, and this point doesn't depend on X expressing the same thought as $H9vH8vH7vH6$. Rather, what my main point depends on is that the *rationality* of having a high credence in some generics (like X) *can* depend on the *rationality* of not having a high credence in certain claims about probabilities (like $H5$). Thus, whenever this dependence relation obtains it will follow that: if one's evidence (and background beliefs) rationally require one to assign a high credence to the relevant claim about the probabilities (like $H5$), then one cannot rationally assign a high credence to the corresponding generic (like X). Put differently, the rationality of believing $H5$ can defeat the rationality of believing X .

There are two independent explanations for thinking such a dependence/defeat relation obtains between $H5$ and X in Solomon's case. First, consider how Solomon came to rationally put a high credence in X in the first place. If we reconstruct the kind of reasoning he implicitly relied on, we get something like this:

Premise 1. Members of my community and I have encountered many intelligent people. Of all the intelligent people encountered very few of them have been women.

Premise 2. I have no reason to think my past encounters with intelligent people, or the past encounters of others in my community, does not accurately represent the intelligence of men and women.²⁸

Conclusion. Therefore, women are less intelligent than men ($=X$).

This argument for X is a fine example of non-deductive reasoning to a conclusion in a generic proposition.

So now let's ask what bearing Solomon's future random encounters with many intelligent women can have on *this* bit of reasoning that he's relying on. Obviously,

²⁷ Leslie (2008), for example, argues against associating generics with probabilities too closely. In his book on the psychology of stereotypes Schneider (2004, p. 206) also warns against treating agents as if their stereotype beliefs (=beliefs with generic content) involved probabilistic content, even though doing so enables efficient mathematical representations. I'm grateful to Begby (2013) for drawing my attention to this.

²⁸ Premise 2 didn't explicitly figure into Solomon's reasoning. But we should grant it. For without Premise 2 Solomon's reasoning to the conclusion would be quite poor. For, other things being equal, encounters with many F s that are G doesn't give one reason to think *most* or *all* or *on average* F s are G if one has some reason think their encounters fail to accurately represent the F s. The fact that Solomon, as imagined in the narrative, is sensitive to this kind of concern is evident in the fact that he takes U to screen off X from E (see above).

having many such encounters is consistent with the conclusion. But the rationality of believing the conclusion on the basis of this argument depends on more than that. The premises must be rational for Solomon to believe. But suppose Solomon randomly encounters 42 intelligent people where about 50% are women. In this case, it becomes rational for Solomon to put a credence of just over .5 in H5 and a credence just below .5 in H9vH8vH7vH6. But if it's rational to put a credence of .5 in H5, it would seem irrational for him to have a high credence in Premise 2. For if it's true that women are just as likely to be as intelligent as men—as H5 implies—then encountering a disproportionately large number of intelligent men will obviously *not* be a sample that accurately represents the intelligence of men and women. So if it's rational for Solomon to be about .5 confident in H5, then he should be about equally confident that Premise 2 is false. But if Solomon's credence in the falsehood of Premise 2 is around .5, then his credence in the conclusion should certainly not be any higher than .5. So in these circumstances it's irrational for Solomon to have a high credence in X on the basis of the inductive argument above. The fact that X does not express the same thought as H9vH8vH7vH6 doesn't change this.²⁹

Here's a second reason for thinking the rationality of believing H5 can defeat the rationality of believing X. Consider the following conjunction:

(Y) Men are not likely to be more intelligent than women, but men are more intelligent than women.

It just seems odd that someone would explicitly assent to this conjunction. It sounds about as bad as believing "It's raining, but it's irrational to believe it." For in both cases it seems like rationally believing the first conjunct would make it irrational to believe the second conjunct. The general issue I'm drawing attention to is a tension between *comparative generic* claims to the effect that *Xs F more (are more F) than Ys* and claims about the objective chances to the effect that *Xs are not likely to F more (be more F) than Ys*.

By contrast, consider the non-comparative generic: ducks lay eggs. Notice that it's not objectively likely that a randomly selected duck lays eggs (half of ducks are male and some females are infertile). But it's still true that ducks lay eggs. So the kind of tension we see in Y is not present in the thought that:

Ducks lay eggs, but it's not likely that a duck lays eggs.

²⁹ When considering further chance hypotheses convergence on H5 will take longer. Does that imply that in such cases an agent like would lack a defeater for Premise 2? Not necessarily, for suppose Solomon did not have a .5 credence in H5, but the following disjunctive chance claim:

(NearH5) The chances that the next intelligent person encountered are a man are either .48 or .49 or .5 or .51 or .52.

Suppose his credence in each disjunct here is spread evenly, or near enough. (That is, he's *not*, say, .49 confident in the hypothesis that says the chances are .52, and thus he *does not* have a high credence in a disjunctive claim that say that one's more likely to encounter an intelligent man than a woman.) In such a case where his .5 credence is spread evenly enough around the cluster of hypotheses specified in (NearH5), it's still intuitive to think that Solomon would continue to have significant reason to think Premise 2 is false and so should not be highly confident in Premise 2. Accordingly, his credence in the conclusion, X, should not be high either.

However, consider the comparative generic claim that *platypi lay more eggs than mallards*. Suppose you believed this because you were told it by someone you took to be in the know about such things. But then suppose you learned that mallards can lay a dozen eggs multiple times a year, while platypi can lay only up to about 3 eggs once a year. So it's false that (G) *platypi lay more eggs than mallards*. But you didn't have to learn the specifics of their egg-laying potentials to acquire evidence that G is false. It would have been enough if you learned the probabilistic fact that (O) platypi are *not likely* to lay more eggs than mallards. Thus, if your evidence makes it rational for you to believe O, it would be irrational for you to believe G even though O and G express different thoughts.³⁰ The same kind of tension that exists between O and G is present between X and H5. So if Solomon's evidence makes it rational for him to have a high credence in H5 (and thus irrational to have a high credence in H9vH8vH7vH6) then it would be irrational for him to believe X even though H9vH8vH7vH6 and X express different thoughts.

4.4 On treating like cases alike

So where do all these points leave epistemic allies? Recall, our concern for epistemic allies was that the generic content of sexist beliefs together with background beliefs seemed to make sexist beliefs rationally resilient in the face of counterevidence. But above we've shown that this resilience is limited. For if H5 is true then sexist thinkers are likely to be in an overall epistemic position where their random encounters with many intelligent women would rationally require them to reject X. This is true despite the fact that X is a generic proposition and the fact that one may have priors that would enable one to explain away a certain subset of their encounters with intelligent women.

Recall that the final problem facing epistemic allies was that they seemed committed to not evaluating like cases alike, and thus seemed to run afoul of familiar thoughts about the supervenience of the normative on the non-normative. For on the one hand they wanted to say that Solomon couldn't rationally explain away his counterevidence to the claim that women are less intelligent than men. But on the other hand they wanted (or should want) to maintain that we could rationally explain away our counterevidence to the claim that dogs have longer tails than chipmunks. Given the Bayesian arguments above for thinking that rationality could quite easily require Solomon to reject X and thus not rationally explain away his counterevidence, one might worry that the same arguments could be used to argue that encountering many tailless dogs would require us to revise our belief about dogs rather than explain away their relevance to that belief.

Happily, this worry is ill-founded for it trades on an implicit ambiguity that should now be apparent. Recall that Solomon's convergence on H5 occurred on three assumptions: (i) Solomon was neither certain of X nor near certain of X; (ii) Solomon encountered many intelligent people where a sufficient percent of them (around 50%)

³⁰ One way of explaining this tension is to argue that generics (or at least comparative generics) have probabilistic truth conditions (cf. Cohen 1999). Another way of explaining the rational tension is to argue that comparative generics, like G, have non-probabilistic truth conditions; but that claims about the objective probabilities, like O, can be evidence that those non-probabilistic truth conditions don't obtain. But we needn't settle on an explanation of the tension to recognize the tension.

were women; and (iii) Solomon had no additional defeating information that functioned like his information about university admissions. Now *if* conditions parallel to (i)–(iii) also held with respect to the belief that dogs have longer tails than chipmunks, then the epistemic ally should say just what they said about Solomon: in such conditions one is rationally required to give up their belief.

So why might we who are epistemic allies correctly be inclined to reject the idea rationality would have us cease believing that dogs have longer tails than chipmunks if we were to encounter many tailless dogs? I believe it's owed to a lack of precision in the question that invites us to fill-in certain details in a way that makes it rational to retain the belief that dogs have longer tails than chipmunks.

For example, it's not implausible that people are actually *certain* or *near-certain* that dogs have longer tails than chipmunks. Introspectively, my level of confidence in this is right up there with claims like *other people exist* and *the earth is more than 5 minutes old*. That is, I am as certain of these things as I am certain of anything. Perhaps this certainty is less than fully rational, but certainty in these cases is not obviously irrational.³¹ Accordingly, when epistemic allies think about encountering tailless dogs they may very well be thinking about it in a context where one is certain or near-certain that dogs have longer tails than chipmunks. So our intuitive reflections on what would be rational to believe upon encountering many tailless dogs may be tracking cases where a counterpart to (i) fails to obtain.

Additionally, even if we were not certain that dogs have longer tails than chipmunks and we were to begin to encounter many tailless dogs, it's not clear that *a sufficiently large proportion of them* would be tailless. But unless we are considering circumstances where a sufficiently large proportion of them are tailless, one's prior confidence that dogs have longer tails than chipmunks will not diminish to the point where one would have to give up the belief. So the question is not phrased in a way that invites thinking about a situation where a counterpart to (ii) clearly obtains.

Lastly, when we think about encountering many tailless dogs, a natural way to imagine this possibility is to think about encountering tailless breeds of dogs at a dog-show or dogs who've been obviously injured (and thus are defective members of their kind).³² But then a condition like (iii) will not clearly obtain. For encountering tailless breeds at a dog show or obviously injured dogs is like Solomon encountering women at university—the relevance of the new information gets screened off. So, yet again, the question is not phrased in a way that invites thinking about a situation where a counterpart to (iii) clearly obtains.

³¹ Notice that having a tail is an easily discernible property of animals; we know more or less just what we're talking about when we are talking about having a tail; and we're very reliable at discerning among the animals we encounter which have tails—especially when it comes to dogs who have rather obvious tails. But our judgements about intelligence are not at all like this. While we have a vague idea of what intelligence is, it's not clear *exactly* what intelligence is; our tests for intelligence are certainly more fallible than our tests for having tails; and the kind of averaging we implicitly engage in thinking about men's intelligence versus women's intelligence is surely very imprecise. So while (I think) a reasonable case can be made for idea that we can be rationally certain (or quite close to it) that dogs have tails, there's no parallel case to be made for thinking anyone in even semi-realistic circumstances can be rationally certain of X.

³² Technically, tailless breeds of dogs typically just have naturally very short tails relative to their body size. We can ignore this empirical wrinkle in the example, however, and treat *having a tail* as *having a sufficiently long tail relative to one's body size*.

Accordingly, the fact that epistemic allies are inclined to issue different judgements between the case involving Solomon and cases involving dogs is, I think, explained by an underlying evidential asymmetry between the cases under consideration. But once the cases are fixed in a way that makes them truly parallel the motivation to issue differing judgements will (or at least should) vanish. This resolves the worry that epistemic allies will end up having to reject normative supervenience principles.

5 Concluding remarks

Above I treated the proposition *that women are less intelligent than men* as a generic claim about the kinds *men* and *women* generally. But some might worry that while epistemic allies might not have a problem with this claim, they may have a problem with the closely related claim *that the most intelligent women are less intelligent than the most intelligent men*. It would be an easy thing to consider Solomon in circumstances where he started out with strong evidence for this claim. But then he encounters many highly intelligent women after leaving for, say, a job at NASA where only highly intelligent people work. There is no problem facing epistemic allies and their judgement that Solomon would be rationally required to give up this belief *so long as* conditions parallel to (i)-(iii) obtain. In such conditions, rationality would require the rejection of the claim that the most intelligent women are less intelligent than the most intelligent men.

So there is no problem in theory for epistemic allies here. But there is a *prima facie* practical problem and it's worth saying a few words about it. Consider that our gauges for very high levels of intelligence typically involve indirect measures, like being in an occupation that effectively weeds-out individuals who are not highly intelligent. For example, being a CEO of a successful multi-million dollar company, working in advanced tech fields, being a scientific researcher who publishes in *Nature*, etc. It's extremely unlikely that one could do these things and fail to be highly intelligent. So these are very reliable indicators that an individual is highly intelligent. But if we look at the proportion of men to women in these areas and we find that women are outnumbered by men, it would *seem* that we would not only fail to gain reason to reject the claim *that the most intelligent women are less intelligent than the most intelligent men* but we would gain strong reason to endorse it.

But, once again, this is where considerations of screening-off kick in. For to the extent that we have reason to think that inegalitarian practices selectively advance highly intelligent men into these roles with a greater frequency than highly intelligent women, we have reason to think that these indirect measures of high levels of intelligence offer us a biased sample of the population of very highly intelligent people. It's like fishing with a net whose holes allow fish smaller than ten inches to easily escape; and then concluding that the lake has few fish smaller than ten inches and many more that are larger after catching many fish, the large majority of which are larger than ten inches. Having significant evidence that you're fishing with such a net makes drawing such a conclusion irrational. Similarly, we have significant evidence that inegalitarian practices do selectively advance highly intelligent men into positions that demand high intelligence more frequently than women. And having such evidence

to this effect makes it irrational for one to be highly confident *that the most intelligent women are less intelligent than the most intelligent men* upon learning that there are more men than women in roles that demand very high intelligence.³³

Recall our target question:

Q1. In normal circumstances is it irrational to maintain one's belief that women are less intelligent than men *upon learning of many intelligent women*?

Epistemic allies answer this in the affirmative, and so long as enough of the “normal circumstances” are circumstances where counterparts to (i)–(iii) obtain they will have no trouble maintaining this position. But we can just circumvent worries about what most normal circumstances are like by slightly re-characterizing the position of epistemic allies. Let's say that epistemic allies are those who think that: In normal circumstances *where conditions akin to (i)–(iii) obtain* with respect to some sexist or racist or similarly problematic belief, it is irrational to maintain that belief upon encountering many counter-instances.

Acknowledgements Endre Begby and Nicholas Fillion were instrumental in getting this paper off the ground. I'm indebted to their kindness in reading rough drafts and offering penetrating comments. Likewise the referees at Synthese were thorough and thoughtful in their remarks and the present paper has been much improved as a result. Many others contributed to this work, including: Dan Singer, Julia Staffel, Karen Frost-Arnold, Alan Hinkle, Rima Basu, Gregory Phelan, Joseph Shieber, Max Lewis, Grace Boey, Kai Draper, and audiences at Washington University in St. Louis, the University of Cologne, the University of Delaware, and the University of Connecticut.

References

- Allport, G. W. (1979). *The nature of prejudice*. New York: Perseus Books Publishing.
- Arpaly, N. (2003). *Unprincipled virtue: An inquiry into moral agency*. Oxford: Oxford University Press.
- Basu, R. (2018). Dissertation. University of Southern California.
- Begby, E. (2013). The epistemology of prejudice. *Thought*, 2, 90–99.
- Benton, M. (2016). Knowledge and evidence you should have had. *Episteme*, 13, 471–479.
- Christensen, D. (2007). Does Murphy's law apply in epistemology? Self-doubt and rational ideals. *Oxford Studies in Epistemology*, 2, 3–31.

³³ One referee asked whether we could “explain the irrationality of believing that *Blacks are more criminal than Whites* on strictly evidentialist grounds. Since Blacks are strikingly overrepresented among both perpetrators and victims of homicide, and are overrepresented at every level of the criminal justice system; while Whites, comparatively, are underrepresented.”

Like all beliefs, when it comes to evaluating the rationality of this racist belief a lot turns on a thinker's over all epistemic situation. That said, I'm *inclined* to think (of course a lot turns on the envisioned details) that the Bayesian explanation of why this racist belief tends to be epistemically irrational for thinkers in the envisioned circumstances will be similar to why it tends to be irrational to believe that *the most intelligent women are less intelligent than the most intelligent men*. Basically, we know (or at least have strong evidence in support of the idea) that in the United States black people are overrepresented in the indicated ways in part because of their differential treatment in socio-economic systems and the justice system. If one has this kind of further information, it will screen-off the relevance of the statistical overrepresentation of black people just as having information about the differential treatment of women screens-off the relevance of the underrepresentation of women in social roles that demand high intelligence. If someone lacks this further information about the differential treatment of black people, then it's information that they arguably *ought to have* and it's this that defeats the rationality of their belief that black people are more criminal than white people (cf. Q2 in footnote 2).

- Cohen, A. (1999). *Think generic: The meaning and use of generic sentences*. Stanford: Stanford University Center for the Study of Language and Information.
- Dovido, J., Glick, P., & Rudman, L. (2008). *On the nature of prejudice*. Oxford: Blackwell.
- Earman, J. (1992). *Bayes or bust? A critical examination of Bayesian confirmation theory*. Cambridge: MIT Press.
- Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (pp. 357–411). New York, NY: McGraw-Hill.
- Fricke, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford: Oxford University Press.
- Goldberg, S. (2016). On the epistemic significance of evidence you should have had. *Episteme*, 13, 449–470.
- Goldin, C., Katz, L., & Kuziemko, I. (2006). The homecoming of American College Women: The reversal of the college gender gap. *Journal of Economic Perspectives*, 20, 133–156.
- Hawthorne, J. (1994). On the nature of Bayesian convergence. *The Philosophy of Science Association*, 1, 241–249.
- Kunda, Z., & Oleson, K. C. (1995). Maintaining stereotypes in the face of disconfirmation: Constructing grounds for subtyping deviants. *Journal of Personality and Social Psychology*, 68, 565–579.
- Leslie, S. J. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, 21, 375–405.
- Leslie, S. (2008). Generics: Cognition and acquisition. *Philosophical Review*, 117, 1–47.
- Leslie, S. J. (2015). Hillary Clinton is the only man in the Obama Administration’: Dual character concepts, generics, and gender. *Analytic Philosophy*, 56, 111–141.
- Leslie, S. J., & Lerner, A. (2016). *Generic generalizations*. Stanford Encyclopedia of Philosophy, Winter 2016 Edition, Edward N. Zalta (Ed.). <https://plato.stanford.edu/archives/win2016/entries/generics/>.
- Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103, 473–490.
- Miller, B. T. (2016). How to be a Bayesian Dogmatist. *Australasian Journal of Philosophy*, 94(4), 766–780. <https://doi.org/10.1080/00048402.2016.1138233>.
- Neisser, U., Boodoo, G., Bouchard, T., Boykin, W., Brody, N., Ceci, S., et al. (1996). Intelligence: Knowns and unknowns. *American Psychologist*, 51, 77–101.
- Pace, M. (2011). The epistemic value of moral considerations: Justification, moral encroachment, and James will to believe. *Nous*, 45, 239–268.
- Pettigrew, R. (2012). Accuracy, chance, and the principal principle. *Philosophical Review*, 121, 241–275.
- Schneider, D. (2004). *The psychology of stereotyping*. New York: The Guilford Press.
- Skyrms, B. (1975). *Choice and chance*. Belmont: Dickinson Publishing Company.
- Staffel, J. MS. *Unsettled Thoughts*. University of Colorado, Boulder.
- Strevens, M. MS. *Notes on Bayesian confirmation theory*. New York University.
- Titelbaum, M. (Forthcoming). ‘Normative Modeling’ in methods in analytic philosophy: A contemporary reader. J. Horvath ed. Bloomsbury Academic Press.
- Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*, 45, 961–977.
- Wedgwood, R. (2014). Rationality as a virtue. *Analytic Philosophy*, 55, 319–338.