

#BelieveWomen and the Ethics of Belief

Renée Jorgensen Bolinger (bolinger@princeton.edu)¹

(penultimate draft, for inclusion in the *NOMOS* Fall 2021 volume on *Truth & Evidence*)

1. Preliminaries

For all its faults—and they are legion—Twitter is a great resource for social epistemology. Particularly over the last five years, the platform has provided a window into our collective practices surrounding testimony of sexual assault and harassment. As women reported the sexual misconduct of various high-profile men in the wake of the #metoo campaign, twitter preserved a public record of a common pattern: a woman accuses a man of misconduct. He denies the accusation, and in the subsequent discussion someone urges the importance of the presumption of innocence and due process: we are told to not ‘condemn him in the court of public opinion’ until the accuser produces “real evidence” sufficient to establish his guilt beyond reasonable doubt. As these patterns recurred on the day after Christine Blasey Ford’s testimony at the Kavanaugh confirmation hearings (September 27 2018), a full-page color ad appeared in the *New York Times*, saying simply ‘BELIEVE WOMEN’.² The phrase also began trending on twitter, as users urged us to #BelieveWomen. Predictably, the same general patterns recurred in response to the hashtag: Op-Eds appeared, worrying that believing women short-circuits due process, strips the accused of the presumption of innocence, or urges us to ignore the facts.³ Similar numbers of editorials sprang up in reply.

There are a cluster of difficult questions here, both interpretive and normative. Among them: how should we understand the social project of the ‘#BelieveWomen’ campaign?⁴ and what

¹ With special thanks to Richard Bradley, Kim Ferzan, Maegan Fairchild, Georgi Gardiner, Heather Young, and William Wells for comments on earlier drafts. I also benefitted from discussion with the participants at the NOMOS conference hosted by Princeton’s University Center for Human Values in Fall 2019, students in Zoë Johnson King’s Spring 2020 seminar on *the Social Turn in Metaphysics and Epistemology* at New York University, participants in the Princeton University Center for Human Values Laurence Rockefeller Research Seminar, members of the Rutgers Value Theory Reading Group in May 2020, and participants in the Social (Distance) Epistemology Series in June 2020, especially Rima Basu, Amy Flowerree, and Sanford Goldberg.

² The ad was sponsored by the dating app Bumble, which is distinctive in requiring women to make the first move, and is marketed as a way to at least somewhat reduce the amount of harassment women face when dating. <https://thehill.com/blogs/blog-briefing-room/news/408946-female-driven-dating-app-bumble-publishes-full-page-ad-in-the>

³ Bari Weiss rejects #BelieveWomen with the explanation “I believe that facts serve feminists far better than faith. That due process is better than mob rule. [...] What we owe all people, including women, is to listen to them and to respect them and to take them seriously. But we don’t owe anyone our unthinking belief.” (Weiss, Opinion, ‘The Limits of Believe All Women’ *New York Times* Nov 28 2017: <https://www.nytimes.com/2017/11/28/opinion/metoo-sexual-harassment-believe-women.html>)

⁴ One important question which I will not spend much time addressing is how we should read ‘women’ in ‘#BelieveWomen’. On one reading, it restricts the scope of the campaign to only testifiers who are female or have a feminine gender, or to only those who suffer identity-prejudicial credibility deficits in virtue of their

does the presumption of innocence even *mean* outside legal contexts? Kimberly Ferzan’s paper in this volume (*#BelieveWomen and the Presumption of Innocence: Clarifying the Questions for Law and Life*) pursues both questions with rigor and clarity, and my contribution will engage closely with hers. Obviously a slogan can be used to mean many things, and undoubtedly there’s substantial variation in what those who tweet #BelieveWomen actually mean. I’m pursuing a partly reconstructive, rather than purely descriptive project: I seek a charitable interpretation that construes the demand made in the most reasonable and well-grounded way consistent with (many but not all) the ways the tag is actually used.

A tempting—and popular—way to understand the hashtag is that it demands that we believe *what women tell us*, when and because they tell us. The most extreme formulation of this is that we *owe it to all women, without exception, to believe what they say without question*.⁵ This seems to be how Margaret Atwood interpreted the hashtag; when asked to comment, she dismissed it saying “Women schwomen—I don’t think you should believe all anything [...] It’s more useful to say listen to all women and take what they’re saying seriously enough to actually do investigations.”⁶ Ferzan leverages the conceptual resources of contemporary epistemology to offer a less extreme formulation of this interpretation, analyzing it “as a call to trust and as a call to non-reductionism [about testimonial justification]. That is, not only do we owe it to women to start with a baseline of trust but this trust will justify our believing them.”⁷

Like Atwood, Ferzan rejects the demand so construed, at least for legal contexts, noting that when the stakes are high we cannot base our beliefs simply on the say-so of women (or anyone, for that matter), but must rather examine the credibility of every witness, and apportion our degree of confidence in what was said to the evidential force we take the speaker’s testimony to

membership in this social group. This reading constrains the scope of the *people* whom we are to believe, but leave the *topics* on which we are to believe them unrestricted. An alternative reading, which I find more attractive, takes the primary aim to be to secure credibility for people who are reporting sexual assault or harassment. Insofar as these reports are *stereotypically* made by women, it is plausible that anyone making such a report is subject to what Emmalon Davis calls a ‘Content-Based Testimonial Injustice’: unjust skepticism or dismissal of a content, stemming from a prejudice against the social identity with which that content is associated (See Davis, ‘A Tale of Two Injustices: Epistemology of the Philosophy Pipeline’, manuscript on file with author.) On this reading, the prejudice to be addressed is principally against women, but the motivating issue is a tendency to dismiss testimony reporting the sexual misconduct of prominent men. What I say should hold for either reading.

⁵ This extreme interpretation is particularly favored by those who think the hashtag makes an unreasonable demand. There are in fact two variants of the tag: ‘#BelieveWomen’ and ‘#BelieveAllWomen’. The latter appears to have emerged later (it rose to prominence when David French re-tweeted Juanita Broadrick’s accusation that Hilary Clinton “tried to silence” her). As compared to the more ambiguous ‘#BelieveWomen’, the ‘#BelieveAllWomen’ tag is used disproportionately often by detractors (see Susan Finaldi, Opinion, ‘#BelieveAllWomen is a Right Wing trap’, *The New York Times* May 18 2020: <https://www.nytimes.com/2020/05/18/opinion/tara-reade-believe-all-women.html>)

⁶ Interview with Sam Gillette, *People Magazine*, Sept. 11 2019

⁷ Ferzan, *#BelieveWomen and the Presumption of Innocence: Clarifying the Questions for Law and Life* (this volume), p. 8. Ferzan is clear that she does not endorse this reading, but thinks it sufficiently plausible to merit spending the bulk of her paper filling in and evaluating. I will postpone discussion until §2, but very briefly: non-reductionism here refers to the view that the evidential justification we receive from testimony does not necessarily reduce simply to the evidence we have about the speaker’s reliability.

provide. Both Atwood and Ferzan are clear that they take the hashtag to be prompted by a real problem: women are persistently under-trusted specifically when they bring first-personal accusations of assault. And both think some epistemic fix is called for: we need to trust women more than we do. Atwood frames the fix as taking their testimony to give us reasons to carefully examine available evidence; Ferzan as giving us reason to intentionally increase our confidence in their testimony in order to cancel out our (probably subconscious) unjustified distrust of women as testifiers.⁸ But both think asking us to *believe* women demands too much.

While I agree that we should not simply believe what women say *because they're women*, I do not think that is the best understanding of the demand to #BelieveWomen. Nor do I think we should understand the demand as a call to be non-reductionists about testimonial justification. In fact, as I will argue in §2, I think the hashtag aims to call attention to a way women are *wronged* when not believed, rather than the fact that their testimony can make belief rational. In §3, I'll argue that in fact the demands of #BelieveWomen are quite modest; it does not require anything like simply accepting whatever women assert as true. We should distinguish between the attitudes of believing *a speaker* and believing *a proposition*, and resist the assumption that believing a speaker entails forming a full belief in the proposition asserted. I outline an alternative model according to which believing a speaker is taking their assurance to provide evidence of the truth of the thing asserted, and orienting one's further inquiry in certain characteristic ways. If we interpret the demand to believe women as a demand to respond to their testimony as I suggest, the demand to *believe women* is fully consistent with investigating the evidence in the ways Ferzan and Atwood advocate.

If this is right, it becomes harder to see why so anyone would think that there is a conflict between believing women and due process or the presumption of innocence. In §4, I suggest that the sense of conflict stems from the pragmatic implicatures generated by the appeals to the presumption of innocence in everyday contexts. Though the literal assertoric content of such appeals does not entail that victims' testimony is not evidence, it is nevertheless often reasonable for hearers to react as though the speaker does in fact dismiss victim testimony. This is because the presumption only concerns what to do when we have yet to receive evidence, and so when invoked *after* hearing victim testimony, implies that the speaker accords the testimony little to no weight.

One of the themes I'll return to throughout is that speakers who offer testimony—at least in the paradigmatic cases for which #BelieveWomen is invoked—have some sort of positive entitlement to being believed. I take this up explicitly in §5, in part to trace out the basis for this entitlement, and in part to demonstrate that the analysis I offer of what it takes fulfill our obligation to believe women does address these interests.

But while I think it is important to recognize that there is a way to interpret #BelieveWomen as a nuanced epistemic instruction, I do not ultimately think this is the *best* way to think about the social project of the hashtag campaign. As part of a large scale social project, #BelieveWomen does more than deliver instructions for correcting individual epistemic bias. In §6, I suggest that it works to call attention to a particularly pernicious trope operating in the

⁸ This is the corrective for testimonial injustice proposed by Miranda Fricker, *Epistemic Injustice: Power and the Ethics of Knowing* (2007, Oxford University Press).

background: that “women lie about assault”. As long as something like this is socially accepted, we are primed to respond to instances of women testifying about assaults with suspicion—an affect that crowds out the trust necessary to afford their testimony appropriate credibility.

My discussion centers on everyday contexts, rather than legal, partly because this is where the collision between the invocations of the hashtag and presumption is most vivid, and partly because I don’t think the epistemic norms for these two domains are so different. I will consequently use ‘testimony’ and ‘evidence’ in their loose everyday sense, rather than with their strict legal senses.

2. Distinguishing Questions of Permission and Obligation

Let’s start by taking a closer look at the subject matter of the call to #BelieveWomen. Ferzan suggests interpreting the hashtag as a call to testimonial non-reductionism because, as she writes, “only to non-reductionists is believing women about *believing speakers*.”⁹ It is true that non-reductionists foreground the speaker as the source of reasons for belief in testimony, but we should distinguish between two threads here, corresponding to two readings of the question “when should we believe testimony?”

The reading central to the reductionist-vs.-non-reductionist debate considers testimony as an input to our rational behavior as epistemic agents. Both theories aim to address a question about what anchors the evidential value of testimony, roughly: “when and how does being told that *p* provide sufficient justification to make belief in *p* rationally permissible?”. Call this the *evidential question*. The reductionist says the evidential value of testimony that *p* is indirect, and reduces to our independent reason to take the testifier to be credible. Non-reductionists maintain that when we receive testimony from someone that *p*, unless there are specific reasons to doubt them, we are directly rationally permitted to base our epistemic confidence in *p* on their assurance; we do not need to first investigate their credibility in order to be justified in believing what they have told us.¹⁰ Both focus on how testimony provides enough evidence to rationally *permit* belief in a proposition *p*; here there is no question of *owing* belief.

But the *reasons* non-reductionists give for taking speakers’ assurances that *p* to be directly justifying highlight a different aspect of the practice of giving and receiving testimony: it is an interaction of trust, a cooperative exchange between responsible moral and epistemic agents.¹¹ In offering testimony, a speaker ‘vouches for’ the truth of her assertion; she stakes her own competence as an epistemic agent as an assurance. There are deep parallels between this sort of exchange and other core exercises of normative agency, including the making of promises and

⁹ Ferzan MS p. 11, emphasis in original

¹⁰ For defenses and more thorough explication of testimonial non-reductionism, see Jennifer Lackey, “A Minimal Expression of Non-Reductionism in the Epistemology of Testimony,” *Noûs* 37 (2003): 706-23; Richard Moran, *The Exchange of Words: Speech, Testimony, and Intersubjectivity* (OUP 2018).

¹¹ Richard Moran, ‘Getting told and being believed’, *Philosophers’ Imprint* (2005) 5:1-29

contracts. The same reactive attitudes seem licensed in response to failures.¹² We feel resentment when we accept testimony but later discover the testifier was wrong, and injury or indignation at having our testimony unfairly rejected; we feel that we were owed more care.¹³ Non-reductionism stresses the ways that agency underwrites the justification provided by testimony, but does not emphasize—and some advocates of non-reductionism deny (see Moran 2005)—that someone who offers their testimony to us isn't just giving us a *permission* to accept their assurances, which we are free to take or leave. They issue an invitation, but it is one that, absent good reason, we cannot decline without wronging them.¹⁴ So, we can ask a different question about testimony: “under what circumstances do we *owe* it to the testifier to believe them?” Let's call this the *ethical question*.

It's important to recognize that this is very different from the evidential question. The focus of the evidential question is on what rational permissions we receive from testimony; the focus of the ethical question is on what we owe to speakers who testify to us. The answers to the two questions are also largely independent: settling what explains how testimony can provide adequate warrant to permit rational belief in a proposition does not settle when we *owe* it to someone to believe them, nor even what cognitive attitude one must adopt in order to count as *believing someone*. There are three distinct issues here:

- (a) Is the justification testimony provides for confidence in *p* direct or indirect?
- (b) When (if ever) do testifiers have an entitlement to be believed?
- (c) Does believing a testifier entail adopting an attitude of full belief toward the propositional content they asserted?

Taking #BelieveWomen to be a call for non-reductionism, I suggest, would subtly conflate (a)—an inquiry into the evidential question—with (b), an instance of the ethical question. When we ask whether we are really obligated to #BelieveWomen, the underlying issue is *what do we owe to women, when they testify?*

Interpreting the hashtag as demanding that we believe *what women tell us* is too quick: it presupposes a 'yes' answer to (c). To make it clear that there is space for a different answer, it will help to introduce another set of distinctions, this time between three different concepts we routinely use the word 'believe' to express. The first two are attitudes we take toward truth-evaluable propositions. The one most familiar to epistemologists, perhaps, is a subjective attitude of full confidence or high credence in the truth of a proposition *p*.¹⁵ The second is more action-

¹² These 'Strawsonian' reactive attitudes are a hallmark of interactions between morally responsible agents. See P.F. Strawson (1962) "Freedom and Resentment." *Proceedings of the British Academy* 48: 1–25.

¹³ Moran ("Getting told") explicitly draws the parallel between the act of assertion involved in testimony and promises. Both actions indicate commitment on the part of the speaker, and change the normative relations between speaker and hearer, generating new obligations and permissions.

¹⁴ Others have made this point as well. See Allan Hazlett (2017) On the special insult of refusing testimony, *Philosophical Explorations*, 20, 37-51; Holton, 'Deciding to Trust, Coming to Believe', *Australasian Journal of Philosophy* 72 (1994), 63-76.

¹⁵ I remain neutral here on the debate over whether the conception of full belief as a subjective attitude of full confidence can be reduced to high credence.

oriented: it is the attitude we take when we close inquiry *whether p*, having concluded that (given our evidence) we are justified in treating *p* as true and using it as a premise in practical reasoning. The third is an attitude taken not toward propositions but toward sources of information: trusting the assurances of source.

To demonstrate: imagine there is a screen that displays a number between 0 and 120. Some proportion of the time—say 74%—the number is the readout from a thermometer in Death Valley; the rest of the time it is randomly generated. Insofar as you treat the readouts on the screen as reliable evidence about the temperature, there is a natural sense in which we should say *you believe the screen*. Perhaps, given that you believe the screen, you will be inclined to treat the fact that it says ‘80’ as decisive evidence and adopt the belief that it is 80° in Death Valley. Or you could respond with more nuanced credences, reasoning from the fact that the screen displays an accurate readout 74% of the time to a credence .74 that it is 80° in Death Valley. Either way, the explanation of why you have the credences and beliefs you do will trace back, in part, to the fact that the screen ‘told’ you that it’s 80°, and you believe the screen.¹⁶

In which sense of ‘belief’ might we owe belief to speakers? One way to get at this is to ask what we do wrong when we refuse someone’s testimony. It is not just that we underestimate their evidential value as an informant, and so end up less epistemically rational than we might have been. The testifier has a different complaint: it is “an insult and an injury to not be believed.”¹⁷ Rejecting someone’s testimony treats them as an unreliable source, ineligible to give adequate assurances—and so as either incompetent or insincere. Treating someone this way is insulting if you lack specific reason to think that the speaker is in fact unreliable. Plausibly, speakers are entitled to *reasonable trust*; roughly, to being believed unless there is specific reason to doubt them. In §5, I’ll examine what interests could ground this kind of entitlement. For now, though, it will be enough to fill in the *content* of the entitlement.

First, notice that it doesn’t seem especially plausible that someone (to whom we do not stand any special relationship) has a right to be believed in either of the first two senses: to our developing full confidence in or closing inquiry on a proposition simply because they asserted *p*. Insofar as speakers are entitled to be believed, then, the better candidate understanding is the third sense: belief as trust in the assurances of a source.¹⁸ So in what follows I will (as much as possible) speak of *trusting* speakers, and *believing* propositions. That is, I embrace Anscombe’s conclusion that “believing someone (in the particular case) is trusting him for the truth—in the particular case.”¹⁹

¹⁶ Incidentally, both of these attitudes are open to you regardless of whether your justification for treating the screen’s readouts this way is that you have a wealth of evidence for the screen’s reliability (reductionist) or simply that you are entitled to treat it as default-reliable absent defeaters (non-reductionist).

¹⁷ G.E.M. Anscombe, “What Is It to Believe Someone?” in C.F. Delaney, ed., *Rationality and Religious Belief* (Univ. Notre Dame 1979):141-151, 150. The same point has been emphasized by several authors working on testimonial injustice, including Hazlett (2017), Pynn (forthcoming).

¹⁸ The suggestion that we frame believing *someone* as involving trust, rather than believing some proposition, is also made by Berislav Marušić and Stephen White, “How Can Beliefs Wrong?—A Strawsonian Epistemology,” *Philosophical Topics* 46 (2018): 97-114.

¹⁹ Anscombe, 151.

Putting this all together, what we have so far is an answer to (b): speakers have an entitlement to be believed—that is, to have their assurances trusted as evidence—when we lack specific reasons to doubt them as reliable testifiers.²⁰

From what we've said so far, this is not such a heavyweight obligation; indeed, one might think it eminently reasonable. If this is all #BelieveWomen asks, it is making a very modest demand. It is not that we owe special credulity to women *qua* women, as a group of testifiers. Rather, since being a woman is not itself reason to doubt a person's credibility, our general obligation to speakers entails that women, too, are entitled to reasonable trust. So far so good, but what are we to do when a speaker whom we trust in this way testifies that *p* is true? One might think (though I do not) that if you believe (trust) a speaker, you must believe *what* they tell you: adopt an attitude of full belief in *p*. We have arrived, then, at question (c).

3. Does Believing a Speaker entail full belief in what they say?

3.1 Three ways to believe someone

The conflict between #BelieveWomen and the presumption of innocence is a tension between two ethical demands on our epistemic activities. #BelieveWomen invokes speakers' entitlements to be believed when they testify. The presumption of innocence—or at least an appeal to it in everyday contexts—emphasizes people's claim to our not believing badly of them on inadequate evidence.²¹ These seem to come into conflict most directly in “she said/he said” cases, when our only or primary evidence of the wrongdoing is the victim's testimony.²² Critics of #BelieveWomen worry that the demand displaces necessary inquiry with something like solidarity or blind faith in the victim's accusation. Ferzan summarizes the concern thus:

²⁰ This view of what it *is* to believe a speaker is *prima facie* at odds with Moran's, on which the act of testimony is a *giving of assurance*, and the fact *that* the speaker offers such assurance provides (non-evidential) reason for the hearer's belief. But this appearance is misleading: Moran is at pains to distinguish the type of justification received from testimony *qua intentional assertion* from evidence one might receive simply from the fact that the person is talking. In taking the speaker's *telling you that p* to give you epistemic reason to increase your credence in *p*, you are responding to the testifier as a responsible agent; it is their assurance that you are treating as evidence, rather than the brute fact of their having produced some set of noises.

²¹ For recent discussions of such a claim, see, e.g., Basu & Schroeder, 'Doxastic Wrongs' in B. Kim & M. McGrath (Eds.), *Pragmatic Encroachment in Epistemology*; Begby (2018) 'Doxastic morality: A moderately skeptical perspective', *Philosophical Topics*, 46(1), 155-172; Gardiner (2018) 'Evidentialism and moral encroachment' In K. McCain (Ed.), *Believing in accordance with the evidence: New essays on evidentialism*; Moss (2018) 'Moral encroachment', *Proceedings of the Aristotelian Society*, 118(2), 177-205; Schroeder (2018) 'When beliefs wrong', *Philosophical Topics*, 46(1), 115-127.

²² I here follow Gardiner ('The “She Said, He Said” Paradox and the Proof Paradox' in *Truth and Trials: Dilemmas at the Intersection of Epistemology and Philosophy of Law*, eds. Zachery Hoskins and Jon Robson. Routledge.) in referring to these cases as '*she* said/he said', because the order of the testimonies—that her accusation comes first, predictably followed by his denial—is significant to the evidential value of each assertion. These cases do not straightforwardly pattern with simple disagreement cases, in which two equally well-placed and equally trustworthy speakers with matched or no incentives to lie disagree about a fact.

“If #BelieveWomen tells us that women have as much claim to be believed as men, that their testimony gives sufficient warrant for belief, then there is no nuance, no credences, and no probabilities. It is a question of which side you believe.”²³

Clearly mere side-picking is inappropriate in a courtroom; it’s not great practice in everyday epistemic dealings, either. The argument underlying this worry can be put like this: *believing a speaker* entails adopting an attitude of full belief in the content of their assertions without further inquiry. That attitude is inappropriate in she said/he said cases, because it conflicts with the presumption of innocence; therefore the demand to #BelieveWomen is unreasonable.

Focus on the first step of this argument. Why should we accept it? It does not follow from either reductionism or non-reductionism about testimonial justification. Grant that (directly or derivatively) a speaker’s testimony gives you some justification for confidence in *p*. If you do not form a full belief in *p*, but you *do* increase your confidence roughly in proportion with the justification received, do you wrong the speaker? Nothing we’ve said so far implies this: what we owe to speakers is to not reject their testimony without specific reason. Failing to respond to testimony that *p* with flat-out belief that *p* constitutes rejecting testimony only if *believing a speaker* entails believing the proposition asserted by the speaker. If it doesn’t, there need not be a conflict between the demands of #BelieveWomen and the epistemic care we should take when the stakes are high (and must take in courts as a matter of procedural justice). So the central question is whether believing *a person* necessarily involves taking an attitude of full belief toward the proposition they assert.

Something like this is a pretty common background assumption, I’ll grant. But I don’t think we should accept it. I won’t so much argue *against* the assumption as present an alternative, and suggest that this alternative has enough going for it that we should interpret #BelieveWomen as invoking it instead. The assumed picture is something like

- (i) To believe someone who asserts that *p* is take the attitude of full belief toward the propositional content of their assertion.

Ferzan does not quite endorse (i); she instead embraces a weaker formulation:

- (ii) To believe someone who asserts that *p* is “to take the claim on board at the same level of credence as the speaker offers it.”

This departs from (i) only in cases where the testifier hedges their claim, saying “I think that...” or “probably...”. But if the testifier flat-out asserts *p*, then (ii) would usually require flat-out believing *p*.

Notice that both of these are specified in terms of *posteriors*, the epistemic attitudes you have *after* you have heard the testimony. We certainly sometimes use phrases like ‘believe me!’ to express a demand for one of these two attitudes. But I think there is also a natural sense of believing someone that is concerned not with posteriors but with *priors*: your dispositions to take the speaker’s assurances to justify raising your confidence in *p*, which you have even *before* you’ve

²³ Ferzan ms, p. 22

heard their testimony or worked through your total evidence to reach your ultimate verdict on p . This is taking them to be a trustworthy epistemic source; someone whose intentional assertions give you evidence not merely that they are talking, but that the propositions they assert are likely true. It is the sense in which you believe your zoologist friend when she tells you an obscure fact about Australian marsupials, but not the breathless five-year-old who relates his adventures fighting an invisible dragon. You believe someone in this sense when you are disposed to treat their assertions as giving you strong evidence in favor of the truth of what they've said. More formally,

(iii-a) To believe someone who asserts that p is to take the fact that they testify that p to give you epistemic reason, greater than a magnitude m , and proportional to your estimation of their trustworthiness, to increase your confidence in p .²⁴

On this interpretation, #BelieveWomen instructs us to ensure that our estimation of the evidential support for p that we receive from hearing a woman testify is at least m , absent specific evidence that she is not trustworthy—and be disposed to increase our credence in p accordingly. This is precisely the disposition that we counted as believing the screen's reports about the temperature in Death Valley. Importantly, even if your estimation of a witnesses' trustworthiness is quite high, believing them in the (iii-a) sense won't always result in having a high enough posterior confidence in p to count as full belief in p . If you started out with a low enough initial degree of confidence, or the rest of your evidence strongly supports $\sim p$, or if belief is not just a matter of high credence, you may well *believe the person who told you p* but not believe p .²⁵ Though certainly this introduces a level of nuance not usually present in twitter exchanges, I think it is still true to the everyday phenomenology of trusting speakers and other sources of testimony.

I expect that some readers will immediately find (iii-a) unsatisfactory; surely *believing someone* involves much more than just increasing my credence in what they've said! I agree, but resist the suggestion that we should fill in the details of this 'much more' by requiring that our posterior confidences be substantially higher (or reach full confidence in p), as (i) and (ii) do. Rather, there are a number of other ways that trusting someone's testimony should restructure our epistemic orientation: changing what we take to be relevant counter-possibilities, framing what questions we

²⁴ m here is a variable that controls how demanding this standard is; I remain neutral on what specific values it should take. Note that this constraint is applied to *priors*, rather than *posterior* credence in p : it requires that the credence you are disposed to have in p conditional just on the speaker's testimony that p be sufficiently greater than your credence in p without the testimony. Expressed more formally: $C(p | \text{'S testifies that } p) - C(p) > m$. This constraint is relatively weak; it does not prescribe any particular posterior degree of confidence in p , since if your total evidence strongly supports $\sim p$, you can satisfy this condition and still ultimately have a final credence in p below 0.5. My thanks to Richard Bradley for discussion on this point.

²⁵ This holds for both of the two concepts of belief as a propositional attitude I mentioned above. Suppose belief is credence above a certain threshold or an *attitude of full confidence* (equivalent to credence 1). Then if you started with a low prior, or had mixed evidence about p , you may end up with a credence too low to count in belief even if you have a lot of confidence in the credibility of your source. Suppose instead that (in line with Ross & Schroeder 2014), whether I believe p is a question about what my reasoning defaults are: if I routinely consider the risks given the remaining probability of $\sim p$, then I don't have a full belief that p —even if my credence in p is extremely high. Conversely, if I routinely assume that p (either consciously or by not even questioning whether p , as one might if ' p ' is 'the coffee is safe to drink'), then I believe p , even if my actual credence in p is pretty low (.58, say. Probably believing when at this credence is irrational, but I'd still say someone who reasons this way believes p). Believing in this sense is *closing inquiry*—and you can have reached quite a high credence in p , based on your source's testimony, without closing inquiry about p .

take to be important, what evidence we look for (and how hard), which working assumptions we adopt, and, of course, which propositions we start treating as default-accepted unless we have specific reason to doubt them.²⁶

Perhaps especially relevant: if you believe the person who testifies but fall short of forming a full-belief in the content of their testimony, you'll still be keeping an eye out for some error possibilities—maybe there's additional evidence that contextualizes what they've said—but you won't be *focused* on whether they're lying. You won't be discounting what they've said in proportion to some statistic about how often people 'like them' have been wrong in the past. If you believe the person, you apply scrutiny in the first instance to your overall evidence on *p*, not to the credibility of the witness. To reflect all this, we should supplement the *evidence-oriented* condition (iii-a) with an additional *inquiry-oriented* condition, something like (iii-b):

(iii-b) [...] *and* to focus your subsequent inquiry on *p* (rather than on the trustworthiness of the testifier).

This doesn't mean we *disallow* evaluations of the speaker's credibility, but they are not the central focus of our inquiry if we believe the speaker.²⁷ It does require treating *p* as a live question, rather than as likely false: it directs us to seek the evidence we would expect to find if *p* is true, rather than spending energy inventing hypotheses to explain how a lying witness would produce the specific testimony received. It also affects how we rank the plausibility of the various hypotheses: believing a witness' statement that *p* gives us reason to take seriously possibilities that render *p* consistent with our other evidence, even if they're complicated, and to demote hypotheses that would imply that our witness is lying, even if they are simpler.²⁸

3.2 *She Said/He Said Revisited: Believing Conflicting Testimony*

If we accept that *believing* (trusting) *a speaker* need not involve *forming the full belief that p*, then we can avoid the worries that many (including Ferzan and Atwood) raise about the call to #BelieveWomen. There need not be any conflict between satisfying this demand, on the one hand, and doing due epistemic diligence and respecting due process, on the other.

²⁶ I owe this insight to Georgi Gardiner, who develops some of these themes in 'The Reasonable and the Relevant', *Philosophy & Public Affairs* 2019, and 'Doubt and Disagreement in the #MeToo Era' (forthcoming) *Feminist Philosophers on #MeToo*, ed. Yolonda Wilson. Routledge.

²⁷ Believing a speaker does not mean we cease looking for defeaters, either. But it does change what sort of defeater we seek: we should remain vigilant for *rebutting defeaters* (facts which indicate that *p* is false), but not seek *undercutting defeaters* (facts that debunk our trust in the speaker). Thanks to Andrew Chignell for discussion on this point.

²⁸ This is in fact the way I read legal theorist Sherry Colb's interpretation of the hashtag, though Ferzan glosses Colb's view as counseling "police to act *as if* they believe a complainant" (Ferzan ms p.24). See Colb, "What Does #BelieveWomen Mean?" Verdict: Legal Analysis and Commentary from *Justicia*, November 7, 2018. <https://verdict.justia.com/2018/11/07/what-does-believewomen-mean>

Consider the problem of conflicting testimony. Ferzan is absolutely right that *if* believing a person requires forming full belief in the proposition they assert, then when speakers assert contradictory contents, we cannot coherently believe them both. Suppose two people (let's call them Alex and Beth) testify to inconsistent contents. It is not possible to form the attitude of full belief in p in response to Alex, *and simultaneously* form the attitude of full belief in $\sim p$, in response to Beth, and be coherent. We can only withhold judgment, or pick a side. So if (i) or (ii) are correct as a model of what it is to believe a speaker, when presented with conflicting testimony you face “a question of which side you believe.”

But even when testimonials conflict, it is perfectly possible to take both testifiers to be *trustworthy*, and so to update your credence in p on both ‘Alex asserts that p ’ and ‘Beth asserts that $\sim p$ ’. If you trust both equally, you may end up with the same credence in p as you had when you began, but if believing someone only requires something like (iii), then you will have believed both of your sources. You will not have rejected the testimony of either, though you will likely need to do more investigating to know *whether to believe p*.

Importantly though, this is *if* your estimation of trustworthiness is roughly equal for both testifiers. In she said/he said cases, as Gardiner points out, we are often not justified in taking the trustworthiness of a denial to be on par with the trustworthiness of the accusation. In general, a speaker's having a strong incentive to assert something—independent of its truth—is a reason to trust their assertion a little less. A person who stands accused of assault always has an incentive to deny it; the same does not hold for *making* an accusation.²⁹ Plausibly the base rates matter here, too. A recent meta-analysis of reporting rates found that only about 5% of rape reports filed are ‘false’—that is, the evidence yielded by investigation *either* on balance suggested the assault did not occur, *or* was insufficient to establish that one did.³⁰ Taking these numbers at face-value, we would be justified in taking an accusation (at least one issuing in a police report) to be reliable 95% of the time—which is a fair bit higher, surely, than we should expect from denials.

3.3 High costs of belief

On their face, both (i) and (ii) are quite demanding interpretations of what it is to believe a speaker. Ferzan takes even (ii) to be ill-suited to the inquiry necessary in legal contexts, though perhaps more amenable to everyday dealings. Still, one might worry that it is sometimes too demanding even there. Believing that a loved one has committed sexual assault sends shockwaves through your life, in a way that believing that it rained while you were in the library does not. Intuitively, we are *at least permitted* to be epistemically cautious, requiring more (or more compelling) evidence before we close inquiry on high-stakes propositions like these. Ferzan seems optimistic that an advocate of (ii)

²⁹ Gardiner, Doubt and Disagreement in the #MeToo Era’.

³⁰ Ferguson and Malouff (2016), ‘Assessing Police Classifications of Sexual Assault Reports: A Meta-Analysis of False Reporting Rates’, *Archives of Sexual Behavior*. Because the police records used for the studies analyzed do not distinguish between investigations concluding with evidence that the assault did not occur, and those concluding with insufficient evidence that one did, these numbers almost certainly over-counts *false* reports, but it is hard to estimate by how much.

could appeal to pragmatic encroachment here, but I think the theory cannot do that work; it is addressed to the wrong question.³¹

Pragmatic encroachment is a theory that explains why a degree of evidential justification that usually suffices to permit rational belief is insufficient when the stakes of mistakenly having that belief are high.³² So if we're addressing the *evidential question*, asking whether testimony gives agents facing high costs sufficient evidence to make belief that *p* rational, pragmatic encroachment can be leveraged to say 'no'. But if that's not the relevant question—if we're instead asking whether respecting the speaker's entitlement to be believed (trusted) still requires the agent to believe *p*, even when that belief would be costly—then pragmatic encroachment only has space to operate if we interpret *believing speakers* my way. If we embrace either (i) or (ii), then believing someone requires at least matching your credence to the strength of their assertion; if the speaker flat-out asserts *p*, then both (i) and (ii) require you to flat-out believe *p* in response, or else you fail to believe the speaker. There is no room here to ask whether your total evidence, including their testimony, is adequately justifying given the stakes you face.

Of course, failing to believe only *wrongs* the testifier if we owed them belief. Maybe there's room for an advocate of (i) or (ii) to argue that speakers are only entitled to be trusted when their testimony in fact gives adequate justification for full belief *given the stakes facing the hearer*. But this suggestion makes the foundations of the speaker's entitlement pretty mysterious. The stakes of belief *for the hearer* are often not transparent, and not obviously of concern, to the testifier. So if what the speaker has a claim to is *being trusted just when the hearer's evidence rationally permits full belief in their context*—to your believing that *p* on their say-so if and only if your evidence suffices to justify belief, given your actual and unknown circumstances—the grounds for this moral claim must be quite abstract.

It's hard to think which interests of the speaker could ground such a claim, or how they could complain that they are wronged by failure. It cannot be that they are wronged whenever we fail to form beliefs which we are justified in having. We do *that* all the time, when we fail to draw inferences on the information we already have, or fail to notice or process evidence in our visual field, etc. The complaint a speaker has against our failing to believe them *qua* testifier is not principally that we have failed to form a justified cognitive attitude.

³¹ Ferzan writes: "To form that belief is life changing for the believer. And in higher stakes situations, we should expect that pragmatic encroachment will require more evidence before you feel comfortable knowing something." (p. 15-16)

³² Fantl & McGrath, 'Evidence, pragmatics, and justification', *Philosophical Review* (2002): 111(1), 67–94; Hawthorne & Stanley, 'Knowledge and Action', *Journal of Philosophy* (2008): 105(10), 571–90; Ross & Schroeder, 'Belief, credence, and pragmatic encroachment', *Philosophy and Phenomenological Research* (2014): 259–288; Weatherson, 'Can we do without pragmatic encroachment?', *Philosophical Perspectives*, (2005): 19, 417–443. For an overview of how different theorists formulate pragmatic encroachment, and how moral stakes might be relevant to the evidence required to justify a belief, see Bolinger 'Varieties of Moral Encroachment', *Philosophical Perspectives* (2020). For a discussion of how the moral costs of failing to form justified beliefs about rape could make skepticism inappropriate, see Bianca Crewe and Jonathan Jenkins Ichikawa, "Rape Culture and Epistemology." In Jennifer Lackey, ed., *Applied Epistemology* (forthcoming OUP).

Pragmatic encroachment is well-suited to explain when an agent should close inquiry on *p*, given their evidence and the comparative costs of relevant errors risked. But the theory is not equipped to tell us how to balance the harms to others of rejecting testimony against our own high stakes to determine when others have a *claim* to our believing. By contrast, the notion of *reasonable trust* is capacious, well-suited to exactly this kind of balancing. As Karen Jones develops the notion, reasonable trust

“...can encompass different kinds of reasons that count in favor of trust, from epistemic, to instrumental, to reasons that follow from commitment to an ideal. In some cases, you will need evidence sufficient for a strongly justified belief that the other is trustworthy. In other cases, trust can be reasonable in response to someone’s explicit invitation to trust (depending on domain and consequences of misplaced trust)...”³³

Reasonable trust in a trustworthy but fallible speaker won’t straightforwardly generate an obligation to match your credence in *p* to the strength with which they asserted it, much less to flat-out believe it. You might well have other evidence that you shouldn’t disregard, or be facing significant practical stakes. If I have reasonable trust in Annie, then when her assertion that *p* gives me adequate evidence in context for full belief (together with my other evidence), I will adopt an attitude of full belief toward *p*. But otherwise I will simply become more confident in *p*.

All the action happens in the ethical question, not the evidential question. We are not concerned that in failing to believe someone’s testimony we will be *irrational*; we are concerned that we will *wrong them*. Our epistemic lives are shot through with ethical considerations, and while recent debates over “moral encroachment” on the epistemic have focused on whether we can wrong someone through *what* we believe about them,³⁴ the important point here is that whether we believe (trust) *them* is also ethically charged. This need not mean that it is in competition with more narrowly epistemic aims; we need not sacrifice truth in particular cases in order to advance the ethical project. Indeed, bringing our actual epistemic practice into alignment with the normative demands of reasonable trust may make us *more* rational by a number of metrics, by correcting ways our inquiry is presently distorted by bias.

Acknowledging these ethical dimensions is important, not just because it appropriately foregrounds that we *owe* it to stakeholders to improve our epistemic practices, but because the distortions we aim to address can be incredibly difficult to spot or fix when approached as purely epistemic problems.³⁵ Distrust does not always manifest as an unreasoned refusal to believe anything someone says. We can distrust someone by failing to attach *adequate* weight to their testimony, or by holding them to inappropriately strict standards (requiring perfect recall of all details from an event years in the past, or being able to rule out *quite* unlikely alternative

³³ Jones, *Counting on One Another: A Theory of Trust and Trustworthiness* (unpublished manuscript) Chapter 4, ms pg. 17.

³⁴ See references in *supra* note 20.

³⁵ Gardiner (‘Doubt and Disagreement in the #MeToo Era’) surveys a range of reasons why the symptoms of unjustified distrust of women can, in these cases, deceptively feel like cognitive achievement, and so be uniquely difficult to identify as rational distortions.

explanations of events, etc.) before we will accept what they say. It can also manifest in seeking defeaters too readily, or focusing disproportionately on the possibility that they are lying or unreliable.

4. Innocent Presumptions and The Presumption of Innocence

If you're persuaded by what I've said so far—that #BelieveWomen urges us to place the same default trust in women that we place in other testifiers, and that this requires us to take the fact that a woman testifies to p to provide evidence in favor of p 's truth—then it is puzzling how appeals to due process, or reminders about the presumption of innocence, conflict with appropriate trust in women's testimony. Is this all just a confusion? Given the demands of #BelieveWomen (as I propose to understand them), is there *anything* inappropriate about reminding interlocutors of the importance of the Presumption of Innocence?

While the obligation to believe women as I have interpreted it is fully consistent with respecting due process and the presumption of innocence, there is still a conflict *implied* by such reminders. To draw this out, let's start by getting more precise about what the 'presumption of innocence' does and doesn't involve, in American criminal law contexts. As Ferzan's discussion makes plain, the content of the presumption is procedural; respecting it involves doing two things:³⁶

- *POI-Default*: setting 'innocent' as the default for what to conclude if the relevant standard of proof is not met.
- *POI-as-No-Evidence*: taking the attitude that one has no evidence concerning whether the defendant is guilty *prior to* the proceedings.

POI-Default is necessary in trials because the court cannot simply *withhold* in the way that an individual can, or even conclude with a hedged "probably the accused is guilty." Because of the structure of a trial, the finder of fact must in the end rule either p or $\sim p$; the standard of proof sets a threshold for the degree of evidential support that will suffice to justify ruling that p , and the presumption of innocence identifies $\sim p$ as the default if that threshold is not cleared.

The *presumption of innocence as no-evidence* requires that our *priors*—our degree of confidence in p before we hear any evidence—should be either concentrated at .5 (reflecting the fact that we think p is as likely to be true as false), or spread evenly across the full space of possibilities 0-1.³⁷ So, prior

³⁶ These correspond roughly to the main functions of the presumption of innocence in American criminal law, as characterized by Ferzan. The *POI-as-No-Evidence* is the '*probatory presumption of innocence*', which Ferzan glosses as meaning "that the jury (or other state actor) starts with the presumption that it simply has no evidence of the defendant's guilt." The *POI-Default* highlights the role of the presumption in allocating the burden of proof; it is "the corollary of proof beyond a reasonable doubt" (Ferzan ms, p. 30).

³⁷ The no-evidence starting point is not necessarily equivalent to taking it to be precisely equally likely that the accused is innocent as that he is guilty. Some argue that jurors ought to be permitted to use their knowledge of background facts about the frequency of the crime to inform their prior credences; I here opt for the less committal suggestion that they spread their credences across the full range. As I'll demonstrate in a moment,

to hearing the testimony, the presumption of innocence requires that we not take *the bare fact of an accusation* as evidence of guilt.

When invoked in informal contexts, the presumption is sometimes glossed as requiring something much stronger than this:

- *Material-POI*: to outright believe *that* the accused is innocent ($\sim p$), until presented with overwhelming evidence to the contrary.

Material-presumption of innocence requires that we start with full confidence in $\sim p$. There are reasons to worry about whether assuming that the accused is *in fact innocent* is coherent in all the contexts in which the presumption of innocence is required.³⁸ But even this way of interpreting the presumption of innocence does not *genuinely* conflict with believing women, if we genuinely respect the evidence we receive.

Here's why: an accusation is made. We start with no evidence, then the accuser offers her testimony that p . If we're good Bayesians, we update by conditionalizing our prior in p on how probable p is given this new evidence. Where this leaves us depends on where we started, how attached we were to that starting point (sometimes called *resilience*), and how much we trust the evidence we got. If we put more stock in the trustworthiness of the testimony than in our starting presumptions—which we should—then no matter our starting point, it will be pretty easy for testimony that p to move us to significant confidence in p . The more attached we are to the starting point (or the less we trust the testimony), the more the difference between the alternative understandings of the presumption of innocence makes a difference to the post-update degree of confidence. If the starting point is resilient, it will take an overwhelming amount of evidence to convince someone who starts out presuming $\sim p$ that in fact p is probably true. But insofar as the presumption of innocence is a stance defined by *the absence of any evidence*, the starting credence it yields should not be resilient, and so should easily shift in response to the weight of evidence with any real probation force.

This fact is actually what helps to make sense of the complaint that invocations of the presumption of innocence in everyday contexts—like twitter exchanges—imply that accusers are lying, or shouldn't be trusted. There are various pragmatic principles, or *maxims*, typically operating in the background of our everyday conversations. To abide by them, you must make your contribution not only true, but efficiently informative and on-topic, given our shared information and the goals of our conversation.³⁹ Reminding your conversational partners about what to do in the absence of evidence is an unhelpful contribution to make immediately *after* receiving evidence. So if you're being cooperative, you must mean something else by asserting it: not just that presuming innocence is a good starting place, but that *even now, having heard the testimony*, we

the differences do not matter for our present purposes. (See discussion in Richard D. Friedman, "A Presumption of Innocence, Not of Even Odds," *Stanford Law Review* 52 (2000): 873-87.)

³⁸ Ferzan discusses several of them in her paper (pp.29-32); I will not rehearse them here.

³⁹ These are the Gricean maxims of Quality, Quantity, and Relevance. See Grice, H.P. (1975) *Logic and Conversation*.

shouldn't stray far from that starting point. This would only be true if you took the testimony to have little to no evidential value.

Here it might be useful to return to the screen example. If >50% of the time the screen just displays a randomly generated number, then seeing what's on it gives you little reason to believe that's the temperature in the desert. But it would be inappropriate to retain this attitude if you knew that the screen is pretty reliable (especially if it was known to report falsely only 5% of the time!). So if, *after* seeing '80°' on the screen, you reprimand someone for saying that's probably the temperature in the desert right now, your assertion pragmatically implicates that we cannot treat the screen as reliable.

Given that respecting testimony entails taking it to be some evidence for p , reasserting either that we should presume that $\sim p$, or that we lack probative evidence as to whether p , is to disregard the testimony we've been offered. When an appeal to the presumption of innocence is used *as a reply* to the claim that victim testimony is probative evidence, the underlying picture thereby forwarded embeds skepticism about the testimony. It pragmatically implies that the testimony either isn't evidence at all, or counts for very little. Both of these possibilities are incompatible with taking the testifier to be a reliable source of information, and so imply that she is either lying or incompetent. Making these appeals, then, *is* in conflict with the demands of #BelieveWomen, though there is no conflict between actually respecting the presumption of innocence and believing women.

5. The Entitlement to *Be Believed*, Revisited

The account that I've offered is reasonably lightweight, as obligations to believe go. It basically only requires being disposed to treat women's testimony as relevant evidence. Hopefully this sounds reasonable, which we should be doing anyway just to be rational. You might suspect that this is because my account is not really an analysis of what it takes to believe someone, or at any rate not what is invoked by #BelieveWomen. So—especially if you were skeptical about the hashtag's demands—you might feel that I've changed the subject. Does the lightweight thing I've outlined really count as *believing a speaker*? And even if it does, is it plausible to think that speakers are actually entitled to that?

To settle these fears, I'll try to make good on the promissory note from earlier to revisit the interests that ground a speakers' entitlement to be believed. If what I've outlined can serve *those* interests, that's some evidence that my proposal should count as a viable interpretation of what it takes to *believe a speaker*.

There are of course some (by now familiar) non-instrumental interests a speaker has in being taken to be a credible source of testimony, at least by some hearers, sometimes. These are the speaker's interests in participating as an agent in good standing in the epistemic community. In rejecting their testimony, we represent them as unreliable. This is both diminishing and excludes

them from the participatory goods of a valuable epistemic practice.⁴⁰ Plausibly these participatory interests at least ground a claim against the type of epistemic injustice that is the primary focus of Fricker (2007): being specially distrusted, *relative to others*, due to an identity prejudice. But this is not yet a claim to *being believed*. It isn't clear that interlocutors who uniformly resist trusting *anyone's* testimony about sexual assault or harassment transgresses any speaker's claim to being counted as a participant in the epistemic community more generally.⁴¹

The simpler explanation for how exactly we *wrong* a speaker when we inappropriately reject their testimony appeals to more instrumental considerations. Let's focus in particular on the social and political role of testimony, and work backward to characterize the testifier's moral interest in being believed.⁴² There are some sorts of facts for which we simply must take someone else's word. The reasons for this forced reliance vary: evaluating the evidence might require expertise we do not have, or access to evidence we cannot get, or the ability to observe phenomena from a particular vantage, point in time, location, or social identity. To have evidence-sensitive beliefs or credences about these facts at all, we will have to trust someone's testimony; we will have to rely on someone vouching for a fact.

When a speaker's concrete interests depend on policy decisions made about facts of this kind, and she *does* have direct epistemic access to those facts, she has a very strong interest in being considered eligible to vouch for them. When she is not—when those who influence the policy decisions are not disposed to treat her assurances that *p* as evidence in favor of the truth of *p*—then there are two possibilities: either they trust *no one to vouch for those facts*, or they deny her (and/or her social group) influence that they grant to others.

If we trust no one to vouch for the facts, we distrust all testimony on the issue, and must throw up our hands saying “who can know?” whenever an accusation is made without enough independent evidence to warrant belief *in the absence* of the victim's testimony. There are special contexts—when the testimonials are equally credible, a thorough investigation has been done, and the evidence remains equivocal—in which we must accept that the facts are murky. But when this is the default position, it is functionally equivalent to declining to enforce victims' rights. Since the accused always have a strong incentive to issue a denial, this policy will lead to systematic failure to enforce the rights of victims or protect them from future victimization.

We are typically at least slightly more trusting: willing to believe the testimonials we find plausible, while ignoring the others. The trouble here is that we evaluate whether testimony sounds trustworthy by quickly checking it against how common we take the type of event reported to be, and how often we think people lie about it. (This is why, for instance, no one believes men on

⁴⁰ Fricker (2007) emphasizes wrongs of this sort, which have been elaborated by subsequent authors writing on testimonial injustice. There is a particularly nice discussion in Pynn (forthcoming), ‘Testimonial Injustice and Epistemic Degradation’ in *New Essays in Applied Epistemology*, ed. Jennifer Lackey.

⁴¹ This form of domain-specific distrust does not count as a testimonial injustice on Fricker's original model. It might be accommodated by the notion of *content-based* testimonial injustice outlined by Emmalon Davis (‘A Tale of Two Injustices’). However, since the grounds for rejecting testimony in this domain is skepticism, rather than an associated identity prejudice, it may not fall under any direct model of testimonial injustice.

⁴² My thanks to Daniel Viehoff for helpful discussion on this point.

Tinder who say they're 6' tall). But if our sense of these relative frequencies isn't itself informed by reliable evidence, we're likely to be woefully inaccurate, and misplace our (mis)trust.

Similarly, when deciding how much weight to give to testimonial evidence, it's well and good to emphasize the need to balance the risks to the accuser of dismissing a true accusation against the risks to the accused of accepting a false accusation. But to do this balancing successfully, we need to have some estimation of the comparative likelihoods of each type of error. We cannot hope to succeed if we procedurally and systematically underestimate the odds of one type of error (e.g., by discounting the testimony of those who are in a position to give evidence).

Unfortunately *all* of these harms are on display in our current approach to accusations of sexual assault. We tend to dismiss claims that we take to be improbable, and police are more likely to decline investigation if they think false reports are common. But this is something of a vicious cycle. There is a significant discrepancy between the actual comparative prevalence of assaults vs. unsubstantiated accusations, on the one hand, and beliefs about these rates held by members of the general public and police departments, on the other. In the US, victims only file police reports in 32-59% of sexual assaults.⁴³ The reported frequency of assault is thus substantially lower than the real frequency of victimisation. Routine non-investigation magnifies the disincentive for victims to report or testify, further widening the gap between the actual rates of assault and the evidence preserved in police records. The all-too-predictable consequence is that the people tasked with *investigating whether a reported assault occurred* approach victim testimony already skeptical about its evidential value, believing assaults to be much less, and false accusations much more, common than they actually are. As Turkheimer (2019) summarizes,

“In one survey of nearly nine hundred police officers, more than half of the respondents stated that ten to fifty percent of sexual assault complainants lie about being assaulted, while another ten percent of respondents asserted that the number of false reports is fifty-one to a hundred percent. Another study found that more than half of the detectives interviewed believed that forty to eighty percent of sexual assault complaints are false.”⁴⁴

As a reminder, the actual rates of false reporting are closer to 5%.

To summarize: we have an interest in our testimony being received as evidence, because we need to be able to offer evidence to motivate action to protect our rights. Thus even though the duty I have outlined is not very demanding, this lightweight duty serves a pretty heavyweight moral interest for the testifier. So very plausibly there *is* a right to being believed in the sense that I have outlined.

To be clear, I do not think that this minimal entitlement is the only form of belief we can owe to someone. I think we do owe this to strangers, and yes, to all women. (Remember, we owe a *qualified* duty, to treat their testimony as reason-giving when we lack specific reason to doubt their reliability.) But we can acquire thicker obligations in the context of closer relationships. We might

⁴³ U.S. Bureau of Justice Statistics, M. Planty and L. Langton, “Female Victims of Sexual Violence, 1994-2010,” (2013). The belief that the police won't do anything, or that the personal costs of reporting are too high, are among the most common justifications given for not filing a report.

⁴⁴ Turkheimer, “Incredible Women: Sexual Violence and the Credibility Discount,” *University of Pennsylvania Law Review* 166 (2017): 1-58, 16; Citing Amy Dellinger Page, ‘Gateway to Reform? Policy Implications of Police Officers’ Attitudes Toward Rape’, 33 AM. J. CRIM. JUST. 44, 55 (2008); Martin Schwartz, *National Institute of Justice Visiting Fellowship: Police Investigation of Rape— Roadblocks and Solutions*, 28 (Dec 2010).

indeed have duties to believe the best of our friends, or accept the testimony of our partners even when it is *unlikely* to be true, given the rest of our evidence. I do not want to rule out the possibility of these duties, nor would I deny that we might remind someone of these obligations by saying they ought to “believe me.”⁴⁵ My aim in this paper has only been to articulate the nature of the entitlement that speakers may have to be believed by people they have never met.

6. The Heart of the Problem: A Social, Affective Deficit

I’d like in closing to briefly float a rather different interpretation of the function of #BelieveWomen. It’s natural to think about testimony as an interaction between individuals: one speaker, one hearer, abstracted from a larger community. But as I emphasized in the previous section, there’s also a social dimension to testimony. This extends to credibility assessments; we do not just rely on each other to vouch for facts. We also vouch for people as sources of information. Rather than calculate the trustworthiness of every speaker *ourselves*, we outsource—or distribute—this labor to others in the community. We do this by relying on social signals about trustworthiness.⁴⁶

If we interpret #BelieveWomen not as an imperative demanding full belief, but as an epistemic instruction aimed to correct the underlying credibility deficit women face when *individuals* assess their evidence, we might think of it as saying, roughly, “Women have a quite low misleading-testimony base rate for first-personal reports of sexual harassment or assault, so assign high credence when they testify.” Formalized, this would presumably be something like: Where p is a first-personal victim’s report of sexual assault, the relative frequency in past cases of [$\sim p$ & a woman testifies that p] is quite rare. So, your credence in $\sim p$, conditional on [a woman testifies that p] should be very low. So all else equal, on getting evidence *a woman testifies that p* , the rational agent should update to a high credence in p . Simplified into advice that your father might follow in his everyday activities, this would be: *trust women (far) more than you do on this topic.*

As a direction to individual believers, it is unlikely that a corrective of this kind will go very far in addressing the harms that motivated the #BelieveWomen campaign. The extent of our distrust—the readiness with which hearers seek defeaters for women’s testimony to assault, and the resilience with which they cling to the presumption of innocence even *after* hearing testimony—suggests that the issue is not simply a collection of individual mistakes about the weight of their evidence. What #BelieveWomen is up against isn’t, I suggest, an individual *cognitive* problem, but an

⁴⁵ My thanks particularly to Johann Frick and Elizabeth Harman for discussion on this point. For further development on the suggestion that we might have relationship-based obligations to believe well of someone, see Sarah Stroud, “Epistemic Partiality in Friendship,” *Ethics* (2006): 498-524; Simon Keller, “Friendship and Belief,” *Philosophical Papers* 33 (2004): 329-351.

⁴⁶ This theme is developed at length in Jones, *Counting on One Another: A Theory of Trust and Trustworthiness* (manuscript).

affective, social one, in which something emotive—the *fear* of false accusation—looms large, and crowds out the trust that underlies confidence in testimony.⁴⁷

Relatively early in her paper, Ferzan considers in passing that “the demand is an attempt to create a positive generic” (ms 4), rather than an epistemic instruction. I think there is something to this, but it isn’t straightforward. A ‘generic’ is a bare plural sentence, like

Ducks lay eggs.
 Mosquitos carry West Nile virus.
 Boys don’t cry.
 Pitbulls are aggressive.

We readily accept such sentences as true, assert them in everyday contexts, and rely on them in reasoning and argumentation. But it’s far from clear what, precisely, we should take the semantic content of such sentences to be. They seem to predicate properties in a way that quantify over whole groups, but the strength of the quantification is implicit—there is no “all”, “most”, or “some”—and it is not apparently reducible to these common quantifiers. Does “ducks lay eggs” claim that *most* ducks lay eggs (or just most female ducks)? That ducks *generally* lay eggs? That *it’s striking that they do so*, or that it’s a *characteristic* property of the kind that they engage in egg-laying? The truth conditions of generics are murky; that’s part of why they’re so interesting to political philosophers of language and linguists.⁴⁸

Generics are pretty close to cognitively basic⁴⁹; we use them to sort the world into manageable categories. But they’re also incredibly slippery: because they aren’t explicitly quantified, when evaluating a generic we tend to read it as ‘[some] pitbulls are aggressive’, and so accept it as true on the basis of just a couple salient or striking cases. But when deploying them in reasoning, we have a strong tendency to treat it as ‘[most/all] pitbulls are aggressive’.⁵⁰ This tendency is particularly pronounced when they concern a dangerous property. What Leslie (2017) terms ‘striking property generics’ can establish fear and distrust of an entire category, based on a few striking instances. On learning of just a handful of cases of Great White sharks attacking surfers off the Australian coast, we move quickly to accept the generic ‘sharks attack bathers’, and subsequently allow it to frame our approach to all sorts of sharks as *dangerous*. There’s an evolutionary advantage to learning quickly, especially about danger—but it overgenerates. The overwhelming majority of sharks will never attack, but this fact does little to unseat the striking property generic that disposes us to view all members of the kind with suspicion.

⁴⁷ Jones, ‘Trust, Distrust, and Affective Looping’, *Philosophical Studies* (2019) 176:955–968.

⁴⁸ See, e.g. Ariel Cohen, “Generics, Frequency Adverbs, and Probability”, *Linguistics and Philosophy*, 22(3):221–253, (1999); David Liebesman, “Simple Generics”, *Nous*, 45(3):409–442, (2011); Sarah-Jane Leslie, “The Original Sin of Cognition: Fear, Prejudice, and Generalization”, *The Journal of Philosophy*, 114(8):393–421 (2017).

⁴⁹ See Leslie, ‘Original Sin’.

⁵⁰ Jennifer Saul, ‘Are Generics Especially Pernicious?’ *Inquiry* (2017): 1–18.

There is a deep resonance between these generics and the social project of #BelieveWomen. Still, generics are still truth-evaluable declaratives, whereas ‘#BelieveWomen’ is an imperative. It makes a demand, not an assertion. To read it as a generic we would have to take it as elliptical for something like “Women are trustworthy”. That aside, one wonders whether a project to *establish* a positive generic of the sort imagined could be expected to succeed in improving the uptake of women’s testimony. There’s a risk, when trying to introduce a positive generic, that you’ll actually only make the negative stereotype that you’re trying to displace more salient, and so entrench *it* instead. Imagine trying to establish ‘men aren’t violent’ or ‘pitbulls are safe’ as generics. If the background context is not already favorable---and especially if it is already *unfavorable*, if it already encodes a pernicious generic---then these are apt to instead cement the opposite belief.

This is why rebranding efforts often circumvent the realm of the rational or truth-evaluable entirely. If there is already widespread acceptance of the pernicious generic ‘pitbulls are dangerous’, you don’t fix that by telling people ‘pitbulls are safe’. You try to change the frame people use to think about pitbulls, with a little branding.⁵¹ You make them seem cute and cuddly by calling them ‘pibbles’. Importantly, this gets in at the level *below* beliefs: it targets the affective frames with which we approach the world, which we use to sort our inputs, to resolve ambiguity, to decide what we are seeing, hearing, etc., and how to feel about it.⁵² I want to suggest that #BelieveWomen is precisely *this* kind of campaign, in which case it’s not (primarily) an epistemic directive for individuals. It seeks to restructure cognitive architecture at a deeper level, to destabilize an affective frame that represents women as prone to making false or specious allegations of sexual assault. The hashtag surfaced in response to a particular pattern: the pattern of *responding* to women’s first-person testimony by implying that we have no evidence, raising vague concerns about false allegations, and pointing to a (very) small number of striking instances of specious accusations.

If Jones (2019) is right that attitudes like fear and suspicion undermine trust—and I think she is—then to actually address the widespread distrust of women’s testimony we need a project in social marketing, aimed to undercut an *already existing* and widespread striking property generic: that women make false or spurious accusations of sexual assault. The target is affective; it does not instruct hearers to believe what is asserted by all, most, or even any women in particular. It does not even instruct them to try to increase their estimates of women’s credibility. Rather, it urges us to interrogate (and abandon) the undeserved affective skeptical frame with which we greet women’s testimony in general and about sexual assault especially. That will require attending carefully to the patterns that manifest in our conversations about sexual assault, how we talk about women’s testimony, and what questions we are disposed to ask when the topic comes up.⁵³

⁵¹ Discussed in, among other places, the *New York Times* (<https://www.nytimes.com/2019/01/04/style/pit-bull-pibble.html>) Marisa Meltzer, “The Pit Bull Gets a Rebrand: Who needs a Goldendoodle when you’ve got a “Pibble”?” Jan 4 2019.

⁵² For more on affective frames, and how they are more explanatory of action and attitude than underlying propositional attitudes like belief (even belief in a generic), see Camp & Flores “‘That’s All You Really Are’: Social Trouble With Cognitive Essentialism”, (unpublished ms, on file with the author).

⁵³ The interpretation I have offered of the #BelieveWomen campaign is supported by recurring editorials that stress the importance of taking women’s testimony seriously, and emphasize that false accusations are much rarer—and much more obvious—than they are believed to be. For instance, in 2017 Sady Doyle emphasized

Wrapping Up

So, to summarize: I've argued that on a plausible understanding of what *being believed* as a speaker amounts to—namely, being trusted—there is no conflict between appropriately believing women who testify and preserving the presumption of innocence and due process. Hence I have suggested that we do in fact have an ethical obligation to believe women, as well as an ethical obligation to avoid believing badly of someone on inadequate evidence. Both obligations are grounded in our duty to avoid imposing the harms and risks that would result from rejecting testimony or making errors arising from our mistaken beliefs. So, while I agree with Ferzan and Atwood that we shouldn't blindly defer to or 'side with' women, there is a plausible reading of #BelieveWomen as an imperative under which it reminds us we ought to have *reasonable trust* in women's testimony. This requires accepting their testimony as evidence in favor of the proposition they assert, and changing our orientation toward error possibilities—in particular, it requires that we avoid centering our inquiry on *whether the testifier is misleading*, focusing instead on what the available evidence supports about the events she recounts.

Though I have spent most of this paper working to demonstrate that we can *believe speakers* without adopting an attitude of full confidence in the content of their assertions, I do not think that #BelieveWomen actually is best understood as calling for a reform of our *individual* epistemic responses to women's testimony. Rather, the correction must be a project in *social* epistemology. We don't just use testimony to vouch for facts; we outsource the work of deciding whom to trust, and how much. We vouch for *people* as trustworthy testifiers within domains of expertise: trust him for fashion advice, her for career advice, etc. In large social networks, we vouch for *types* of people (don't trust politicians or used car salesmen; do trust scientists, etc). When this works well, it is a cognitively efficient way of dividing epistemic labor. But it can go—and has gone—badly. A group of agents have been branded as characteristically unreliable in their testimony about an area of policy concern that is testimony-dependent. It's not that people believe that *all* women lie, no more than that *all* sharks attack bathers. But just a few salient bad cases, blown to mythic proportions and deeply embedded in cultural tropes, are sufficient to brand all members of the kind as untrustworthy, ineligible to vouch for the facts in that domain.

that “The phrase is “believe women”—meaning, don't assume women as a gender are especially deceptive or vindictive, and recognize that false allegations are less common than real ones.” (Doyle, ‘Despite what you may have heard, “Believe Women” has never meant “ignore facts”’ *Elle* Nov 29 2017

(<https://www.elle.com/culture/career-politics/a13977980/me-too-movement-false-accusations-believe-women/>). More recently, Monica Hess critiqued the substitution of #BelieveAllWomen for the original tag, noting that ““Believe women” was a reminder, not an absolute rule; the beginning of a process, not an end. It was flexible enough to apply to various contexts: Believe women . . . enough to seriously investigate their claims. Believe women . . . when they tell you about pervasive indignities — catcalling, leering — that happen to them and their friends when you're not around.” (Hess, ‘Believe Women as a slogan; Believe All Women is a strawman’, *Washington Post* May 11 2020. https://www.washingtonpost.com/lifestyle/style/believe-women-was-a-slogan-believe-all-women-is-a-strawman/2020/05/11/6a3ff590-9314-11ea-9f5e-56d8239bf9ad_story.html). Susan Finaldi made the comparison even more explicitly: “the preferred hashtag of the #MeToo movement is #BelieveWomen. It's different without the “all.” Believing women is simply the rejoinder to the ancient practice of #DoubtWomen.” (Finaldi, Opinion, ‘#BelieveAllWomen is a Right Wing trap’, *The New York Times* May 18 2020. <https://www.nytimes.com/2020/05/18/opinion/tara-reade-believe-all-women.html>).

When she testified, Blasey-Ford exemplified many of the properties we use to indicate trustworthiness: she leads a conservative lifestyle, in general obeys social norms, has advanced degrees, has a stable income, spoke in a calm and organized matter, answered questions directly, and *still* the first questions asked (and the majority of subsequent public discussion) focused on whether *she* was trustworthy. What the hashtag calls attention to is that our social system of figuring out whom to trust is broken: it unfairly excludes women as sources of reliable epistemic testimony, *especially* on the topics where their testimony is crucial to understanding the scope of the problem.

As I said at the outset, this is a reconstructive (and perhaps partly revisionist) gloss of the hashtag; I expect plenty of users *do* mean things by #BelieveWomen which I have ruled out as implausible. But I think there is value in carefully tracing what it could mean, at its *best*, and outlining the nature of the underlying social problem which provoked it. Of course, this is all a bit complicated for twitter, so it's understandable if people make a hash of it when limited to 280 characters (pun intended).