



<http://social-epistemology.com>  
ISSN: 2471-9560

Response to Joshua Mugg's "How Not to Deal with the Tragic Dilemma"

Charles Lassiter, Gonzaga University, [charles.lassiter@gmail.com](mailto:charles.lassiter@gmail.com)

---

Lassiter, Charles. 2020. "Response to Joshua Mugg's 'How Not to Deal with the Tragic Dilemma'." *Social Epistemology Review and Reply Collective* 9 (3): 44-49. <https://wp.me/p1Bfg0-4Us>.

Before getting to my response to Joshua Mugg’s paper “How Not to Deal with the Tragic Dilemma” (2020), I’d like to thank Jim Collier for his invitation to continue the discussion in *SERRC*. I’ve often enjoyed the back-and-forth found in *SERRC* and I’m just tickled pink to contribute. Also, while Nathan Ballantyne and I co-wrote the paper (2017) to which Mugg replies, the response here is my own. I’ll make clear when I’m referencing arguments that Ballantyne and I give in our paper and when I’m giving my own views.

### **Localism**

Some time ago, Tamar Gendler (2011) brought analytic epistemologists’ attention to a problem. As a result of living in a society structured by race and sex, discriminatory stereotypes worm their way into our patterns of thinking: Black men are violent, women are bad at math, poor people are lazy, east Asians are good at math. Sometimes we’re interested in forming beliefs based on social categories, but the existence of these stereotypes dooms us to an epistemic Gehenna. Suppose we encode the information (and also that we’re not overtly racist, sexist, poverty-phobic, etc). If that information remains implicit and subpersonal, then we have a kind of irrationality: we implicitly endorse what we explicitly disavow. If that information becomes explicit and personal, then we spend mental energy to pay the cognitive costs of subverting the discriminatory effects on our thinking.

But suppose we *don’t* encode that information. There are plenty of times when that information is useful. For example, suppose we want to assign insurance premiums that reflect the relevant level of risk. Self-identified liberals are inclined to deflate premiums upon learning that they would affect largely Black communities (Tetlock et al. 2000). This is an instance of base-rate neglect: failure to account for background rates in estimating likelihoods. But ignoring base-rates is an epistemic failure *par excellence*. So encode or not, you’ll end up falling short of the demands of rationality.

Mugg points out that there is a parallel version of the dilemma. We could be rational, using the available information, but that can further disadvantage an already-marginalized population. Or we could refrain from using the relevant information, thereby helping out populations that have experienced harms historically, but then we fail to abide by relevant epistemic norms. Enact justice or rationality; but you can’t have both. He calls the dilemma in Gendler’s paper the ‘epistemic-epistemic dilemma’ and the latter the ‘epistemic-moral dilemma.’

Ballantyne and I argued that we can avoid the epistemic-epistemic dilemma by paying attention to *activation* of information and not simply the *encoding* of it. If the stereotype information isn’t made salient, then it won’t affect the token reasoning process. We suggested this is possible by means of what we called ‘localism’ (and what Mugg calls ‘contextualism’)<sup>1</sup>: local environments can be assembled in such a way as to mitigate the effects of implicit biases on a token instance of reasoning. If the decision-making

---

<sup>1</sup> I’m going to use ‘localism’ over ‘contextualism.’ The latter carries a lot of baggage.

environment fails to trigger the stereotype—or mitigates the effects of the stereotype being triggered by other stimuli—then it’s possible to avoid Gendler’s dilemma.

### Mugg’s Objections and Some Replies

Mugg’s most recent contribution (2020) suggests that this view is wrong. He offers three arguments. I’ll sketch each and then offer a reply. The page numbers are in reference to the online version of the paper.

First, our view “overestimates the extent to which we can control access to the stereotypical information” (4): the stereotype can be activated in circumstances where we least suspect it. To make matters worse, not only do we not *presently* know all relevant triggers for all stereotypes, we have good reason to think that there will be an *ever-growing list* of triggers for stereotypes! Mugg doesn’t make this second point, but it’s (i) consistent with what he says, (ii) suggested by our best social science (that cultures are dynamic and evolving), and (iii) strengthens his argument. He says (5): “[t]o avoid the tragic dilemma by blocking access, I need to know how to prevent epistemologically useful but racially tainted information from manifesting in *every* case” (emphasis his).

In reply, this seems to confuse what localism is committed to. Localism says that one can construct local conditions to mitigate the effects of stereotype information. We should expect our capacity to construct the local environments to grow as our knowledge of triggers grows. There might be practical limits—people with full-time jobs and families and without PhD’s can’t reasonably be expected to keep up with the latest social psychology research<sup>2</sup>—but in principle, our trigger-prevention knowledge grows in parallel with our knowledge of triggers.

Mugg, however, predicts this response. He says that it’s not enough to have an appropriate strategy in the future. We need solutions now. But this is confusing: the solution localism provides is one that *can* be implemented now. The *APA Monitor*, for example, suggests that racial bias in police shootings can be reduced if officers get to know the people in the neighborhoods where they walk their beats (<https://www.apa.org/monitor/2016/12/cover-policing>). This is something police officers and their superiors might implement now. Here’s another: want to help reduce transgender prejudice? Engage others in a short conversation encouraging them to take the perspective of the transgender person (Broockman and Kalla 2016). These are things that people can do *now* to reduce the effects of pernicious stereotypes on thinking. And people can do these things now even while new information about stereotypes and biases continue to roll in.

One source of the disagreement might be that Mugg is thinking of localism as a one-and-done strategy: we figure out how to prevent triggering stereotypes and then we do the implementation. I think of localism as an on-going project: the more we learn about how to reduce prejudicial attitudes, the more we adjust our local environment. Ballantyne and I suggest as much when we talk about a “ratchet effect” (Tomasello 1999): cultural knowledge accumulates, which prevents slippage to previous cultural states. This is a difficult and

---

<sup>2</sup> Plus, who’s covering the costs of accessing journals?

arduous task, to be sure. But even Satan knows that “Long is the way and hard, that out of Hell leads up to light.”

Here’s Mugg’s second argument. Sometimes we have to decide whether stereotype information is relevant in some case. But merely considering the relevance can activate the stereotype, which cascades into the moral and epistemic problems described above. Mugg illustrates this point with a story. Spencer read somewhere that black diners typically tip less than white diners—an important fact to know in the US since servers depend almost entirely on their tips for their wages.<sup>3</sup> Spencer sees Jamal, a black diner, enter the restaurant and recalls that Jamal left a lower-than-average tip the last time he was in. Spencer now has to decide whether Jamal’s being black is relevant to his choice to be Jamal’s server (or to let someone else take Jamal’s table and get the lower tip).

In reply, there’s a conceptual confusion in Mugg’s argument. In describing the case, it sounds as if Spencer is thinking to himself whether Jamal’s being black is relevant to his future actions, as though mentally rehearsing, “I know Jamal is black and that black diners tend to give lower tips. Is Jamal’s being black relevant to my choice of being his server or not?” Mugg’s argument and vignette suggest that considering whether information about race is relevant is a *personal* as opposed to a *subpersonal* process. This seems wrong. Determinations of relevance for cases of implicit bias are made at the subpersonal rather than personal level. Otherwise, it wouldn’t be *implicit* bias. It’s not that *Spencer* considers whether race is relevant but rather subpersonal processes in Spencer that would deploy the stereotype information or not.

In that case, Mugg’s second argument reduces to this: in some cases, information about race is tagged as relevant by subpersonal processes. But this *just is* the very problem of implicit bias! Mugg’s second objection doesn’t shed light on any extant problems. So if localism is a viable strategy in the face of the epistemic-epistemic dilemma, then it’s a viable strategy for Spencer’s case.

Here’s Mugg’s third argument. He identifies a new way to reach the tragic dilemma, with Antony’s (2016) insights as the launching pad. Antony observes that (1) biases are epistemically useful (cf. Gigerenzer 2008) and (2) some markers of properties-of-interest are markers that have been used to marginalize. These facts give us another tragic dilemma. Mugg illustrates with an example of predicting who will win a math contest: A is a member of race<sub>1</sub> and B a member of race<sub>2</sub>. Race<sub>1</sub> is stereotypically bad at math; race<sub>2</sub> is stereotypically good. If these stereotypes are made salient to A and B, then A will likely perform worse than if the stereotypes had not been made salient. B will likely perform better.

---

<sup>3</sup> In only four US states do waiters and waitresses earn an hourly wage of at least \$10 dollars. In 23 states is there a law requiring that salary *plus* tips results in an average hourly wage that can range from \$7.50/hour (Nevada) to \$12.00/hour (Colorado). In 21 states, there is no law requiring servers to make a minimum salary. Live in Texas or North Carolina? You’re making \$2.13/hr. Spencer’s decision seems more important if he’s living in Texas than Washington. [https://en.wikipedia.org/wiki/Tipped\\_wage](https://en.wikipedia.org/wiki/Tipped_wage)

Mugg argues that we ought<sub>E</sub> take stereotype-threat into account when forming a belief about who will win the math competition. Namely, I ought to predict that B will do better than A if the stereotypes are made salient to each. But if we do this, then we're in moral trouble, for stereotype-threat is a "self-fulfilling prophecy." While we don't cause B to perform better than A, our believing "reinforce[s] the structures in place," creating a feedback loop in which race<sub>2</sub> continues to be better at math than race<sub>1</sub>.

But if we don't take stereotype-threat into account? Then we're failing to be epistemically rational, since both the stereotype and stereotype-threat are epistemically relevant in this case.<sup>4</sup> Call this the 'New Tragic Dilemma.'

What does this mean for localism? Mugg argues that if localism is to work, there can't be any cases where epistemic and moral oughts diverge. But the New Tragic Dilemma is such a case. Therefore, localism isn't going to work.

Notice an important assumption needed for the anti-localist argument to work, found in footnote 5: if localism is a solution to the epistemic-epistemic dilemma, then it's also a solution to the epistemic-moral dilemma. The conditional is needed because localism is a solution to the epistemic-epistemic dilemma; the New Tragic Dilemma is an epistemic-moral dilemma. So if the New Tragic Dilemma is to be a problem for localism, then localism has to solve the kind of problem realized in the New Tragic Dilemma. Now the conditional seems *prima facie* plausible, but digging a bit shows the plausibility to be superficial. Why? Localism is the wrong sort of solution for the problem manifested in the New Tragic Dilemma. This isn't a strike against localism but an acknowledgement of problems for which it isn't suited.

Let me explain.

One of the prerequisites for adopting a localist solution to Gendler's dilemma is this: questions about competing values are settled and we're interested in enabling agents to act in accordance with those values. Call these 'Clear Values' cases. For instance, when we're interested in arriving at true beliefs, come what may, we have a Clear Values case. If our job is to assess risk for investing in a neighborhood, we want as accurate a picture as possible of crime rates and weather-related catastrophes. The value to be manifested is clear and uncontested: we are to increase our stock of true beliefs.

But the world isn't always so kind. We face cases where different values can't be jointly satisfied. Call these 'Competing Values' cases. Consider, for example, poverty and undergraduate admissions. Poorer folks simply don't have the same time or financial resources in applying to undergraduate programs as their wealthier peers: SAT prep is time and money, applications are on average \$77 each<sup>5</sup>, and revising writing samples takes time that students with full- or part-time jobs don't have. Compounding matters, this is just at the

---

<sup>4</sup> Observe that the case is under-described. Mugg says that we only know the races of A and B. But if I ought<sub>E</sub> to predict B will outperform A, then I would need to know that the relevant stereotypes had been activated for A and B. Otherwise, the only information I have is that each participant is a member of a race that is reputed to be better or worse at math. I don't think this is hugely worrisome, but it might make one skeptical about the possibility of rationally predicting an outcome with such a poverty of evidence.

<sup>5</sup> In case you're curious, Stanford is the highest at \$90: <https://www.usnews.com/education/best-colleges/the-short-list-college/articles/colleges-with-the-highest-application-fees>.

end-stage of applying to undergrad. Gifted and talented programs all through the pre-college years, for instance, are better at picking out wealth than intelligence (Grissom, Redding, and Bleiburg 2019). So imagine that we're on the admissions committee for an undergraduate program. We see an application that's borderline, but reading the student's personal statement we see that they've come from a life of poverty and have had to work at least half-time since they were 15 (often getting paid under-the-table). Higher poverty negatively correlates with higher academic achievement (see Lacour and Tissington 2011 for review). So we should expect that this student would perform poorly in a more rigorous academic environment; but also, an undergraduate degree often confers a higher salary, which can help lift this student (and any family they might have) out of poverty. What the data predict of this student and what we morally want for them are in tension. What do we do?

This is a Competing Values case, and localism is of little help. Do we manipulate our local environment so that we pay attention to lifting students out of poverty or so that we make the most rational decision possible given the evidence? In this case, we can't jointly satisfy moral and epistemic demands. But this has nothing to do with localism and implicit biases and everything to do with values whose joint satisfaction is near-impossible in this case.

But why can't localism help with Competing Values cases? Localism is an empirically-informed strategy for helping agents manifest some value. Importantly, it is *not* a strategy for navigating competing values. How on earth could it be? Figuring out what matters for us is the task of deep personal and philosophical reflection: is our obligation to predict accurately student success on the basis of the information we have or is it to offer poor students an opportunity to better their situations?

We could imagine a Super Localism, in which I construct my environment to remind me to make choices manifesting one set of values over another. I might fill my office with pictures of Descartes, each captioned with "Epistemology is First Philosophy" or perhaps of Levinas captioned with "Ethics is First Philosophy." But *even then* the normative matter of which values to privilege would have been settled. Super Localism would be just as ineffective as localism in navigating Competing Values cases.

There are, of course, some cases in which moral and epistemic values converge on the same course of action: suppose our hypothetical, poor student was actually very promising. In that case it's epistemically and morally right to offer the student admission; and if the committee were biased against the poor student because he's poor, then they would be epistemically and morally in the wrong. But that's not a case of *Competing Values*. Call these 'Aligning Values' cases. In a Kantian mood, we might wonder which value was motivating us or which *would* motivate us if circumstances were different. But we can thank our lucky stars in Aligning Values cases that we don't have to make that choice.

I previously said the conditional "if localism is a solution to the epistemic-epistemic dilemma, then it's also a solution to the epistemic-moral dilemma" is false. Now we're in a position to see why. Localism can, and ought to, provide a solution to the epistemic-epistemic dilemma, insofar as such dilemmas are Clear Values cases. But localism is not, and

ought not be, a solution to epistemic-moral dilemmas. Every tool has its limits; a hammer's no good for a leaky pipe, nor should it be.

## Conclusion

I've argued that Mugg's paper doesn't succeed in undermining localism as a strategy to avoid Gendler's epistemic-epistemic dilemma and that it's the wrong tool for making headway on epistemic-moral dilemmas. Localism is a strategy that works when we've already settled on the underlying values.

One complication is when, within some normative domain and for some problem, there are competing values. This is common in moral philosophy: respecting autonomy or pursuing the best ends—then they don't coincide, which do we pursue? In epistemology, a similar dilemma is: when we have to choose, do we opt for believing only truths or avoiding all falsehoods as the overriding value? Localism is silent on these issues too.

So what *is* localism good for? Ballantyne and I argued that it's good medicine for the epistemic pessimism served up by Gendler and now Mugg. We ended our paper on a melioristic note: understanding how the environment contributes to our cognitive processes empowers us to remake our epistemic worlds for the better. In a society pockmarked by social stratification, localism points a way out of Gehenna and into the light.

## References

- Antony, Louise. 2016. "Bias: Friend or Foe? Reflections on Saulish Skepticism" In *Implicit Bias and Philosophy* edited by Michael Brownstein and Jennifer Saul, 157-190. New York: Oxford University Press.
- Broockman, David, and Joshua Kalla. 2016. "Durably Reducing Transphobia: A Field Experiment on Door-to-Door Canvassing." *Science* 352 (6282): 220-224.
- Gendler, Tamar S. 2011. "On the Epistemic Costs of Implicit Bias." *Philosophical Studies* 156 (1): 33-63
- Gigerenzer, Gerd. 2008. *Rationality for Mortals: How People Cope With Uncertainty*. Oxford University Press.
- Grissom, Jason A., Christopher Redding, and Joshua F. Bleiberg. 2019. "Money Over Merit? Socioeconomic Gaps in Receipt of Gifted Services." *Harvard Educational Review* 89 (3): 337-369.
- Lacour, Misty, and Laura D. Tislington. 2011. "The Effects of Poverty on Academic Achievement." *Educational Research and Reviews* 6 (7): 522-527.
- Lassiter, Charles and Nathan Ballantyne. 2017. "Implicit Racial Bias and Epistemic Pessimism." *Philosophical Psychology* 30 (1-2): 79-101.
- Mugg, Joshua. 2020. "How Not to Deal with the Tragic Dilemma." *Social Epistemology* 1-12 (online). <https://doi.org/10.1080/02691728.2019.1705935>.
- Tetlock, Philip E., Ori V. Kristel, S. Beth Elson, Melanie C. Green, and Jennifer S. Lerner. 2000. "The Psychology of the Unthinkable: Taboo Trade-Offs, Forbidden Base Rates, and Heretical Counterfactuals." *Journal of Personality and Social Psychology* 78: 853-870.