



Short Communication

Deliberately prejudiced self-driving vehicles elicit the most outrage

Julian De Freitas^{*}, Mina Cikara

Harvard University, United States of America



ARTICLE INFO

Keywords:

Moral judgment
Autonomous vehicles
Driverless policy
Moral outrage

ABSTRACT

Should self-driving vehicles be prejudiced, e.g., deliberately harm the elderly over young children? When people make such forced-choices on the vehicle's behalf, they exhibit systematic preferences (e.g., favor young children), yet when their options are unconstrained they favor egalitarianism. So, which of these response patterns should guide AV programming and policy? We argue that this debate is missing the public reaction most likely to threaten the industry's life-saving potential: moral outrage. We find that people are more outraged by AVs that kill discriminately than indiscriminately. Crucially, they are even more outraged by an AV that *deliberately kills a less preferred group* (e.g., an elderly person over a child) than by one that *indiscriminately kills a more preferred group* (e.g., a child). Thus, at least insofar as the public is concerned, there may be more reason to depict and program AVs as egalitarian.

1. Introduction

Each year, we kill 1.25 million people and injure 20 million more in car accidents, 90% of which are caused by human error (Singh, 2015). Self-driving cars, or autonomous vehicles (AVs), are projected to prevent these deaths (Fleetwood, 2017) yet one of the biggest hurdles in the way of this future is getting AVs to behave ethically. How should we do this? One popular proposal is to inform AV policy with crowdsourced responses to so-called 'driverless dilemmas', in which respondents read hypothetical scenarios involving an AV that is forced to kill one social group versus another (e.g., a child or an elderly person), and make a choice on the AV's behalf (Awad et al., 2018). These experiments reveal several systematically biased social preferences in people's choices, such as preferring to save women over men, executives over the homeless, athletes over overweight individuals, the young over the old, and the lawful over the unlawful. The revelation of these social preferences is not new (Cikara, Farnsworth, Harris, & Fiske, 2010), but considering whether they are relevant to AV policy has become a matter of great interest to some social scientists and philosophers.

2. Discrepancies between results with forced-choice versus free-choice paradigms in AV dilemmas

Driverless dilemmas already assume that AVs may need to take into account each person's social group in order to decide whom to save. For this reason, participants in these experiments are *forced* to choose

between members of different social groups. Yet a forced-choice question places a demand on participants to indicate a preference, ignoring the possibility that they actually prefer that AVs not be programmed with any social preferences in the first place. Supporting this conclusion, recent work finds that fewer than 20% of people think it is a good idea for AVs to discriminate between social groups in such dilemmas (De Freitas, Anthony, Censi, & Alvarez, 2020), and another study reveals that most people choose to treat social groups equally (i.e., have the AV choose at random whom to save) if they are provided the option to do so (Bigman & Gray, 2020).

This presents a new conundrum: should we allow AV policy to be influenced by people's forced-choice preferences, or their stated preferences for equality? The answer depends on the goal.

Inevitably, AVs will end up killing some people that belong to the 'preferred' social groups, raising the question of how people will react when they do, and what implications that reaction will have for the AV industry and associated policy.

3. Arbitrating between equality and group-biased social preference with moral outrage

One relevant public reaction that has not been measured is *moral outrage* (Tetlock, 2003), a powerful emotion that plays a key role in coordinating multiple parties against a wrongdoer, e.g., via public protests or shaming on social media (Crockett, 2017; De Freitas, Thomas, DeScioli, & Pinker, 2016; Spring, Cameron, & Cikara, 2018). If

^{*} Corresponding author at: Department of Psychology, Harvard University, 33 Kirkland Street, United States of America.

E-mail address: defreitas@g.harvard.edu (J. De Freitas).

moral outrage surrounding AV-related harm is too high, then it could catalyze collective action that stunts or even halts the growth of the industry. Given that AVs are projected to save millions of lives per year, such delays would constitute a significant public health concern.

Returning to our conundrum, if people are more outraged by AVs that harm socially preferred groups than unpreferred ones, then this indicates that people’s underlying preferences might be worth considering in discussions of AV policy (even if people express egalitarian preferences when given the opportunity). Alternatively, if people are more outraged by AVs that deliberately harm in a prejudiced fashion rather than indiscriminately, then this indicates that policy-makers can ignore people’s social preferences. In fact, we would have a particularly strong reason to ignore these social preferences if it turns out that people are even more outraged when an AV deliberately harms an ‘unpreferred’ group than indiscriminately harms a ‘preferred’ group, e.g., deliberately kills an elderly person versus randomly kills a child.

4. To minimize public outrage against AVs, program equality instead of group-biased social preference

In order to resolve this conundrum, we told 826 participants about an AV that is programmed to either deliberately or randomly kill a ‘preferred’ or ‘unpreferred’ social group when faced with a tradeoff between the two (e.g., young vs. old person). We focused on social categories that elicited conflicting preferences (group-biased preference vs. egalitarian) depending on whether the choices in previous studies were forced or unconstrained (Awad et al., 2018; Bigman & Gray, 2020; De Freitas, Anthony, et al., 2020): age, gender, socioeconomic status, physical fitness, and lawfulness. Each participant read just one of the possible conditions (e.g., about an AV that deliberately kills an elderly person rather than a child). They then reported how much they blamed each party (the human who owns the AV, the AV itself, and the AV manufacturer), how much they were outraged by the manufacturer (Jordan & Rand, 2020), how willing they were to take collective action against it (Ford, Feinberg, Lam, Mauss, & John, 2018), and how worried they would be for their own safety and that of others. We were most

interested in respondents’ attitudes toward the AV manufacturer, given that coordinating against manufacturers could stunt the growth of the industry. We analyzed each outcome measure in a mixed effects linear regression model, with AV programming (deliberate vs. random) and the target killed (‘preferred’ vs. ‘unpreferred’) as fixed effects, and social category of the target (age, gender, status, fitness, and lawfulness) as a random intercept.

Overwhelmingly, AVs that harmed deliberately elicited more negative reactions toward the manufacturer than AVs that harmed indiscriminately: people blamed the manufacturer more ($\beta = 19.76, p = .003$), and were more outraged ($\beta = 36.93, p < .001$) and willing to take collective action against it ($\beta = 10.72, p = .049$). They were also more likely to blame the human who purchased the AV ($\beta = 31.90, p < .001$), but not the AV itself ($\beta = 2.59, p = .761$)—presumably because they recognized that the AV’s behaviors were pre-programmed by the manufacturer, making the manufacturer the morally responsible entity worth targeting.

These effects were evident both when collapsing across social groups (Fig. 1), and at the level of each individual social category (Fig. 2). Further, interaction effects revealed that participants blamed the human vehicle owner more if the AV deliberately killed the preferred rather than unpreferred group, but not if it did so randomly. Similarly, although participants consistently exhibited greater outrage when the AV killed the preferred rather than unpreferred group (whether it did so deliberately or randomly), this difference was twice as large if the AV deliberately targeted a particular social group (for more details, see <https://osf.io/8mke6/>). Most important for addressing the question of how to program AVs, people were more outraged by AVs that deliberately killed the ‘unpreferred’ group than indiscriminately killed the ‘preferred’ group (Fig. 1, Table S1). For instance, they thought it was worse for an AV to deliberately kill an elderly person than randomly kill a child.

5. Conclusions

These findings offer a resolution to the previous deadlock (Awad

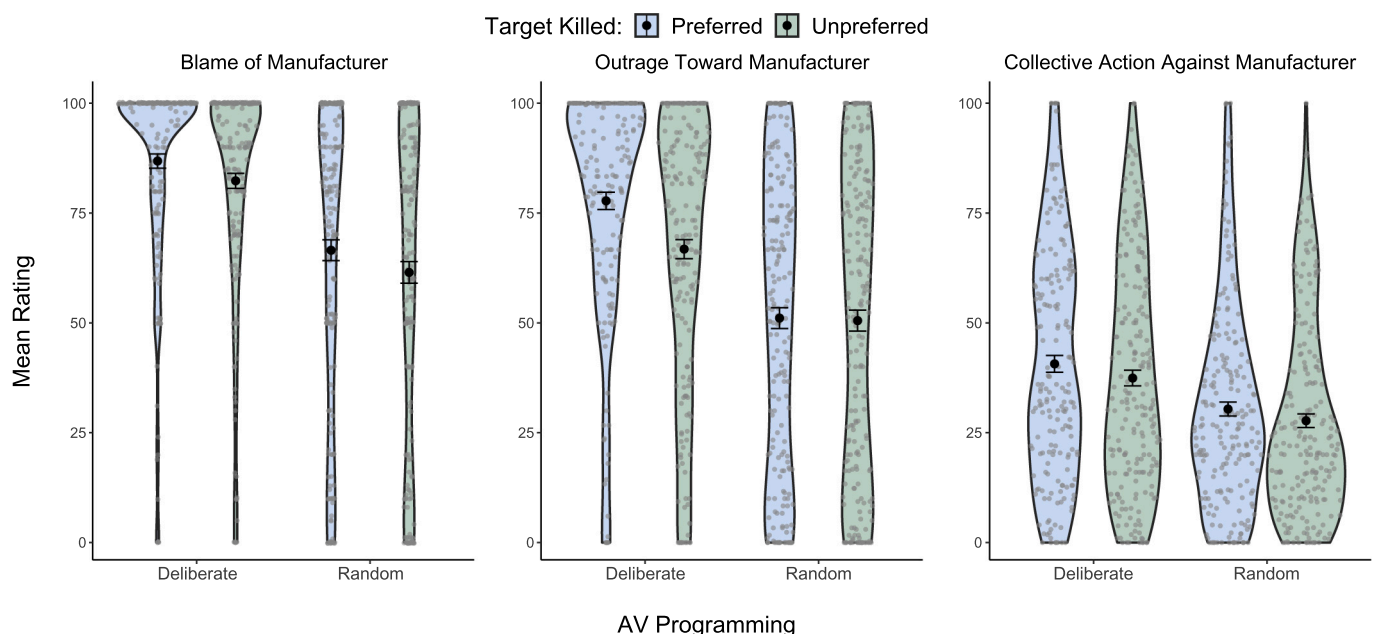


Fig. 1. Mean blame, outrage, and likelihood of taking collective action against the manufacturer, depending on AV programming (deliberate vs. random) and target killed (preferred vs. unpreferred).

Note. Across various measures, people cared more about whether an autonomous vehicle deliberately harmed a person than about the social group of the person. Notably, they were even more outraged by an AV that deliberately harmed an ‘unpreferred’ group than by one that indiscriminately harmed a ‘preferred’ group (the difference between the second and third bars). Error bars indicate standard error of the mean, and width of violin plots represent density of data points.

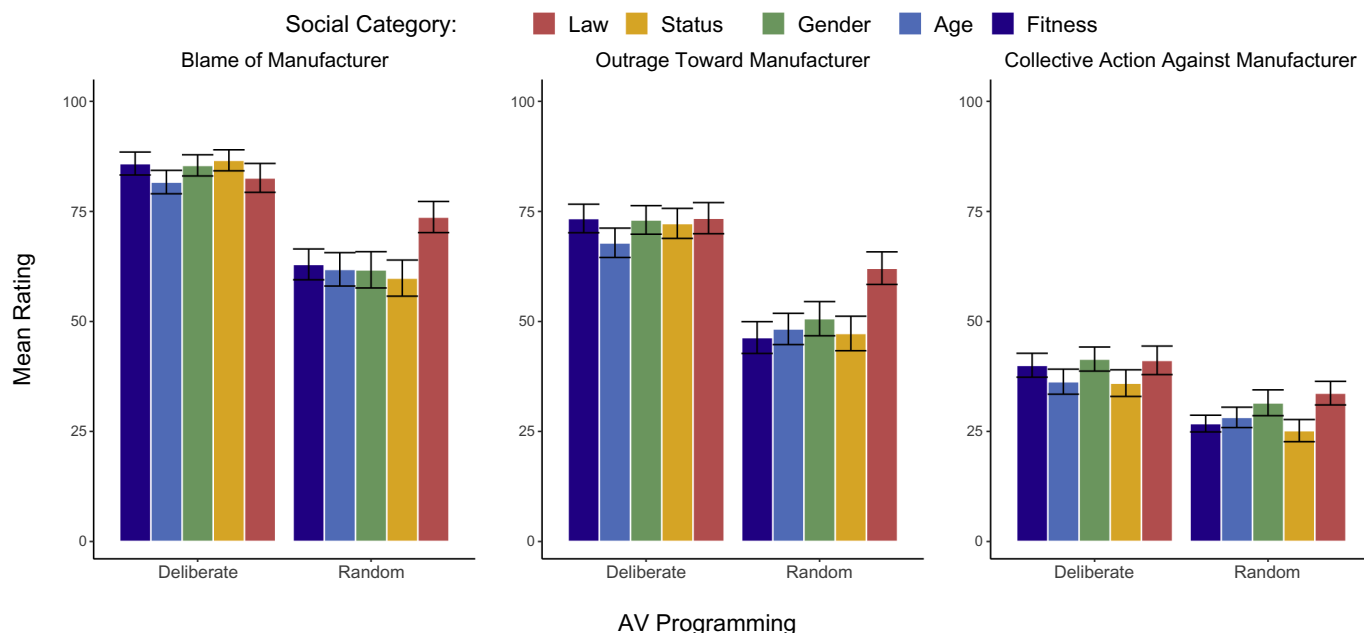


Fig. 2. The effect of AV programming (deliberate vs. random) on manufacturer-relevant outcome measures (blame, outrage, and likelihood of taking collective action) for five social categories.
Note. Error bars represent standard error of the mean.

et al., 2018; Bigman & Gray, 2020) regarding whether, from a public perception standpoint, we should favor AVs with egalitarian versus social preferences. We have suggested that one potential arbitrator is moral outrage, the moral emotion most likely to mobilize the kind of large-scale collective action that could threaten the growth of the AV industry. Provided the goal is to minimize public outrage in response to accidents involving AVs, our findings suggest that AVs should be portrayed as making egalitarian decisions (Huang, Greene, & Bazerman, 2019). Of course, it would only be ethical to portray them this way if this is, in fact, how they have been programmed, suggesting that there may also be an argument to actually program them to be egalitarian— provided it is not unsafe to do so.

Importantly, just because the current findings tilt the evidence on human preferences in favor of egalitarianism, does not necessarily mean that AV policy should be solely determined by public preferences or moral outrage specifically. Many acts that most of society now deems harmless have historically elicited outrage (e.g., racial integration, or same-sex marriage), and supreme laws like the constitution exist in large part to assure adherence to ‘eternal’ principles rather than the current sentiments of some portion of the population. In the case of AV regulation, data on public preferences should be considered in tandem with ethical and technical considerations (De Freitas, Censi, De Lillo, Anthony, & Frazzoli, 2020; Savulescu, Kahane, & Gyngell, 2019). With that said, even if the industry makes an informed decision to program AVs with some group-biased social preferences, it still needs to anticipate the social repercussions of such decisions. Outrage is arguably the most important reaction to keep in mind, given its industry-threatening potential. Finally, programming AVs without social preferences would not alone be a panacea for all ethical questions surrounding AV behavior, such as whether to program AVs to avoid all harm or just harm *at fault* (Censi et al., 2019); and whether to allow human drivers to switch control between themselves and the AV or require that they relinquish control to AVs once this is safer (Smith, 2020).

Follow up studies could test the extent to which the current results are affected (or not) by cultural, demographic, and personal factors, as when an elderly person has a personal stake in whether an AV has social preferences based on age. In line with Bigman and Gray (2020), we recruited a large sample from an online panel, rather than the millions of

global participants recruited by Awad et al. (2018). While increasing sample size alone would likely not change the effects, which are large, statistically well-powered (see sensitivity analyses in supplementary materials, Tables S3, S4), consistent across all five social categories, and in harmony with recent driverless dilemma studies (Bigman & Gray, 2020; De Freitas, Anthony, et al., 2020), future work should consider recruiting a representative sample. Another question is whether people’s reactions would change if the AV’s decisions were framed as deliberately ‘saving’, rather than harming, one social group at the expense of another.

Finally, it is tempting to dismiss the current issues altogether by pointing out that the very idea of driverless dilemmas is silly in the first place (De Freitas, Anthony, et al., 2020). Yet this has not stopped some people from arguing that policy-makers should take them seriously (Awad et al., 2018, 2020), and it is not unrealistic to think that AVs could be programmed with settings that more subtly favor certain social groups with clearly perceptible characteristics. Further, the media has the freedom to portray AVs in various ways, including as personified machines that have social preferences. Our findings suggest that, from a public perception standpoint, such discriminatory portrayals are likely to be industry-threatening, as are efforts to inform AV policy with crowdsourced discrimination.

CRedit authorship contribution statement

Julian De Freitas: Data curation, Formal analysis, Writing - original draft. **Mina Cikara:** Conceptualization, Writing - review & editing.

Acknowledgements

Thanks to Pechthida Kim and Nathan Hunt for assistance designing materials and collecting data, and Andrea Censi, Sam E. Anthony, and Steven Pinker for feedback.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2020.104555>.

References

- Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., ... Rahwan, I. (2018). The moral machine experiment. *Nature*, *563*(7729), 59–64.
- Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., ... Rahwan, I. (2020). Reply to: Life and death decisions of autonomous vehicles. *Nature*, *579*(7797), E3–E5.
- Bigman, Y. E., & Gray, K. (2020). Life and death decisions of autonomous vehicles. *Nature*, *579*(7797), E1–E2.
- Censi, A., Slutsky, K., Wongpiromsarn, T., Yershov, D., Pendleton, S., Fu, J., & Frazzoli, E. (2019). Liability, ethics, and culture-aware behavior specification using rulebooks. In *2019 international conference on robotics and automation (ICRA)* (pp. 8536–8542).
- Cikara, M., Farnsworth, R. A., Harris, L. T., & Fiske, S. T. (2010). On the wrong side of the trolley track: Neural correlates of relative social valuation. *Social Cognitive and Affective Neuroscience*, *5*(4), 404–413.
- Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, *1*(11), 769–771.
- De Freitas, J., Anthony, S. E., Censi, A., & Alvarez, G. (2020). *Doubting driverless dilemmas*. *Perspectives on psychological science*.
- De Freitas, J., Censi, A., De Lillo, L., Anthony, S. E., & Frazzoli, E. (2020). *From driverless dilemmas to more practical ethics tests for autonomous vehicles*. PsyArXiv. Retrieved from <https://psyarxiv.com/ypbve>.
- De Freitas, J., Thomas, K., DeScioli, P., & Pinker, S. (2016). Recursive mentalizing and common knowledge in the bystander effect. *Journal of Experimental Psychology: General*, *145*(5), 621–629.
- Fleetwood, J. (2017). Public health, ethics, and autonomous vehicles. *American Journal of Public Health*, *107*(4), 532–537.
- Ford, B. Q., Feinberg, M., Lam, P., Mauss, I. B., & John, O. P. (2018). Using reappraisal to regulate negative emotion after the 2016 US presidential election: Does emotion regulation trump political action? *Journal of Personality and Social Psychology*, *117*(5), 998–1015.
- Huang, K., Greene, J. D., & Bazerman, M. (2019). Veil-of-ignorance reasoning favors the greater good. *Proceedings of the National Academy of Sciences*, *116*(48), 23989–23995.
- Jordan, J. J., & Rand, D. G. (2020). Signaling when no one is watching: A reputation heuristics account of outrage and punishment in one-shot anonymous interactions. *Journal of Personality and Social Psychology*, *118*(1).
- Savulescu, J., Kahane, G., & Gyngell, C. (2019). From public preferences to ethical policy. *Nature Human Behaviour*, *3*(12), 1241–1243.
- Singh, S. (2015). *Critical reasons for crashes investigated in the national motor vehicle crash causation survey*. National Highway Traffic Safety Administration. Retrieved from <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812115>.
- Smith, B. W. (2020). Ethics of artificial intelligence in transport. In M. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of ethics of artificial intelligence*.
- Spring, V. L., Cameron, C. D., & Cikara, M. (2018). The upside of outrage. *Trends in Cognitive Sciences*, *22*(12), 1067–1069.
- Tetlock, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in Cognitive Sciences*, *7*(7), 320–324.