

Building social science knowledge on public attitudes and autonomous weapons

Michael C. Horowitz (@mchorowitz), Associate Professor of Political Science and Associate Director of Perry World House, University of Pennsylvania

Last week, Charli Carpenter published an [important piece](#) advancing the conversation about public attitudes, public conscience, and autonomous weapons. Her post critiqued my recent article in *Research & Politics* on public opinion and autonomous weapons. As a former Duck contributor, I am excited to return and further the dialogue (for a longer version of this post that contains more detailed responses to some other parts of Carpenter's piece, go [here](#)). [JOSH INSERT A LINK HERE]

Carpenter's [path-breaking survey](#) on public opinion and autonomous weapons was the departure point for my *Research & Politics* article. She persuasively showed widespread public opposition to autonomous weapon systems (AWS), in principle. My research builds on hers, as I seek to find out how likely political frames would affect public opposition to autonomous weapons. Carpenter and I actually agree on a lot about public attitudes concerning autonomous weapons, and even about what my own research shows, though we have some disagreements about survey methodology.

My paper demonstrates what Carpenter says it shows: that public attitudes about AWS will vary depending on the “primes” or “frames” used to provide context for the discussion. In particular, when AWS are presented as more necessary, defined in several different ways, public support rises. This is consistent with decades of political science research on how elite framing of issues shapes public attitudes. It is different from measuring public attitudes in a vacuum, and I hope that my original paper did enough to highlight this distinction. As I state on [page 1](#), the article's takeaway is that “public opposition to autonomous weapons is contextual.” (additional quotes on pages 2 and 7 say similar things). This type of inquiry is straight from the standard playbook on public opinion survey and field experiments, as research by [Jamie Druckman](#), [Diana Mutz](#), and [others shows](#).

The use of experimental conditions to better understand important humanitarian topics includes studies on [drones](#), [climate change](#), [same-sex marriage](#), [race and tolerance](#), [poverty](#), [immigration](#), [health care](#), [and torture](#), among other topics (the links above go to research that uses experimental conditions to evaluate public attitudes; the torture pieces even build on the [Gronke et al.](#) research Carpenter cites). My results similarly show that public attitudes concerning autonomous weapons vary based on the frames used in the discussion.

I then [conclude](#) that “More directly, from a policy perspective, these results suggest that it is too early to argue that AWS violate the public conscience provision of the Martens Clause because of public opposition.” The logic behind my statement is that, since presenting AWS in a positive light leads to significantly more public support for their development, it means public opinion on this topic is still malleable. If context shapes attitudes in the case of autonomous weapons, as it does in so many other areas, then we need to be careful when asserting what public attitudes “are.” More research is necessary.

Thus, my paper did not make the claim that Carpenter objects most to – that I was directly measuring the public conscience. I was instead testing whether context shapes attitudes in the case of autonomous weapons by comparing experimental frames to an unframed condition, as scholars

do in many other areas. What Carpenter and I disagree about most, it seems, is actually survey methodology.

More broadly, highlighting context in surveys of public attitudes is important because, in the real world, as [Baum and Potter](#) show, public opinion on foreign policy issues does not exist in a vacuum. Individual citizens tend to know little about (and pay little attention to) foreign policy issues. Cleanly measuring publicly attitudes is also tricky, especially on an issue such as autonomous weapons, which are not a hot button issue for the American public. Survey respondents are likely to have low degrees of information, meaning their opinions may represent what [John Zaller](#) calls non-attitudes. How elites (whether government, NGO, or other) raise and discuss issues, due to their informational advantage over the public, shapes public attitudes. Public attitudes are essentially formed in a process that involves framing.

In an environment where survey respondents have low levels of information, one can measure the public's instinctive reaction with a "clean" survey and get a snapshot of public attitudes. However, public attitudes on many topics also change over time due to exposure or information (for example, [same-sex marriage](#)). Thus, experimental conditions that reflect potential elite argumentation show how real world debates may influence public attitudes (for another example on drones, see [this](#)). This is also one reason why my paper [tested](#) how the level of information respondents have about autonomous weapons influences their opinions (see Figure 2).

Moreover, there are limits to clean surveys with open-ended prompts, the type Carpenter prefers, as well. Nearly all survey questions, [including Carpenter's](#), also bias respondents in their framing. And asking people about autonomous weapons, for example, means the responses conflate survey respondent views on autonomy with survey respondent views on weapons. Someone that is opposed to all weapons, for example, might oppose the development of autonomous weapons not because they are autonomous, but because they are weapons.

Additionally, since many people have unformed opinions, open-ended prompts can lead to [response instability](#). A [lack of eloquence](#) or sufficient [frame of reference](#) to answer an open-ended question can complicate the results. In particular, moral attitudes are [often unconscious](#), making them particularly hard to articulate. Methodological pluralism is therefore helpful to get multiple takes on public attitudes.

It is also important to note that my study was not, as [Carpenter states](#), "biased in favor of killer robots" because the experimental conditions are "drawn purely from pro-AWS hypotheticals. The experimental conditions focus on issues such as protecting US troops and military effectiveness, with some respondents, for example, told that AWS are more effective than other weapons, and others told they are not more effective (all question wording, data, and replication information is available [here](#)).

It is true that I did not add a condition where AWS were definitively worse than other weapons, but ironically, it is because Carpenter's research persuasively showed baseline public opposition to AWS. I strongly suspect that conditions where AWS were definitively worse than other weapons would increase public opposition (though it is worth testing in future research). Thus, any bias was due to the strength of her initial survey.

Given Carpenter’s finding of public opposition to AWS, my experimental design attempted to test the scope of that opposition. After all, if frames where AWS were more effective than other weapons or would protect US troops did not lead to greater public support for AWS, that would be especially strong evidence about the level of public opposition. Instead, the results show that context matters.

The next step could be follow up research looking at how pro and anti AWS moral, humanitarian, and military frames influence public attitudes about AWS (as well as more representative surveys outside the United States, building on research by the [Open Roboethics Initiative](#), and surveys comparing attitudes about AWS to other weapons). Our understanding of public opinion in this area is still limited. Combined with the open-ended survey prompts that Carpenter prefers, this could further build knowledge. This is consistent with my [conclusion](#) that “[M]ore research is necessary to scope out how the context influences public support for and opposition to autonomous weapons.”

Conclusion

Carpenter [concludes](#) that “Horowitz’ own data on autonomous weapons proves his assertion wrong.” Since I only claim is that public opinion is context-dependent on the topic of autonomous weapons, and that Carpenter agrees with that point in another part of her post, I actually think we agree more than we disagree. Interestingly, my baseline results replicate, in part, Carpenter’s initial findings of baseline opposition to AWS, as she states. This should increase confidence in the experimental results, because it means that, even though the study utilized a convenience sample, not a representative population sample, it generated similar results to Carpenter’s survey.

My paper does not support or oppose a ban on autonomous weapons. It is social science research – survey results that attempt to better understand the conditions in which the very opposition Carpenter identifies may or may not hold.

If policymakers can present AWS in ways that make them seem more (or less) attractive and by doing so impact public support, this has policy and empirical social science implications. The bar for speaking for humanity by claiming any weapon violates (or does not violate) the global public conscience should be high. As stated above, if public opinion was not malleable even when respondents were presented with information suggesting the effectiveness and utility of AWS, it would say something unique about the level of public opposition. My results show, however, that as with many other issues, public opinion on AWS depends on the way the issue is presented.

For those who support a ban on “killer robots,” these results also have implications. There may be nascent public opposition to autonomous weapons, but political reality is fragile. Public conscience, of which public attitudes are a part (as [Peter Asaro](#) argues), might not be manifest in the moment societies decide whether to develop or use autonomous weapons, depending on the scenario. If the goal is to “ban killer robots,” therefore, my results show it is useful to better understand the interaction between elite framing and public attitudes.

Finally, Carpenter’s piece also raises crucial methodological questions. Do normative beliefs exist in a vacuum, and what is the best way to measure them? Carpenter prefers “clean” surveys with open-ended prompts for respondents to report why they have the attitudes they have. Yet open-ended surveys are not perfect, as I describe above. Moreover, as autonomous weapons become more politically salient, elite discourse will shape viewpoints more. It therefore makes sense to also

understand the way different approaches to explaining the issue of autonomous weapon systems may influence those public attitudes. Methodological pluralism is critical, therefore, to move the discussion forward and build scholarly knowledge about public attitudes on autonomous weapons.

Additional Thoughts

Since Carpenter raises several specific questions about my research, in addition to the broad points I just made, this additional section responds to some points not addressed elsewhere.

-Carpenter argues in several places that context-based surveys are “inappropriate” for measuring the public conscience, meaning we should not prime people at all. Carpenter’s point raises the question of how to measure the part of the public conscience reducible to public opinion (As Peter Asaro argues, it is critical to keep in mind that public attitudes are only one part of the public conscience).

Consistent with decades of research in multiple subfields of political science, using experimental conditions is not “constructing” public attitudes, as opposed to letting the public speak for itself; it is researching how different contexts shape public attitudes, as Diana Mutz shows. If the purpose of the research is investigating how those contexts shape attitudes, as mine was, survey experiments are appropriate. One benefit to survey experiments is that, by holding everything constant across conditions except for the specific wording you are testing, you can be confident that the difference between the conditions reflects what you are testing, not other potential wording bias. Experimental frames actually can be particularly important in the area of science communications.

-Carpenter also argues that the experimental conditions in my research do not concern “normative beliefs”, meaning they are irrelevant from the perspective of the public conscience provision of the Martens Clause.

I am not a lawyer, so I will generally leave this topic for legal scholars. However, there is no provision of the Martens Clause that actually states public conscience can only be measured through normative beliefs. Moreover, what counts as a normative belief for measuring the public conscience is complicated. Carpenter raises the example of some thinking autonomous weapons will protect civilians better than other weapons. Some might view that argument morally, but others might view the protection of civilians instrumentally, as part of the legitimacy of a conflict.

Moreover, determining the public conscience on the basis of normative considerations does not exclude cost-benefit analysis to draw those conclusions, unless one further reads the Martens Clause as implying a Kantian imperative to avoid making ethical determinations on the basis of the likely consequences of actions.

-For the purpose of clarity, in no way did I mean to imply that we should “dismiss” public conscience and the Martens Clause. Morality is, as Carpenter says, “a constraint against the untrammelled pursuit” of interests “at all costs.” That is precisely why it is important to understand the conditions in which public attitudes about autonomous weapons are likely to vary.