

Introducing AI to an Online Petition Platform Changed Outputs but not Outcomes

Corresponding Author: Ms Isabel Corpus

Version 0:

Decision Letter:

5th November 2025

Dear Ms Corpus,

Thank you once again for your manuscript, entitled "Introducing AI to an Online Petition Platform Changed Outputs but not Outcomes", and for your patience during the peer review process.

Your Article has now been evaluated by 3 referees. You will see from their comments copied below that, although they find your work of potential interest, they have raised quite substantial concerns. In light of these comments, we cannot accept the manuscript for publication, but would be interested in considering a revised version if you are willing and able to fully address reviewer and editorial concerns.

We hope you will find the referees' comments useful as you decide how to proceed. If you wish to submit a substantially revised manuscript, please bear in mind that we will be reluctant to approach the referees again in the absence of major revisions. We are committed to providing a fair and constructive peer-review process. Do not hesitate to contact us if there are specific requests from the reviewers that you believe are technically impossible or unlikely to yield a meaningful outcome.

In particular, Reviewer #2 raises concerns about the interpretation of the results, given potential unobserved effects of the introduction of the LLM on user behaviour or platform dynamics. It will be important to address these concerns, ideally with additional data.

Finally, your revised manuscript must comply fully with our editorial policies and formatting requirements. Failure to do so will result in your manuscript being returned to you, which will delay its consideration. To assist you in this process, I have attached a checklist that lists all of our requirements. If you have any questions about any of our policies or formatting, please don't hesitate to contact me.

If you wish to submit a suitably revised manuscript, we would hope to receive it within 4 months. I would be grateful if you could contact us as soon as possible if you foresee difficulties with meeting this target resubmission date.

With your revision, please:

- Include a "Response to the editors and reviewers" document detailing, point-by-point, how you addressed each editor and referee comment. If no action was taken to address a point, you must provide a compelling argument. When formatting this document, please respond to each reviewer comment individually, including the full text of the reviewer comment verbatim followed by your response to the individual point. This response will be used by the editors to evaluate your revision and sent back to the reviewers along with the revised manuscript.
- Highlight all changes made to your manuscript or provide us with a version that tracks changes.

Please use the link below to submit your revised manuscript and related files:

Link Redacted

Note: This URL links to your confidential home page and associated information about manuscripts you may have submitted, or that you are reviewing for us. If you wish to forward this email to co-authors, please delete the link to your homepage.

Thank you for the opportunity to review your work. Please do not hesitate to contact me if you have any questions or would like to discuss the required revisions further.

Sincerely,



REVIEWER COMMENTS:

Reviewer #1 (Remarks to the Author):

This paper investigates the effects of genAI integration on an international platform that aggregates proposals and facilitates democratic governance and inquiry. The authors use the plausibly exogenous introduction of ChatGPT, leveraging the staggered adoption of genAI in Australia and specifically the platform's integration of it in Australia.

The paper is very well done: it's clean, clear, and contains all the relevant parallel trends and related diagnostics. I don't think there's any reason to make this a long drawn-out process; the paper is effectively ready to go, as far as I can tell. The authors do a good job setting up the prospective mechanisms. I also appreciated with the DID estimator that you focused on different sub-geographies in Australia. The only two minor comments:

1. Can you talk more about the institutional reasons for why Australia had a later adoption?
2. To further parse the mechanisms, you could always stand up a survey on Prolific.co, but you provide good enough results already. You could, however, look at the text of the proposals and see if the words that are used are more dis-passionate or less specific (reflecting the homogenizing effects that genAI can have on proposals).

Great work!

Reviewer #2 (Remarks to the Author):

Review of "Introducing AI to an Online Petition Platform Changed Outputs but not Outcomes"

I read this paper with interest. It empirically analyzes a change in tools offered to petition writers on a prominent petition hosting platform, Change.org ("the platform"). The platform invites petitions from authors, then invites signatures and comments on petitions, and then invites owners to declare 'success' when petitions succeed. I learned from this paper that large majorities of petitions do not attract any comments or indicate success.

The platform integrated its petition writing tools with an early large language model in 2023. It first introduced the change in Australia before its other English-speaking markets. After ten weeks it extended the change to other markets. This is a promising setting for a quasi-experimental econometric application. In general I think authors' applications are detailed, competent, clear, and mostly avoid overinterpretation. They examine robustness to alternate approaches, all great attributes to have.

My primary questions are about aims, hypotheses and contribution. These are first-order questions whose answers remain unclear to me after reading the paper and supplementary information. Their answers may influence the measurements and econometric designs, or they may not. Without those answers I have limited information about the optimal research design. Therefore I will not comment extensively on the empirics.

1. Aims

The platform apparently tested the LLM integration in Australia before extending the integration to other countries. The criteria for such a test are left unstated. They may be simple, such as did the integration succeed technically? However, I strongly suspect they extended to user-related outcomes as well. We are not credibly informed about the platform's goals of the integration, nor how the test performed on those goals.

My naive guess is that the platform sought to increase the completion rate of draft petitions. Authors often might start a petition and then abandon it due to the writing cost. Offering writing assistance may help to increase the completion rate and thereby increase the number of petitions available on the platform. Or LLM integration might help to reverse a declining trend on the platform. Or perhaps the platform noticed a decline in petition writers' writing quality during the pandemic and thought the LLM integration could help to reverse that.

The authors have chosen some metrics that are partially aligned with these claims. It remains unclear what the platform intended or how it evaluated success.

2. Hypotheses.

The authors' causal logic appears to simplify to the following:

1. Writing quality increases petition success markers
2. LLM assistance increases writing quality
3. Therefore, LLM assistance should increase petition success markers

There are several weaknesses in this causal chain.

2a. The first claim remains unproven. Authors note in the text that writing quality predicts petition success markers, and they show this in SI Figure 17. However, this is not a test of a causal claim. It may easily be that people who write well also choose petition topics that are likely to succeed. Without the first link in the chain, we have no reason to believe that LLM assistance should increase petition success markers.

2b. Return to point 1. Suppose the primary purpose of LLM integration is to increase draft petition completion rates. If uncompleted petitions tended to be low quality, then LLM integration would reduce the average quality of petition on the platform. We then should expect petition success markers might even decrease on average. In fact this appears consistent with figure 1d. But it seems to reverse the hypotheses that the authors seem to maintain.

2c. Completed petitions compete in a ranking system for platform user attention. There is a default ranking system and users may sort or filter petitions in some limited ways. We do not know how petition writing quality affected the default ranking system, therefore we do not know how writing quality may change user attention to completed petitions. This is a very common issue in empirical work in platforms, that interventions often affect platform algorithm inputs before and then subsequently affect users, so it becomes difficult or impossible to separate the effect of the intervention on the ranking system from the effect of the intervention on the users. For example, this issue is the difference between “lift tests” and “AB tests” in Facebook advertising; the former is a true experiment, whereas the latter combines the effect of an ad on the platform algorithms with its effect on the exposed users. No general solution is available to this issue.

3. Contribution

Given the uncertainty about the platform’s original intention, the appropriate measures to evaluate effects, and the mechanism, I find it hard to know what I can draw from this article that could inform other settings.

I do not think it would be reasonable to believe that LLM assistance would generally worsen existing creators’ outcomes in other platforms.

I find it generally reasonable to think that LLM assistance might increase creation by increasing creation accessibility, while simultaneously worsening average creation quality. How those changes affect the resulting equilibria would further depend on how high-quality creators respond to low-quality creations, how high-quality creators integrate LLM within their own creations, how users respond to the changes, and how the platform itself responds (via, e.g., its search rankings or presentation algorithm).

Following comments are more minor relative to the foregoing

4. Intervention. I am unclear whether the LLM was tuned to predict petition writing features that led to successful petitions in the pre-LLM era. I am also unclear how the rapid recent innovations in LLM quality may affect the durability of the authors’ results. I think they should comment on this, even if it is only speculative.

5. If the authors have connections at the platform, I think their research would be much stronger if they could report an experiment that increases the barrier to petition creation. If it is “too easy” to create a petition, the platform may have too much noise, and the best petitions may be seen suboptimally infrequently. A possible intervention could be a delay, or even better, a random intervention in petition visibility.

I thank the authors for their research and I hope to read a future version of the article which clarifies these points. I encourage them to develop their research further and hope these comments may be helpful.

Reviewer #3 (Remarks to the Author):

This paper describes a difference-in-difference natural experiment on change.org to study the impact of in-platform AI access on petitions’ lexical features and petitions’ comment outcomes. Overall, I thought the paper was well-written, the study to be rigorously done, and the findings interesting. I recommend accept.

That said, I do have some suggestions for authors to consider:

1. I find the framing of AI access to be a bit overly negative. While I’m not surprised that AI did not improve petition signing, it doesn’t necessarily mean there is no value for AI access. What is not factored in this current study is the impact of AI access on petition-writing work. If it helped reduce some effort there, a positive spin of the results could be that the same outcome could be achieved with lower cost. Given that effort/cost is not factored in this work, it might be good to at least acknowledge that potential tradeoff.

2. The part where I find the most interesting is the findings on outcomes. However, one thing I think is valuable to note in the paper is that it is unclear how much of the observed difference is due to the impact of AI access on the quality/persuasiveness of individual petitions, versus the impact of AI access on a more macro-level, on the general perceptions of petitions on change.org. While I believe in the findings on the causal impact of AI access, the underlying mechanism is less clear. The paper seems to make the argument that it’s because of the impact on the petition. I’m not confident about that. I think a more macro-level, longer-term effect could help explain why the effects are stronger for %petitions with 10+ signatures than %petitions with 1+ comment when limiting that to 30 days, and the difference between first and second petitions for repeat petition writers when both are post-AI. I

don't quite agree with the authors' argument to discount "AI suspicion". Perhaps there's no explicit indication that any single petition was AI generated, but AI access is clearly communicated on change.org (at least in the screenshot provided). I think the paper can be more precise about this and acknowledge while some of the findings are clearly about the petitions themselves (e.g., lexical features and petition homogeneity), the effect of AI access on outcomes may have been impacted by different mechanisms – e.g., how AI access is changing the petition text, as well as how its introduction is shaping reader's perceptions of change.org petitions in general.

Version 1:

Decision Letter:

Our ref: NATHUMBEHAV-25094355A

16th March 2026

Dear Dr. Corpus,

Thank you for submitting your revised manuscript "Introducing AI to an Online Petition Platform Changed Outputs but not Outcomes" (NATHUMBEHAV-25094355A). It has now been seen by the original referees and their comments are below.

As you can see, the reviewers find that the paper has improved in revision. We will therefore be happy in principle to publish it in Nature Human Behaviour, pending minor revisions to satisfy the referees' final requests and to comply with our editorial and formatting guidelines.

We are now performing detailed checks on your paper and will send you a checklist detailing our editorial and formatting requirements within two weeks. Please do not upload the final materials and make any revisions until you receive this additional information from us.

Please do not hesitate to contact me if you have any questions.

Sincerely,

[REDACTED]

[REDACTED]

Nature Human Behaviour

Reviewer #1 (Remarks to the Author):

Good work - thanks for the added institutional detail and labeling exercise.

Reviewer #2 (Remarks to the Author):

The paper has improved. It's quite helpful to learn more about the authors' connections to former change.org employees, and for them to analyze petition contents and headline concreteness. I thank the authors for their responsiveness, and I agree with them that they should not have to require platform connections to publish empirical platform research. I would advise them to disclose the private information that influences their analysis within the paper, as withholding it leaves the reader to guess what you know that you're not saying, but this is their choice.

Unfortunately, I remain uncertain what the paper teaches us. The single geographical control calls the quasi-experimental identifying assumptions into greater question than usual. Subject to the identifying assumption, we learn that Gen AI made petitions better while reducing their effectiveness. This is quite a puzzle! Usually, higher quality leads to higher success, but in this case higher writing quality reduced success.

Why did this happen? Authors speculate about three mechanisms:

1. Quality superficiality: AI may have made low-quality petitions look like high-quality petitions, removing writing quality as a signal that readers could use to discern petition quality.
2. AI suspicion: Readers may have been suspicious about AI use, and avoided supporting petitions as a result.
3. Reduced petition author effort: Authors may have over-relied on gen AI tools and failed to invest to develop petitions' qualitative content sufficiently.

I asked about a fourth possibility which the authors dismissed. Maybe by enabling low-quality petition completeness, Gen AI injected noise into the system that generates petition success. I asked about the competition for attention that occurs on the

platform. Authors replied that petition marketing occurs off of change.org. Yet they do not address the possibility that the success of off-platform petition marketing may depend on petition characteristics. I think it is entirely realistic to think that off-change.org algorithms (e.g., Facebook, Instagram, LinkedIn, etc.) may evaluate the content of external links and use those content evaluations to determine how widely to distribute the link.

Ultimately I agree with the authors that no causal mechanism evidence is available here. Their discussion and speculations and descriptive evidence are generally very reasonable, but doesn't go very far to help us understand why this very surprising main effect occurs. So I am left with some doubt as to whether the unverifiable necessary conditions for quasi-experimental identification are met, I am left with doubt as to what other settings this might apply to, and I am left with doubt as to what I can learn from this.

Clearly I am out of step with the other reviewers. I will understand if the journal decides to publish the article without my endorsement. For authors' sake, I encourage the journal to make an up-or-down decision at this stage.

I wish the authors well and thank them for their research.

Reviewer #3 (Remarks to the Author):

Overall, I appreciate the authors' engagement with the feedback. I believe the issues I raised have been appropriately addressed, and the paper is now more precise in its discussion of some of the findings.

Version 2:

Decision Letter:

Dear Ms Corpus,

We are pleased to inform you that your Article "Introducing AI to an Online Petition Platform Changed Outputs but not Outcomes", has now been accepted for publication in Nature Human Behaviour.

Authors may need to take specific actions to achieve compliance with funder and institutional open access mandates. If your research is supported by a funder that requires immediate open access (e.g. according to [Plan S principles](https://www.springernature.com/gp/open-science/plan-s-compliance) or the [NIH public access policy](https://www.springernature.com/gp/open-science/us-federal-agency-compliance)) then you should select the gold OA route, and we will direct you to the compliant route where possible. Because authors warrant under our subscription licensing terms that they haven't committed to licensing any version of their article under a licence inconsistent with the terms of our agreement – including the applicable embargo period – publication under the subscription model isn't suitable for authors whose funders require no embargo.

Once your manuscript is typeset and you have completed the appropriate grant of rights, you will receive a link to your electronic proof via email with a request to make any corrections within 48 hours. If, when you receive your proof, you cannot meet this deadline, please inform us at rjsproduction@springernature.com immediately. Once your paper has been scheduled for online publication, the Nature press office will be in touch to confirm the details.

Acceptance of your manuscript is conditional on all authors' agreement with our publication policies (see <http://www.nature.com/nathumbehav/info/gta>). In particular your manuscript must not be published elsewhere and there must be no announcement of the work to any media outlet until the publication date (the day on which it is uploaded onto our web site).

If you have posted a preprint on any preprint server, please ensure that the preprint details are updated with a publication reference, including the DOI and a URL to the published version of the article on the journal website.

An online order form for reprints of your paper is available at <https://www.nature.com/reprints/author-reprints.html>. All co-authors, authors' institutions and authors' funding agencies can order reprints using the form appropriate to their geographical region.

We welcome the submission of potential cover material (including a short caption of around 40 words) related to your manuscript; suggestions should be sent to Nature Human Behaviour as electronic files (the image should be 300 dpi at 210 x 297 mm in either TIFF or JPEG format). Please note that such pictures should be selected more for their aesthetic appeal than for their scientific content, and that colour images work better than black and white or grayscale images. Please do not try to design a cover with the Nature Human Behaviour logo etc., and please do not submit composites of images related to your work. I am sure you will understand that we cannot make any promise as to whether any of your suggestions might be selected for the cover of the journal.

You can now use a single sign-on for all your accounts, view the status of all your manuscript submissions and reviews, access usage statistics for your published articles and download a record of your refereeing activity for the Nature journals.

To assist our authors in disseminating their research to the broader community, our SharedIt initiative provides you with a unique shareable link that will allow anyone (with or without a subscription) to read the published article. Recipients of the link with a

subscription will also be able to download and print the PDF.

As soon as your article is published, you will receive an automated email with your shareable link.

In approximately 10 business days you will receive an email with a link to choose the appropriate publishing options for your paper and our Author Services team will be in touch regarding any additional information that may be required.

If you have any questions about our publishing options, costs, Open Access requirements, or our legal forms, please contact ASJournals@springernature.com

We look forward to publishing your paper.

With best regards,





Nature Human Behaviour

P.S. Click on the following link if you would like to recommend Nature Human Behaviour to your librarian
<http://www.nature.com/subscriptions/recommend.html#forms>

** Visit the Springer Nature Editorial and Publishing website at http://editorial-jobs.springernature.com?utm_source=ejp_NHumB_email&utm_medium=ejp_NHumB_email&utm_campaign=ejp_NHumB > www.springernature.com/editorial-and-publishing-jobs for more information about our career opportunities. If you have any questions please click here. **

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source. The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Editor:

Reviewer #2 raises concerns about the interpretation of the results, given potential unobserved effects of the introduction of the LLM on user behaviour or platform dynamics. It will be important to address these concerns, ideally with additional data.

Response: Thank you for your detailed and helpful feedback. We provide a detailed response to reviewer concerns in this document. We include additional data to support our analysis: in particular, we ran a labeling task in which human participants labeled petitions (N=417) for writing quality and persuasiveness. Additionally, we include additional analysis in the SI Appendix of petition text sentiment, title concreteness, and the distribution of petition topics. In the following comments we provide clarity on the relationship between the in-platform AI and user behavior/platform dynamics. Furthermore, we have edited the writing of our Introduction, Discussion, and SI Appendix to consider points raised by the reviewers and lend clarity where necessary. We appreciate the editor and referees for their careful reading and thoughtful comments on the paper.

REVIEWER COMMENTS:

Reviewer #1 (Remarks to the Author):

This paper investigates the effects of genAI integration on an international platform that aggregates proposals and facilitates democratic governance and inquiry. The authors use the plausibly exogenous introduction of ChatGPT, leveraging the staggered adoption of genAI in Australia and specifically the platform's integration of it in Australia.

The paper is very well done: it's clean, clear, and contains all the relevant parallel trends and related diagnostics. I don't think there's any reason to make this a long drawn-out process; the paper is effectively ready to go, as far as I can tell. The authors do a good job setting up the prospective mechanisms. I also appreciated with the DID estimator that you focused on different sub-geographies in Australia.

Response: We thank the reviewer for their support.

The only two minor comments:

1. Can you talk more about the institutional reasons for why Australia had a later adoption?

Response: Our knowledge of the delayed AI release in Australia stems from private communications with two Change.org employees who were directly involved with the development of the AI tool. We learned from these conversations that Change.org typically conducts segmented launches, in which some countries receive A/B testing and product access

earlier than other countries. We further learned that, in this case, Australia had delayed AI tool release because of uncertainty around legal counsel. As a result, while the U.S., Great Britain, and Canada had A/B testing over the summer of 2023 and full access starting in October 2023, the Change.org legal team only met with Australian counsel during the second week of December 2023 and so Australia only first received access in late December, 2023.

We have added a paragraph to the SI Appendix introduction to the natural experiment, under the “[Change.org](#) and the Write with AI Tool” section, in which we speculate on the reasons for the later adoption (i.e., tech platforms typically conduct staggered releases, legal uncertainty in 2023 due to evolving AI governance across countries), but do not directly quote the private communications.

2. To further parse the mechanisms, you could always stand up a survey on Prolific.co, but you provide good enough results already. You could, however, look at the text of the proposals and see if the words that are used are more dis-passionate or less specific (reflecting the homogenizing effects that genAI can have on proposals).

Response:

We appreciate the suggestion and took three actions to address your points.

First, to further parse the mechanisms, we ran a labeling task in which participants labeled a sample of 417 petitions (203 written pre-AI access, 214 written post-AI access) for writing quality and persuasiveness. We ran this task to have stronger descriptive evidence of the features that characterize petition text when written with or without AI access. Through this analysis, we are able to strengthen our hypotheses for the mechanisms that drive the trends observed in the main text. In particular, we observe that petitions written with AI access have higher writing quality and persuasiveness ratings, on average, than petitions written without AI access. This point suggests that our main findings of unchanged (or worsened) petition outcomes are not driven by degraded writing quality or persuasiveness of the text. We describe our findings in the Results section of the main text, under “Petition Text Quality”, refer to this finding in our discussion of mechanisms in the Discussion section, and detail the methods and results in the SI Appendix, in the section “Petition Writing Quality and Persuasiveness.”

Second, to consider if language is more dis-passionate, we used the VADER (Valence Aware Dictionary and sEntiment Reasoner) Python tool to label the sentiment of petition text. The VADER tool takes in text, and outputs the ratio of that text which falls into one of three categories: positive, neutral, and negative. If the petition text became less passionate, we would expect to see an increase in the share of a petition that is attributed the “neutral” tag. We observe that the ratio of sentence text assigned the “neutral” tag decreases when petitions are written with AI access. In contrast, we observe an increase in the distribution of text assigned the “positive” tag. This pattern is aligned with previous work suggesting that generative AI models have a positivity bias. We write about the methods and the results of this analysis in the

SI Appendix in the section titled “Sentiment Analysis” and show the distribution of petition sentiment in Figure 11.

Third, to consider if petitions are “less specific” we label the concreteness of petition titles using an automated scale of headline concreteness proposed by Aubin Le Quere and Matias.¹ We find descriptive evidence that the distribution of petition title concreteness shifts toward lower concreteness scores when petitions are written with AI access, with modestly lower median title concreteness. We show the distribution of title concreteness in Figure 12 of the SI Appendix. We write about the methods and analysis of the title concreteness analysis in the SI Appendix in the section titled “Title Concreteness”, and refer to the hypothesis of generic language as influential to petition outcomes under the Discussion section of our main text.

Great work!

Response:

Thank you!

Reviewer #2 (Remarks to the Author):

Review of “Introducing AI to an Online Petition Platform Changed Outputs but not Outcomes”

I read this paper with interest. It empirically analyzes a change in tools offered to petition writers on a prominent petition hosting platform, Change.org (“the platform”). The platform invites petitions from authors, then invites signatures and comments on petitions, and then invites owners to declare ‘success’ when petitions succeed. I learned from this paper that large majorities of petitions do not attract any comments or indicate success.

The platform integrated its petition writing tools with an early large language model in 2023. It first introduced the change in Australia before its other English-speaking markets. After ten weeks it extended the change to other markets. This is a promising setting for a quasi-experimental econometric application. In general I think authors’ applications are detailed, competent, clear, and mostly avoid overinterpretation. They examine robustness to alternate approaches, all great attributes to have.

My primary questions are about aims, hypotheses and contribution. These are first-order questions whose answers remain unclear to me after reading the paper and supplementary information. Their answers may influence the measurements and econometric designs, or they may not. Without those answers I have limited information about the optimal research design. Therefore I will not comment extensively on the empirics.

¹ Aubin Le Quéré, M., Matias, J.N. When curiosity gaps backfire: effects of headline concreteness on information selection decisions. Sci Rep 15, 994 (2025). <https://doi.org/10.1038/s41598-024-81575-9>

Response:

We appreciate the reviewer's interest in the paper and characterization of the Change.org platform, AI tool, and setting for natural experiment and the reviewer's suggestion that we clarify our aims, hypotheses, and contribution. We hope the following responses can provide clarity on first-order questions in the responses that follow.

1. Aims

The platform apparently tested the LLM integration in Australia before extending the integration to other countries. The criteria for such a test are left unstated. They may be simple, such as did the integration succeed technically? However, I strongly suspect they extended to user-related outcomes as well. We are not credibly informed about the platform's goals of the integration, nor how the test performed on those goals.

Response:

We'd like to clarify that the tool was launched first in the United States, Canada, and Great Britain, with a delayed release in Australia. The reason for the delayed release in Australia was largely logistical — the Change.org team required additional legal counsel before releasing the tool in Australia, and did not schedule that legal counsel meeting until mid-December. As such, the tool was released in Australia in late December. We detail the reason for the delayed release above, in response to Reviewer #1, and added speculation for the delayed release in the SI Appendix ("Change.org and the Write with AI Tool").

My naive guess is that the platform sought to increase the completion rate of draft petitions. Authors often might start a petition and then abandon it due to the writing cost. Offering writing assistance may help to increase the completion rate and thereby increase the number of petitions available on the platform. Or LLM integration might help to reverse a declining trend on the platform. Or perhaps the platform noticed a decline in petition writers' writing quality during the pandemic and thought the LLM integration could help to reverse that.

The authors have chosen some metrics that are partially aligned with these claims. It remains unclear what the platform intended or how it evaluated success.

Response:

We appreciate the reviewer's suggestion that we elaborate on the technical success of the integration and user-related outcomes. We would like to use this point to clarify that we did not seek to replicate the platform's internal metrics. The metrics we selected reflected key features of petition text and the outcomes that enable petitions to be successful tools of change. We did not seek to replicate the platform's metrics, as we did not intend to suggest that the platform had failed to meet their internal evaluation of the tool. Instead, we suggest that regardless of the platform's intent, platform-level changes occurred that may be interesting to consider, as they may present unintended harms (e.g. homogenization, reduced outcomes, the deterioration of

signals of underlying quality). While it would be interesting to study additional metrics related to user-specific outcomes, such as completion rates, unfortunately, we conducted this study without platform collaboration. Furthermore, we note that these internal platform metrics would be private, and therefore not reproducible, even if our team did receive internal data access. We add discussion of the constraints of studying online platforms without collaboration or internal access in a paragraph on limitations in the SI section “Change.org and the Write with AI feature”. We are happy to add commentary to the main text, if recommended.

We also note in our SI Appendix (Figures 37, 38) that we did not observe a shift in the trend of petition production after the tool was released. We used Bayesian Structural Time Series modeling to predict petition production using pre-AI trends, and found that the actual volume of petitions produced did not differ significantly from the predicted volume. However, as the reviewer notes, it is still possible that the counterfactual would have been a declining volume of petitions.

To clarify some of the points raised by the reviewer’s comment above, we further:

1. Add speculation regarding the potential positive impact of LLMs on completion rates in the Introduction. “As suggested by prior research, it is possible that this lower barrier for participation could increase petition production and completion rates on the platform, while improving the quality of content produced by low-skill writers.”

2. Suggest a positive interpretation of our results in the Discussion section: Our results, and prior research, suggest that it may be the case that with AI, petition-writers could expend less effort to produce a petition of higher quality to a petition written without AI [1]. Prior literature, as well as our results on petition writing quality and persuasiveness, suggests that the average quality of petition text and arguments would improve with AI assistance [1, 21, 22, 23].

3. Add the following lines to the “Change.org Write with AI Feature” section of the SI Appendix: “Furthermore, this message (“Create a compelling petition in minutes!”) suggests that the platform’s intent in releasing the AI tool was to enable petition-writers to create high quality petitions with reduced effort. By lowering the barrier to petition-creation, we might speculate that the platform sought to increase the volume of petitions produced on the platform or increase petition completion rates.”

2. Hypotheses.

The authors’ causal logic appears to simplify to the following:

- 1. Writing quality increases petition success markers*
- 2. LLM assistance increases writing quality*
- 3. Therefore, LLM assistance should increase petition success markers*

There are several weaknesses in this causal chain.

Response:

We appreciate the opportunity to clarify the causal logic proposed in the paper.

We indeed suggest that this chain *may* occur. However, we do not claim to prove this chain. Instead, we propose that this chain may be worth considering, given the descriptive and causal evidence that we are able to show without platform collaboration.

The causal claims we make are as follows:

- In-platform AI access significantly changed platform-level petition lexical features (i.e., word count, lexical diversity, and reading difficulty).
- In-platform AI access significantly increased the homogeneity of petition text.
- In-platform AI access did not improve (and by some measures, worsened) platform-level outcomes.

Having said that, we do believe that the causal chain proposed by the reviewer is an interesting speculation. We have added new results that lend more strength to the chain proposed by the reviewer, based on the reviewer's suggestions below.

2a. The first claim remains unproven. Authors note in the text that writing quality predicts petition success markers, and they show this in SI Figure 17. However, this is not a test of a causal claim. It may easily be that people who write well also choose petition topics that are likely to succeed. Without the first link in the chain, we have no reason to believe that LLM assistance should increase petition success markers.

Response:

We added new (non-causal) analysis of petition quality and analyses that control for petition topics that we hope alleviate some of these concerns.

We agree that we should consider writing quality more directly, though a causal claim is hard to prove without platform access. To address this question, we ran a labeling task in which human labelers annotated a sample of 417 petitions for writing quality and persuasiveness. We show that petitions written with access to the AI tool have, on average, higher writing quality and persuasiveness. We also find that writing quality (and persuasiveness) has a strong predictive relationship with success outcomes (probability of receiving 1+ comment in 30 days) when petitions were written before AI access. When petitions are written with AI access, this predictive relationship is significantly weaker, or reversed. We agree and explicitly state that this is not a test of a causal claim. We include the results of this analysis in the new Results section "Petition Text Quality", and in the SI Appendix (Figures 6-9, Tables 4-8). Furthermore, we

demonstrate that the lexical features we study in our main text are positively correlated with writing quality (Figure 10).

Furthermore, we appreciate the reviewer's recommendation that we consider the petition topic as a confounder. To better understand this point, we labeled petition by topic (methods detailed in SI Appendix section Petition Topic). In SI Appendix Figure 30 we demonstrate that the distribution of petition topics remained consistent during the Pre-AI, A/B test, and Post-AI periods. In SI Appendix Figure 31, we also demonstrate for the repeat-writer analysis that the share of petition authors that write about the same topic in their first and second petitions remains consistent across the three cohorts (Pre/Pre, Pre/Post, and Post/Post), at about 35%. Finally, in Figures 32 and 33 we provide a robustness check for our main analysis, in which we only consider a subset of topics. In particular, in Figure 32, we conduct the same difference-in-differences analysis for civic engagement topics; in Figure 33, we repeat the analysis for "other" topics (not related to civic engagement). We observe consistent results across these analyses.

We point to these analyses in the main text under Materials and Methods: "We do not include any covariates in our main analysis; the SI Appendix provides robustness checks with additional covariates, such as day of week and petition topic."

2b. Return to point 1. Suppose the primary purpose of LLM integration is to increase draft petition completion rates. If uncompleted petitions tended to be low quality, then LLM integration would reduce the average quality of petition on the platform. We then should expect petition success markers might even decrease on average. In fact this appears consistent with figure 1d. But it seems to reverse the hypotheses that the authors seem to maintain.

Response:

We hope that our new data from labeling petition writing quality and persuasiveness, as detailed above, helps alleviate this concern, by demonstrating that writing quality is *higher*, on average, when petitions are written with access to the in-platform AI tool.

We would like to re-assert that we do not make any causal claims about the effect of LLM integration on petition writing quality. With our labeling task, we have non-causal evidence that the average writing quality of petitions written with AI access is significantly higher than for petitions written without AI access. We find a similar relationship for persuasiveness. We should then expect petition success markers to increase on average, if the predictive relationship between writing quality and outcomes remains consistent with in-platform AI access. However, as we show in Figures 8 and 9 and Tables 5-8, there is a significant interaction effect between the treatment period and writing quality in predicting outcomes. While writing quality had a positive predictive relationship of petition outcome pre-AI, this relationship is negative post-AI.

Taken together, we believe we have provided strong evidence that the diminished outcomes are not due to lower quality petitions created by LLMs.

2c. Completed petitions compete in a ranking system for platform user attention. There is a default ranking system and users may sort or filter petitions in some limited ways. We do not know how petition writing quality affected the default ranking system, therefore we do not know how writing quality may change user attention to completed petitions. This is a very common issue in empirical work in platforms, that interventions often affect platform algorithm inputs before and then subsequently affect users, so it becomes difficult or impossible to separate the effect of the intervention on the ranking system from the effect of the intervention on the users. For example, this issue is the difference between “lift tests” and “AB tests” in Facebook advertising; the former is a true experiment, whereas the latter combines the effect of an ad on the platform algorithms with its effect on the exposed users. No general solution is available to this issue.

Response:

We appreciate the reviewer’s suggestion that we consider algorithmic ranking as a confounder for this study. While we acknowledge that it is impossible to discount the role of algorithmic systems when studying online behavior, as is always the case in computational social science studies, we do not believe that these factors have significant impact on our findings. Unlike the case of online advertising, usage patterns on Change.org are fundamentally different from the usage patterns of Facebook or other such social networks. Specifically, as detailed below, (1) user behavior on Change.org is not based on browsing, but rather from external promotion, and (2) the home page on Change.org only surfaces already-successful petitions.

First, prior work has detailed how engagement with petitions is largely sourced from external platforms off of Change.org. Few petition signatures come from users browsing on Change.org. Indeed, most petition signees are routed to petitions from individuals or organizations that they trust. As such, the algorithms that drive ranking of Change.org website content (e.g. home page) are not consequential to the petition outcomes we measure. To address this point we added this text to the Methods and Data section of the main text: “Prior research into online petitions has demonstrated that e-petitions are dependent on off-platform promotion through social networks, news outlets, or direct shares to gain signatures [42, 43, 44, 45, 46]. Crowdfunding platforms, which operate similarly to Change.org, have been found to display “friendfunding”, by which campaigns tend to be funded by friends and family, unless shared broadly on external social media platforms [47, 48]. Based on this extensive prior work, we hypothesize that few readers browse on Change.org in pursuit of petitions, but instead find petitions off-platform through social networks, news coverage, or direct shares from friends.”

Second, while we acknowledge that the Change.org home page helps featured petitions gain signatures, we expect that our analysis would not be impacted by this feature. The petitions featured on the home page during the time of the natural experiment were featured because

they were either declared as “victories”, or trending. In both cases, petitions featured would have already garnered far more than 10 signatures and 1 comment (oftentimes, trending petitions have thousands of signatures) required for “successful outcome” in our analysis. We include a screenshot of the home page at the time (via the Internet Archive Wayback Machine) in the SI Appendix (Figure 1), and at the end of this response document. As such, we consider our use of minimum thresholds to provide robustness against these confounds.

3. Contribution

Given the uncertainty about the platform’s original intention, the appropriate measures to evaluate effects, and the mechanism, I find it hard to know what I can draw from this article that could inform other settings.

I do not think it would be reasonable to believe that LLM assistance would generally worsen existing creators’ outcomes in other platforms.

I find it generally reasonable to think that LLM assistance might increase creation by increasing creation accessibility, while simultaneously worsening average creation quality. How those changes affect the resulting equilibria would further depend on how high-quality creators respond to low-quality creations, how high-quality creators integrate LLM within their own creations, how users respond to the changes, and how the platform itself responds (via, e.g., its search rankings or presentation algorithm).

Response:

We appreciate the reviewer raising these points to help us clarify the contribution.

We believe this study’s findings will be meaningful for Change.org and many other advocacy and social platforms, because of the widespread trend of integrating AI tools into online platforms, and lack of existing understanding around how such integrations will impact outcomes.

While we cannot, as external researchers, fully control for all the factors involved, we believe (as detailed above) we have provided significant evidence that, overall, the introduction of AI to Change.org resulted in diminished platform-level outcomes that are not due to changed participation (SI Appendix Figures 37, 38) or reduced quality (SI Appendix Figure 7). Our paper does not seek to estimate the impact of AI access on individual users, but instead on the dynamics of the platform at large. As online platforms involve deeply interconnected networks of users, causal estimates of interventions on individual-level outcomes may suffer from bias, as the reviewer notes, and require industry connections.² Instead, we estimate the causal

² Bak-Coleman, Joseph B., Stephan Lewandowsky, Philipp Lorenz-Spreen, Arvind Narayanan, Amy Orben, and Lisa Oswald. "Moving towards informative and actionable social media research." arXiv preprint arXiv:2505.09254 (2025).

impact of an intervention to the system. In our Discussion section, our speculation regarding social mechanisms outlines a few possible future work directions that may be fruitful for understanding the mechanisms at play but are beyond the scope of current work.

Following comments are more minor relative to the foregoing

4. Intervention. I am unclear whether the LLM was tuned to predict petition writing features that led to successful petitions in the pre-LLM era. I am also unclear how the rapid recent innovations in LLM quality may affect the durability of the authors' results. I think they should comment on this, even if it is only speculative.

Response:

We appreciate the reviewer's suggestions. Indeed, changes to prompting and LLM innovation could improve the text of the petitions, to make them increasingly readable, persuasive, or any other attributes that drive petition success. However, as we describe in our Discussion section, we do not believe that the quality of the text alone explains shifts in outcomes. Indeed, the results of our newly added labeling task lead us to speculate that the current prompting strategy used by Change.org does generate higher quality petitions, on average, than were written by humans alone prior to access to the tool. However, outcomes do not improve with this improvement to text quality. As such, we do not see evidence to suggest that the trends demonstrate a *technological* shortcoming that would change with more powerful LLMs – instead, we suspect that social mechanisms (e.g. reader AI suspicion, writer commitment to petitions) are at play, and might be important to consider in future work. Furthermore, these social mechanisms are likely durable over time and not tied to any single LLM innovation or degree of technological capacity: for example, AI suspicion has been attributed to social penalties from 2023³ to 2025.⁴

We address the technological advancement of LLM's in our discussion section: "In particular, the technological capabilities of AI to produce believable, high-quality text might be insufficient to explain the outcomes of content produced on social platforms, as evaluations are determined by readers' dynamic, subjective, and social perceptions. Specifically, it is possible that petitions, with their real-world, context-specific missions, fall outside of AI's capabilities for augmentation because the task requires personal and community-specific knowledge and engagement."

5. If the authors have connections at the platform, I think their research would be much stronger if they could report an experiment that increases the barrier to petition creation. If it is "too easy" to create a petition, the platform may have too much noise, and the best petitions may be

³ M. Jakesch, J.T. Hancock, & M. Naaman, Human heuristics for AI-generated language are flawed, Proc. Natl. Acad. Sci. U.S.A. 120 (11) e2208839120, <https://doi.org/10.1073/pnas.2208839120> (2023).

⁴ J.A. Reif, R.P. Larrick, & J.B. Soll, Evidence of a social evaluation penalty for using AI, Proc. Natl. Acad. Sci. U.S.A. 122 (19) e2426766122, <https://doi.org/10.1073/pnas.2426766122> (2025).

seen suboptimally infrequently. A possible intervention could be a delay, or even better, a random intervention in petition visibility.

Response:

We thank the reviewer for this suggestion! Unfortunately, we do not have a collaboration with the platform; all data used for analysis was scraped from the public-facing petition pages on Change.org. Information about internal processes is drawn from private communications with informants at Change.org, however, we do not have internal access to data or ability to run experiments. To address this concern, we add a paragraph to the SI Appendix section “Change.org and the Write with AI Tool” in which we detail the challenges of studying online platforms without collaboration.

I thank the authors for their research and I hope to read a future version of the article which clarifies these points. I encourage them to develop their research further and hope these comments may be helpful.

Response:

Thank you again; your comments helped us more clearly articulate our assumptions and our causal arguments.

Reviewer #3 (Remarks to the Author):

This paper describes a difference-in-difference natural experiment on change.org to study the impact of in-platform AI access on petitions’ lexical features and petitions’ comment outcomes. Overall, I thought the paper was well-written, the study to be rigorously done, and the findings interesting. I recommend accept.

Response:

We appreciate the reviewer’s comments and suggestions.

That said, I do have some suggestions for authors to consider:

1. I find the framing of AI access to be a bit overly negative. While I'm not surprised that AI did not improve petition signing, it doesn't necessarily mean there is no value for AI access. What is not factored in this current study is the impact of AI access on petition-writing work. If it helped reduce some effort there, a positive spin of the results could be that the same outcome could be achieved with lower cost. Given that effort/cost is not factored in this work, it might be good to at least acknowledge that potential tradeoff.

Response:

We appreciate the reviewer's feedback, and agree that we could better note the potential benefits of the AI tool. We have added to the Introduction and Discussion of the main text to emphasize the potential positive effects of an in-platform AI tool.

We initially stated in the Introduction: "Conceivably, integration of AI could bolster the effectiveness of the Change.org platform, as well as individual petitions, by lowering the barrier to participation and enabling users to produce high quality content with ease." We have added the following sentence: "As suggested by prior research, it is possible that this lower barrier for participation could increase petition production and completion rates on the platform, while improving the quality of content produced by low-skill writers."

We have amended the framing of the discussion to acknowledge the point that reducing average user effort while maintaining the same average level of outcomes (and increasing average writing quality, persuasiveness) might be a positive outcome. We have added the additional lines:

"On one hand, this finding suggests that with AI, petition-writers could expend less effort to produce a petition with similar outcomes to a petition written without AI (and more effort). Indeed, we observe that petitions written with AI access have, on average, higher writing quality and persuasiveness scores than petitions written without AI access. This descriptive finding suggests AI could increase the writing quality and persuasiveness of petitions, especially those written by low-effort or low-skill authors."

2. The part where I find the most interesting is the findings on outcomes. However, one thing I think is valuable to note in the paper is that it is unclear how much of the observed difference is due to the impact of AI access on the quality/persuasiveness of individual petitions, versus the impact of AI access on a more macro-level, on the general perceptions of petitions on change.org. While I believe in the findings on the causal impact of AI access, the underlying mechanism is less clear. The paper seems to make the argument that it's because of the impact on the petition. I'm not confident about that. I think a more macro-level, longer-term effect could help explain why the effects are stronger for %petitions with 10+ signatures than %petitions with 1+ comment when limiting that to 30 days, and the difference between first and second petitions for repeat petition writers when both are post-AI.

Response:

We agree, and our new results on quality and persuasiveness further support the fact that other factors are involved. We agree that this is exactly what makes our contribution interesting: despite the rise in quality, there are diminished outcomes, calling for other explanations which we bring out in the discussion.

With regards to the two outcomes (signatures and comments), we clarify that the distinction between time-boxing comments and not time-boxing signatures was a decision made during pre-registration before looking at the data. We suspected that the time-boxed comment

outcome would prevent discounting the outcomes of more recent petitions. We would have time-boxed signatures as well if we had had access to this data (the data we scraped only contained the total number of signatures as of the time of crawl). However, we do not anticipate that there would be much difference in time-boxed and non-time-boxed signature data. Prior longitudinal studies of petition outcomes have shown petitions likely gain a significant share of their signatures in the first month of activity.⁵ As we describe in the Materials and Methods section, “Provided that less than 1% of petitions gained their first comment after the first 30 days, we expect that most petitions also reach 10 signatures within the first 30 days, if at all.” Furthermore, when we analyzed the effect of in-platform AI access on total comments (not time-boxed), we observed similar results to the time-boxed comment outcome metrics. We show these results in the SI Appendix (SI Figure 14, writeup in SI Appendix Difference-in-Differences Robustness Checks).

I don't quite agree with the authors' argument to discount “AI suspicion”. Perhaps there's no explicit indication that any single petition was AI generated, but AI access is clearly communicated on change.org (at least in the screenshot provided).

Response:

Thank you for the suggestion; we have qualified our discounting of the AI suspicion hypothesis. However, we note that during our period of differential access, the Change.org home page did not clearly communicate AI was being used. Figure 2 in the SI Appendix, exposing the AI use, is from 2025. We added Figure 1 in the SI Appendix that includes a (non-AI) screenshot of the home page as of November 1, 2023, during the period of differential AI access. Furthermore, we do not think (see comments above in response to Reviewer #2, bullet point 2c) that most readers discover petitions via the main page, but instead through circulation off-platform.

To further qualify, we added the following sentence to the Discussion section to strengthen our consideration of AI suspicion: “Despite the absence of platform-provided signals of AI use, readers may have heightened AI suspicion driven by the lexical style, paragraph structure, and vocabulary of petitions written with AI access. Such shifts could potentially increase suspicion among petition readers.”

I think the paper can be more precise about this and acknowledge while some of the findings are clearly about the petitions themselves (e.g., lexical features and petition homogeneity), the effect of AI access on outcomes may have been impacted by different mechanisms – e.g., how AI access is changing the petition text, as well as how its introduction is shaping reader's perceptions of change.org petitions in general.

⁵ Julia Proskurnia, Przemyslaw Grabowicz, Ryota Kobayashi, Carlos Castillo, Philippe Cudré-Mauroux, and Karl Aberer. 2017. Predicting the Success of Online Petitions Leveraging Multidimensional Time-Series. In Proceedings of the 26th International Conference on World Wide Web (WWW '17). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 755–764. <https://doi.org/10.1145/3038912.3052705>

Response:

We agree; see comments above to this and other reviewers' comments. The Discussion goes deeply into various other mechanisms. Specifically, regarding the point that in-platform AI access may change perceptions of Change.org in general, we add the following line in Discussion: "Our second proposed mechanism is AI suspicion: readers' perceptions of petitions, and the Change.org platform, might deteriorate if they suspect the use of AI."

Figures

change.org

Start a petitionMy petitionsBrowseMembership

Log in

The world's platform for change

516,966,365 people taking action. **Victories every day.**

Start a petition

VICTORY



UPS drivers get A/C

Five years and 1.3 million signatures later, Theresa's petition to get A/C in UPS delivery trucks won. The Teamsters officially ratified a contract with historic heat protections. Theresa's petition is now one of the biggest victories in Change.org's history.

Theresa Kle...

United States

1,352,713

Supporters

Featured in

Forbes



UPS drivers get A/C



Iran removed from UN Commission on the Status of Women



Brittney Griner is released from Russia



Kanye West's honorary degree rescinded



UN Passes Survivors' Bill of Rights



Fraternity held accountable

What's happening on Change.org

Trending in the movement: **Prevent US Pedestrian Deaths**

See more

Dangerous cars are killing our children

In 2021, my 5-year-old daughter was hit and killed by a vehicle. Allison was one of 42,939 Americans violently taken from us by traffic violence that year. Allie was doing something any kid might do: riding her bike after school. She was riding her bicycle in a crosswalk – in a school zone – just a block away from our home. Her 40-pound body was no match for the transit van that struck and killed her.... [Read more](#)



Jessica Hart

36,310 supporters

Trending in **Economic justice**

See more

Raise and extend Unemployment benefits for those affected by the WGA and SAG strikes.

Crew members, vendors, and other affiliated workers in the entertainment industry have been forced out of work for over seven months now with no end in sight. We are facing financial ruin with no opportunity to work. The EDD benefits max amount of \$450 a week and are ending for most people. That amount is not ... [Read more](#)



Farah Bunch

25,063 supporters

Trending in **Human rights**

See more

Unite Against Discrimination: Stop SB 264 and End AAPI Hate in Florida!

The time has come to confront the deep-rooted discrimination faced by the Asian American and Pacific Islander (AAPI) community in Florida. We cannot stand idly by while Florida's unjust property law, SB 264, perpetuates hateful treatment. It's time for action! On May 8, 2023, Governor Ron DeSantis signed SB 264 into ... [Read more](#)



Jenny Liu

44,433 supporters

FEATURED TOPICS

- Iranian revolution
- Abortion access
- Twitter
- Human rights iran
- Women's rights
- Free speech
- Congressional integrity
- Entertainment

We thank the reviewers for their thoughtful attention to our revised paper, and for their positive feedback about our added analysis and data. We have incorporated the suggested edits to the main text by adding a limitations section (described below), and note that some numbers have shifted slightly as we fixed minor issues in the code (none that change interpretation or results).

Reviewer #1 (Remarks to the Author):

Good work - thanks for the added institutional detail and labeling exercise.

Thank you for your support! We are confident that your feedback has strengthened the paper.

Reviewer #2 (Remarks to the Author):

The paper has improved. It's quite helpful to learn more about the authors' connections to former change.org employees, and for them to analyze petition contents and headline concreteness. I thank the authors for their responsiveness, and I agree with them that they should not have to require platform connections to publish empirical platform research. I would advise them to disclose the private information that influences their analysis within the paper, as withholding it leaves the reader to guess what you know that you're not saying, but this is their choice.

We appreciate the suggestion that we include information about our personal communications with Change.org employees, which we have included in the SI Appendix section "Change.org and the Write with AI Tool".

Unfortunately, I remain uncertain what the paper teaches us. The single geographical control calls the quasi-experimental identifying assumptions into greater question than usual. Subject to the identifying assumption, we learn that Gen AI made petitions better while reducing their effectiveness. This is quite a puzzle! Usually, higher quality leads to higher success, but in this case higher writing quality reduced success.

Why did this happen? Authors speculate about three mechanisms:

- 1. Quality superficiality: AI may have made low-quality petitions look like high-quality petitions, removing writing quality as a signal that readers could use to discern petition quality.*
- 2. AI suspicion: Readers may have been suspicious about AI use, and avoided supporting petitions as a result.*
- 3. Reduced petition author effort: Authors may have over-relied on gen AI tools and failed to invest to develop petitions' qualitative content sufficiently.*

I asked about a fourth possibility which the authors dismissed. Maybe by enabling low-quality petition completeness, Gen AI injected noise into the system that generates petition success. I asked about the competition for attention that occurs on the platform. Authors replied that petition marketing occurs off of change.org. Yet they do not address the possibility that the success of off-platform petition marketing may depend on petition characteristics. I think it is entirely realistic to think that off-change.org algorithms (e.g., Facebook, Instagram, LinkedIn,

etc.) may evaluate the content of external links and use those content evaluations to determine how widely to distribute the link.

We appreciate your suggestion that we consider the influence of algorithmic systems on the platforms where petitions might be externally posted. We have added a paragraph to our discussion section that explicitly includes the limitations you described, including off-platform algorithmic amplification, the limitations of using a single geographic control, and our lack of causal mechanism evidence.

Ultimately I agree with the authors that no causal mechanism evidence is available here. Their discussion and speculations and descriptive evidence are generally very reasonable, but doesn't go very far to help us understand why this very surprising main effect occurs. So I am left with some doubt as to whether the unverifiable necessary conditions for quasi-experimental identification are met, I am left with doubt as to what other settings this might apply to, and I am left with doubt as to what I can learn from this.

We understand that our work provides speculation, rather than concrete causal evidence, into the mechanisms that drive the observed high-level trends. In the new limitations paragraph of the Discussion section, we have added a suggestion that future work use online lab experiments (in the absence of industry connections) to test the specific mechanisms that we speculate might drive our high-level trends.

Clearly I am out of step with the other reviewers. I will understand if the journal decides to publish the article without my endorsement. For authors' sake, I encourage the journal to make an up-or-down decision at this stage.

I wish the authors well and thank them for their research.

We thank you for your detailed and thoughtful feedback. We understand that you are left with some doubt. We are very grateful to the suggestions you have provided throughout the review process, which have pushed us to improve our writing as well as analysis (e.g. the labeling exercise). We thank you sincerely for your close reading and feedback.

Reviewer #3 (Remarks to the Author):

Overall, I appreciate the authors' engagement with the feedback. I believe the issues I raised have been appropriately addressed, and the paper is now more precise in its discussion of some of the findings.

We thank you for your attention to our revisions and for your support; your feedback and suggestions throughout the process have strengthened the paper.