The Second Order Effects of LLMs and Generative AI on Political Communication *

Kaylyn Jackson Schiff† Daniel S. Schiff‡ Bryce Dietrich§

August 27, 2024

Abstract

How will generative AI systems and large language models (LLMs) like ChatGPT impact political participation and communication? Prior work shows that legislators largely cannot distinguish between machine-generated and human-generated communication (Kreps and Kriner 2023). Moving beyond these persuasive impacts of machine-generated messages, we seek to explore broader *second order* effects of generative AI, examining its impact on political attitudes, communication, and behavior. For example, does awareness of the possible use of generative AI to draft individual emails or even to facilitate systematic use in large advocacy campaigns change the way that members of the public think about political participation? Does it alter how the public feels regarding whether their communication is valued by their elected political representatives, or their intention to reach out to politicians? To study this, we administer an online survey experiment of American adults that randomly varies the presentation of information about how LLMs like ChatGPT can be used in political communication. We examine impacts of generative AI on the perceived effectiveness of and support for using email or other communication strategies for contact with representatives, among other outcomes.

Keywords: artificial intelligence, generative AI, large language models (LLMs), survey experiment, political participation, political communication

*Thanks to participants at the Yale CSAP Summer Workshop Series for feedback on the experimental design.

†Assistant Professor, Department of Political Science, Purdue University, schiffk@purdue.edu
‡Assistant Professor, Department of Political Science, Purdue University, dschiff@purdue.edu
§Associate Professor, Department of Political Science, Purdue University, bjdietri@purdue.edu

1 Motivation and Background

The growing social penetration of generative AI—especially large language models (LLMs)—has spawned an array of public and scholarly investigations into the effects of the systems. Amongst the most pressing concerns surround the impact of generative AI on labor, privacy, intellectual property, cybersecurity, and more (Anderljung et al. 2023; Solaiman et al. 2023; Weidinger et al. 2021). In the context of political communication, risks of misinformation receive substantial attention due to the ease of use of generative AI for creating realistic but false text, images, audio, and video (Goldstein et al. 2023; Mustak et al. 2023; Vacarelu 2023). Accordingly then, a growing research agenda considers issues like incentives for creating and sharing generative AI for political purposes, technical detection methods, platform regulation efforts, effects on and attitudes of members of the public and politicians, and mitigation strategies (Kreps et al. 2022; Lee and Bissell 2023; Sharma et al. 2023). In the political arena, this work has demonstrated how AI-generated misinformation is increasingly convincing, and capable of shifting attitudes in some cases (Barari et al. 2021; Groh et al. 2022; Schiff et al. 2023).

Another critical recognition is that generative AI systems used for political purposes can have indirect, second-order, or systemic effects, in addition to their direct effects like persuading individuals that false content is real. For example, generative AI's impact on the information environment could lead individuals to be less willing to view any content as authentic, be less trusting of institutional actors, or change how they engage with co-partisans and anti-partisans (Chesney and Citron 2019; Pashentsev 2023; Schiff et al. 2023). The purpose of the study is to investigate these "second order effects" of generative AI, particularly large language models and AI chatbots like ChatGPT, on political communication. We focus on how awareness by members of the public of generative AI systems might affect their attitudes and intended behaviors regarding outreach to their representatives.

For example, does knowledge of the possible use of LLMs to generate constituent outreach messages affect the perceived ease of contacting representatives or one's intention to do so? Does it change which modes of communication (e.g., online versus in-person) are preferred, or the sense of members of the public that politicians will view their messages as authentic and meaningful? Further and in recognition of the fact that generative AI is still in its early days, we investigate whether the portrayal of LLM usage might affect attitudes. For example, does an explicit warning that these models can be used to mass generate misinformation alter public receptivity relative to a more positive framing about generative AI?

To examine these questions, we administer an online survey experiment to 2,891 American adults that randomly varies the presentation of information about how chatbots like ChatGPT or Bard can be used in outreach to politicians. We consider whether awareness of 1) isolated, individual use of machine-generated emails or 2) their systematic use in large-scale advocacy campaigns affects public perceptions and behavior regarding ease of outreach, intention to reach out, support for use of generative AI in political communication, and other outcomes.

The "Isolated Use" treatment explains (and demonstrates) that individual citizens can use generative AI to automatically create messages to send to their elected representatives. The "Systematic Use" treatment explains that political advocacy and lobbying groups can also use these tools to generate thousands of these messages easily. Finally, the "Explicit Warning" treatment builds on the systematic use treatment with a direct statement cautioning that this may lead representatives to doubt whether messages are authentic. The treatments all include the same example message composed by ChatGPT and depicted to participants about support of or opposition to an education bill. We compare the three treatments to two control conditions: 1) a pure control with no information provided before the outcome questions, and 2) the provision of a template message but without the information or indicators that it was machine-generated (providing a reference point for existing advocacy strategies that use human-generated templates and form letters to help supporters create messages).

We find that 1) across treatments, individuals support and express interest in using generative AI in political communication, 2) they anticipate that this may make email communication with representatives less effective, and 3) they are unwilling to use alternative, higher-cost contact methods. Additionally, we generally find a lack of heterogeneous effects by LLM familiarity and techno-optimism, but those with greater prior contact with representatives and greater trust in government are more opposed to machine-generated correspondence with politicians. The results encourage further investigation into the sources of support for generative AI in political communication and efforts to preserve effective pathways for constituent-representative dialogue.

2 Theory

2.1 Technology, Access, and Equity in Political Communication

Since the early days of the web, advocates have suggested that digital communication methods, such as the Internet broadly, email, and social media platforms, could expand access to politicians, increase interaction between citizens and their representatives, and effectively close longstanding equity and representation gaps (Burns et al. 2018; Hemphill and Roback 2014; Strömberg 2001). Notably, compared to the most prominent form of participation—voting—individualized constituent outreach to politicians has the advantage of precision; citizens are able to articulate the specific issues and policy positions on their mind, rather than rely on a binary vote as a rough proxy of their preferences (Schlozman et al. 2018).

Digital communication with representatives also reduces the costs, primarily time and effort, required to send messages. While emails or social media messages can be sent quickly and easily from multiple devices at any time of the day, phone calls often require time taken out of the work day and letters require physical resources (paper, envelope), time for writing, and the cost of a stamp. The downside to this, however, is that according to theories of costly signaling (Gause 2022; Hill 2022; Lohmann 1993; Spence 1973), representatives would deem messages that are more costly to send (i.e., through letters or phone calls) to be more informative. Therefore, technology may cheapen political communication.

Digital communication methods present an additional and related difference from more traditional channels: less-clear signals of identity. While letters contain addresses and phone calls display area codes, messages sent through digital methods often have less obvious indicators of identity, and some are even anonymous. This can lead public officials to wonder whether some messages come from non-constituents, political opponents, or even bots.

Furthermore, to the extent that digital participation or constituent outreach are more informative than voting or simple polls, unequal utilization could exacerbate democratic gaps as well, either directly or indirectly through advantages procured through increased access to digital tools (Nisbet 2008). A systematic review of more than 20 studies in the US, UK, Canada, and France found that more educated and younger individuals, and women, are consistently more likely to use digital communication methods to contact politicians relative to offline methods (Boulianne 2023). Moreover, recent work on public meetings via Zoom during the pandemic reveals that lower-cost participation options still often result in greater engagement by individuals that are white, older, and homeowners (Einstein et al. 2022). As such, digital communication methods may indeed facilitate closing participation gaps for areas like education and age, though other key demographic divides like gender, race, and income stubbornly continue to manifest participatory inequalities (Gallego 2014) and evidence on the overall efficacy of digital communication is mixed (Bimber 1999; Schlozman et al. 2010).

2.2 Impacts of Generative AI

It is within this context that an additional technological development—the rise of generative AI—is poised to further shape political communication. In particular, generative AI can exacerbate the costly signaling and identity challenges described above. First, generative AI provides an option for citizens to easily generate the *content* of their messages, further reducing costs over and above the benefits of increased digital access. Second, and relatedly, the possibility of machine-generated content means that the receivers of messages may be concerned about not only the identity of the sender but also about whether the content of their message is *authentic* and representative of their own voice and perspective. Therefore, while generative AI has the potential to boost political participation through significantly lowered barriers, it also has the potential to significantly cheapen communication through inauthentic, mass content.

In recent work, Kreps and Kriner 2023 employ a field experiment on state legislators to investigate the persuasiveness of machine-generated emails, and whether they can pass for human-generated. They find that legislators largely cannot distinguish between machine-generated and human-generated content. Beyond these first order, direct persuasive effects, machine-generated emails may also have broader, second order effects once individuals begin to doubt whether all emails are authentic, due to the challenges of generative AI described above. We tackle this possibly in this paper and extend related work on the second order effects of misinformation on political trust and accountability (Schiff et al. 2023; Ternovski et al. 2022). Critically, it does not have to be the case that messages are actually misinformation or that emails are actually machine generated for individuals' concerns and suspicions to alter their behavior. We devote particular attention to impacts on behavior within the realm of political communication.

2.3 Hypotheses

We test the following hypotheses related to generative AI and political communication. We use a survey experiment with three treatments corresponding to different stages or levels of the rollout of AI chatbots for communicating with representatives. The Isolated Use treatment considers isolated, individual use of LLMs for citizens to draft messages to their representatives. The Systematic Use treatment considers broader-scale usage, including through mass advocacy and lobbying campaigns. The Explicit Warning treatment builds on the Systematic Use treatment to consider a stage at which concerns are expressed regarding the authenticity of emails that representatives receive. Below, we state our hypotheses: Hypothesis Family 1 pertains to the Isolated Use treatment, Hypothesis Family 2 pertains to the Systematic Use treatment, and Hypothesis Family 3 pertains to the Explicit Warning treatment.

The hypotheses specifically consider potential benefits—reduced barriers and difficulty, increased intentions to reach out—as well as the potential drawbacks surrounding authenticity and the cheapening of communication. Furthermore, we explore support for these tools in order to address public opinion on the emerging topic and to potentially inform evolving policy conversations. As a potential remedy or solution, we also investigate individuals' willingness to engage in alternative, higher-cost communication methods. In the hypotheses below, the "A" hypothesis in each family describes an expected mechanism (loosely defined) related to the benefits and drawbacks of machine-generated emails that we expect to drive support changes.

Hypothesis 1A: The Isolated Use treatment will reduce perceptions of the difficulty of political communication relative to both control conditions.

Hypothesis 1B: The Isolated Use treatment will increase intentions to reach out to elected representatives relative to both control conditions.

Hypothesis 1C: The Isolated Use treatment will increase support for use of generative AI in political communication relative to both control conditions.

Hypothesis 2A: The Systematic Use treatment will increase perceptions that elected representatives' email inboxes contain mass generated messages, as opposed to emails personally created by individuals, relative to both control conditions.

Hypothesis 2B: The Systematic Use treatment will increase willingness to use other communication methods (beyond email) relative to both control conditions and Isolated Use.

Hypothesis 2C: The Systematic Use treatment will decrease support for generative AI in political communication relative to both control conditions and Isolated Use.

Hypothesis 3A: The Explicit Warning treatment will reduce perceptions that elected representatives view emails as authentic and valuable constituent communication, relative to both control conditions.

Hypothesis 3B: The Explicit Warning treatment will increase willingness to use other communication methods (beyond email) relative to both control conditions and Isolated Use.

Hypothesis 3C: The Explicit Warning treatment will decrease support for generative AI in political communication relative to both control conditions and Isolated Use.

3 Experimental Design

3.1 Ethics

The study was approved by the Institutional Review Boards of Yale University and Purdue University. To ameliorate deception, we include a debrief at the end of the survey to explain (especially for those in the second, template control condition) that the message that respondents saw was written by ChatGPT. To improve digital literacy, we also include a link to more information about generative AI, how it works, and what it can be used for.

3.2 Sample and Power Considerations

We recruited a sample of 2,891 respondents through Prolific to take our online survey. This sample size was selected to maximize power given budget constraints. We use Prolific as our preferred survey sampling organization due to high levels of respondent attentiveness and high quality open-ended responses measured in our prior survey experiment work, and because other prior research has found that Prolific provides high quality data (Peer et al. 2021).

We conducted a pilot study of 100 respondents in early June 2023 through Prolific. Based on responses to a few key outcome questions in the pilot study, we estimated minimum detectable effects (MDEs). Given the pilot survey design and a sample size of 2,891, we have 80% power to detect effects as small as 0.22 standard deviations for the outcome about support for ChatGPT, and as small as 0.19 standard deviations for the outcome about difficulty contacting elected representatives. This suggests that we are able to detect small to medium effects. Moreover, given preliminary results from the pilot study, we suspected that the treatments would be relatively powerful. In the pilot, we found standardized treatment effects of 0.38 to 0.92 standard deviations across treatments for the support outcome and 0 to 0.5 standard deviations across treatments for the difficulty of contact outcome. Appendix Section A.1 provides additional information about the power analysis.

3.3 Treatments

Respondents were randomly assigned to one of five groups:

- Pure Control: did not receive any information about political outreach or generative

AI before the outcome questions

- Template Control: received a template message, without information that it was machine-generated (by implication human-generated), for citizens to send to their elected representatives
- Isolated Use: received information with an example about how individual citizens can use ChatGPT to automatically generate messages to send to their elected representatives
- Systematic Use: received information with an example about how political advocacy and lobbying groups can also use these tools to generate thousands of messages easily
- Explicit Warning: received information with an example that builds on the systematic use treatment with a direct statement that this may lead representatives to doubt whether messages are authentic and to pay less attention to public outreach

The Template Control uses the following treatment wording:

> Sometimes people send **messages to their elected representatives** about topics that they care about.
>
> As an example, here's a message that someone could send to their representative:
>
> [Template Control Message Image]

The Isolated Use treatment uses the following treatment wording:

> AI tools such as ChatGPT can be used to **automatically generate written messages, like emails**. Based on a prompt with just a few words, ChatGPT and similar tools can create compelling and detailed responses that can pass as human-written.
>
> **One possible use of these AI tools is to automatically generate messages from citizens to their elected representatives.**

> As an example, here's a message written by ChatGPT based on a human-provided prompt:
>
> [Treatment Message Image]

The Systematic Use treatment uses the following treatment wording:

> AI tools such as ChatGPT can be used to **automatically generate written messages, like emails**. Based on a prompt with just a few words, ChatGPT and similar tools can create compelling and detailed responses that can pass as human-written.
>
> **One possible use of these AI tools is for political advocacy and lobbying groups to engage in mass automatically-generated advocacy campaigns sent to elected representatives. It would be possible to generate thousands of these messages easily without much human input.**
>
> As an example, here's a message written by ChatGPT based on a human-provided prompt:
>
> [Treatment Message Image]

The Explicit Warning treatment uses the following treatment wording:

> AI tools such as ChatGPT can be used to **automatically generate written messages, like emails**. Based on a prompt with just a few words, ChatGPT and similar tools can create compelling and detailed responses that can pass as human-written.
>
> **One possible use of these AI tools is for political advocacy and lobbying groups to engage in mass automatically-generated advocacy campaigns sent to elected representatives. It would be possible to generate thou-**

> **sands of these messages easily without much human input.**
>
> **Some are concerned that this will lead representatives to doubt whether messages are authentic and to pay less attention to public outreach.**
>
> As an example, here's a message written by ChatGPT based on a human-provided prompt:
>
> [Treatment Message Image]

Figures 1 and 2 show the example messages used in the template control and three treatment conditions, respectively. The three treatment conditions all use an identical image. Additionally, the template control and treatment images are the exact same, except that the template control image removes the human-generated prompt and ChatGPT logo that would indicate that the message was produced by ChatGPT.

Dear [State Representative's Name],

I am writing to express my strong support for the education funding bill that is currently under consideration in our state legislature, and to urge you to vote in favor of this important piece of legislation.

As you know, education is the foundation of our society, and it is essential that we provide our children with the best possible education in order to ensure that they have the skills and knowledge necessary to succeed in today's complex world. Unfortunately, our state's education system is currently facing a serious funding crisis, and many schools are struggling to provide the quality education that our children deserve.

The education funding bill that is currently before the legislature would help to address this crisis by providing much-needed funding to our schools. This funding would help to ensure that our schools have the resources they need to hire the best teachers, purchase up-to-date textbooks and instructional materials, and provide students with access to the latest technology and other educational resources.

In addition, the education funding bill would help to reduce class sizes, which is critical for ensuring that every student receives the individual attention and support they need to succeed. By providing more funding to our schools, we can ensure that our children are prepared for success in the 21st century economy, and that they have the skills and knowledge they need to thrive.

I urge you to support this important piece of legislation, and to do everything in your power to ensure that it is passed into law. Our children's future depends on it.

Thank you for your attention to this matter.

Sincerely,

[Your Name]

Figure 1: Template Control Message

I would like to write a letter to my state representative asking them to support the education funding bill. Please include details about why my representative should support the bill.

Dear [State Representative's Name],

I am writing to express my strong support for the education funding bill that is currently under consideration in our state legislature, and to urge you to vote in favor of this important piece of legislation.

As you know, education is the foundation of our society, and it is essential that we provide our children with the best possible education in order to ensure that they have the skills and knowledge necessary to succeed in today's complex world. Unfortunately, our state's education system is currently facing a serious funding crisis, and many schools are struggling to provide the quality education that our children deserve.

The education funding bill that is currently before the legislature would help to address this crisis by providing much-needed funding to our schools. This funding would help to ensure that our schools have the resources they need to hire the best teachers, purchase up-to-date textbooks and instructional materials, and provide students with access to the latest technology and other educational resources.

In addition, the education funding bill would help to reduce class sizes, which is critical for ensuring that every student receives the individual attention and support they need to succeed. By providing more funding to our schools, we can ensure that our children are prepared for success in the 21st century economy, and that they have the skills and knowledge they need to thrive.

I urge you to support this important piece of legislation, and to do everything in your power to ensure that it is passed into law. Our children's future depends on it.

Thank you for your attention to this matter.

Sincerely,

[Your Name]

Figure 2: Treatment Message

3.4 Outcome Measures

We measure the following outcomes, designed to allow for tests of our main hypotheses and also allow for some exploratory analyses:

1. In your opinion, how likely are elected representatives to **pay attention to constituents** who reach out through the following contact methods? [Rank in-person meeting, hand-written letter, phone call, email, social media message]
2. In your opinion, how **effective** are these contact methods in **influencing representatives' opinions or decisions**? [Rank in-person meeting, hand-written letter, phone call, email, social media message]
3. It is **easy to contact** elected political representatives. [Strongly disagree to strongly agree]
4. About what percentage of **elected political representatives' emails do you think are mass generated messages**, as opposed to personally created by individuals? [Slider from 0 to 100]
5. Elected political representatives **view the emails they receive as valuable input for their work**. [Strongly disagree to strongly agree]
6. In the future, how often do you **plan to reach out** to your elected political representatives (mayor, city council member, state representative, or member of Congress)? [Never to Frequently]
7. Which of the following communication methods **would you consider using** to reach out to your elected political representative? [Select all that apply: in-person meeting, hand-written letter, phone call, email, social media message]
8. To what extent do you **support members of the public using AI chatbots** like ChatGPT for communication with elected political representatives? [Strongly oppose to strongly support]
9. To what extent do you **support advocacy and lobbying groups using AI chat-**

bots like ChatGPT for communication with elected political representatives? [Strongly oppose to strongly support]

10. **Would you consider using AI chatbots** like ChatGPT to communicate with your elected political representatives? [Definitely not to definitely yes]
11. What are your thoughts on the use of AI chatbots like ChatGPT for communicating with politicians? [Open-ended]

4 Analysis Strategy

4.1 Hypothesis Testing

Table 1 describes the specific treatments and outcomes that we use to test our primary hypotheses, referencing numbers in the prior list of outcome questions.

Hypothesis	**Treatment**	**Outcome(s)**
1A	Isolated Use	3 - Ease of contact
1B	Isolated Use	6 - Intentions to reach out to representatives
1C	Isolated Use	8,9 - Support for use of ChatGPT
2A	Systematic Use	4 - Mass generated emails
2B	Systematic Use	7 - Use of other outreach methods
2C	Systematic Use	8,9 - Support for use of ChatGPT
3A	Explicit Warning	4,5 - Mass generated emails, value input via email
3B	Explicit Warning	7 - Use of other outreach methods
3C	Explicit Warning	8,9 - Support for use of ChatGPT

Table 1: Hypothesis Testing

We use the following regression specification for each outcome of interest, focusing on the main treatment(s) of interest per Table 1 above, and using either the pure control or the template control as the reference group (two separate specifications):

$$Outcome = \beta_0 + \beta_1 template + \beta_2 isolated + \beta_3 systematic + \beta_4 explicit + \gamma \mathbf{X} + \epsilon \qquad (1)$$

where *Outcome* is the outcomes of interest; *template*, *isolated*, *systematic*, and *explicit* refer to the treatments; $\mathbf{X}$ is a vector of covariates, and ϵ is the error. In this example, the pure control group is the reference group.

For covariates, we include:

- A pre-treatment measure of prior contact with elected political representatives: "How often do you reach out to your elected political representatives (mayor, city council member, state representative, or member of Congress)?"
- A measure of LLM familiarity: "Before today, how familiar were you with large language models (LLMs), or AI chatbots, like ChatGPT, Bing, or Bard?"
- A measure of perceived risks/benefits of AI: "Thinking about society generally, the benefits of artificial intelligence (AI) outweigh the risks."
- A measure of political knowledge: "How long is the term of a United States Senator?"
- A measure of trust in government: "How often can you trust government to do what is right?"
- Standard demographic questions: gender, age, race, education, income, region, and political party

4.2 Exploratory Analyses

We explore heterogeneous effects for all of our primary hypotheses by 1) familiarity with LLMs and 2) prior contact with elected political representatives. For the hypotheses related to support for ChatGPT (Hypotheses 1C, 2C, and 3C), we additionally explore heterogeneous effects by perceived risks/benefits of AI.

5 Results

5.1 Treatments Compared to Pure Control

We find that the Isolated Use treatment—describing how *individual* citizens can use AI chatbots like ChatGPT to generate messages to send to their representatives—significantly increased support for both individual and mass use of LLMs for political communication, by 0.44 and 0.34 standard deviations, respectively. However, this increase in general support was not accompanied by changes in outcomes more closely related to political behavior. That is, the Isolated Use treatment did not increase perceptions that it is easy to contact political representatives ($p = 0.729$), nor did it significantly increase individuals' willingness to reach out to their representatives ($p = 0.221$).

Therefore, while one promising possibility is that LLMs could serve as as a new technological solution to low democratic participation, the evidence suggests generative AI may fall short of those aims. Moreover, the increased support appears to not have been generated through the channel of increased contact ease, nor by perceived effectiveness in terms of persuading representatives. This contrasts with core elements of the technology acceptance model (TAM) (Marangunić and Granić 2015), which identifies perceived ease of use and usefulness as key factors driving use intention. More work is needed to understand the drivers of individuals' support for generative AI, in political communication.

As for the Systematic Use treatment emphasizing mass LLM use in advocacy and lobbying campaigns, we find several unexpected results. Most notably, the treatment decreased, rather than increased, individuals' perceptions of the proportion of representatives' emails that are mass generated (0.43 standard deviations, $p < 0.001$). It is possible that the treatment changed respondents' views on what constitutes "mass generated" emails, or that respondents anticipated that increased detection and filtering efforts would accompany these changes. This may also explain the additional unexpected result that support for both individual and mass use of LLMs for political communication significantly increased after this treatment by 0.24 and 0.20 standard deviations, respectively. Note, however, that these support gains are smaller than for the isolated use treatment. Finally, we also find that the systematic use treatment did not change respondents' willingness to use other methods (letters, in-person meetings, or phone calls) to contact representatives ($p = 0.685$).

The Explicit Warning treatment highlighting concerns about authenticity did shift responses related to hypothesized "mechanism" outcomes in the anticipated direction—the treatment reduced perceptions that representatives pay attention to emails from constituents (0.18 standard deviations, $p = 0.003$) and also reduced perceptions that emails are effective in influencing representatives' opinions (0.13 standard deviations, $p = 0.048$).[1] However, as with the Systematic Use treatment, this did not translate into reduced support for LLMs in political communication.

That is, despite concerns that generative AI would reduce politician attentiveness and email effectiveness, individuals in the Explicit Warning treatment group still supported individual and mass use of LLMs to a greater extent that those in the pure control group. With effect sizes of 0.18 and 0.21 standard deviations, respectively, the impacts are again smaller than for the Isolated Use treatment. Nevertheless, even when individuals received explicit information about risks, which in part led them to identify some downsides of LLMs, they still overall exhibited increased support for use of generative AI compared to control. Finally, we again find no effect on respondents' willingness to use alternative contact methods ($p = 0.824$), providing little overall evidence for substitution between communication modes.

The primary regression results are presented in Table 2, and Table 3 reports results for our main hypothesis tests using the pure control group as the comparison group. Of note, all

[1] Interestingly, the Explicit Warning treatment increased perceptions that representatives value emails, against expectations and in contrast to the negative impacts on email attention and effectiveness. Respondents may have interpreted this question as referring to only authentic emails and/or only personal emails from constituents.

of the treatments increased individuals' stated willingness to personally use AI chatbots to reach out to their representatives, with the Isolated Use treatment generating the largest effects, largely in favor of LLM use. Combined with the results above, this suggests that while 1) individuals anticipate generative AI may engender at least some downsides related to email communication with representatives (either decreasing politician attention to emails, or decreasing email effectiveness), they 2) broadly support and may be willing to use generative AI in political communication, and 3) are unwilling to substitute for alternative, higher-cost contact methods. These tensions create a quandary for political communication in the age of generative AI, and suggest the need for further efforts to preserve and build effective effective pathways for constituent-representative dialogue that minimize problematic compromises.

	Email Attention	Email Effective	Contact Ease	Email Spam	Value Email	Contact Freq	Other Contact	Support (Indiv)	Support (Mass)	Personal Use
(Intercept)	1.899***	2.201***	1.328***	63.778***	0.628*	0.385**	0.473***	1.296***	0.725**	0.174
	(0.260)	(0.268)	(0.293)	(6.684)	(0.279)	(0.121)	(0.129)	(0.253)	(0.270)	(0.286)
Template Control	−0.068	−0.073	0.011	0.058	0.090	0.034	−0.015	0.009	0.006	0.077
	(0.065)	(0.060)	(0.073)	(1.334)	(0.062)	(0.035)	(0.029)	(0.064)	(0.061)	(0.068)
Isolated Use	−0.089	−0.065	−0.025	−5.040***	0.114+	0.042	−0.015	0.496***	0.381***	0.469***
	(0.065)	(0.062)	(0.074)	(1.426)	(0.061)	(0.035)	(0.030)	(0.061)	(0.062)	(0.070)
Systematic Use	−0.155*	−0.096	0.012	−9.485***	0.051	0.103**	0.012	0.269***	0.221***	0.281***
	(0.065)	(0.061)	(0.073)	(1.421)	(0.063)	(0.036)	(0.029)	(0.063)	(0.063)	(0.070)
Explicit Warning	−0.197**	−0.129*	0.024	−10.323***	0.109+	0.065+	−0.007	0.204**	0.237***	0.277***
	(0.067)	(0.065)	(0.072)	(1.457)	(0.062)	(0.035)	(0.030)	(0.063)	(0.063)	(0.069)
Num.Obs.	2481	2418	2525	2520	2525	2525	2525	2525	2525	2525
R2	0.041	0.029	0.082	0.057	0.147	0.487	0.057	0.255	0.252	0.295
Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Treatment Effects Relative to Pure Control

Hypothesis	**Treatment**	**Outcome**	**Prediction**	**Coefficient**	**p-Value**
1A	Isolated Use	Ease of contact	+	-0.025	0.729
1B	Isolated Use	Reach out to reps	+	0.042	0.221
1C	Isolated Use	Support for indiv. LLM use	+	0.496	< 0.001
1C	Isolated Use	Support for mass LLM use	+	0.381	< 0.001
2A	Systematic Use	Mass generated emails	+	-9.485	< 0.001
2B	Systematic Use	Other contact methods	+	0.012	0.685
2C	Systematic Use	Support for indiv. LLM use	-	0.269	< 0.001
2C	Systematic Use	Support for mass LLM use	-	0.221	< 0.001
3A	Explicit Warning	Email attention	-	-0.197	0.003
3A	Explicit Warning	Email effectiveness	-	-0.129	0.048
3B	Explicit Warning	Other contact methods	+	-0.007	0.824
3C	Explicit Warning	Support for indiv. LLM use	-	0.204	0.001
3C	Explicit Warning	Support for mass LLM use	-	0.237	< 0.001

Table 3: Results of Hypothesis Testing, Pure Control Comparison

Our experiment also employed an alternative comparison group in which some respondents were provided with the same template message but absent any mention of LLMs and their potential use for political communication. Yet, we find no significant differences between this template control group and the pure control group, as shown in Appendix A.2. This increases our confidence that the treatment effects identified here are driven by information about the potential uses of generative AI, rather than the mere presentation of a template email.

5.2 Heterogeneous Effects

In addition to testing our main hypotheses, we also explored heterogeneous effects, given the possibility that the treatments about generative AI could produce differential effects related to, for example, digital literacy. First, we examined how prior familiarity with LLMs impacts the effects of the treatments on respondents, on the assumption that increased familiarity could mediate individual comfort using generative AI for political outreach, either positively or negatively. Individuals in the Systematic Use and Explicit Warning treatment groups exhibit few differential effects. However, individuals in the Isolated Use group with greater prior familiarity did exhibit increased support of and intention to personally use generative AI, as shown in Figure 3. This could result from, for example, a better appreciation or interest in individual use of these tools. Overall, however, the uniformity across subgroups for most outcomes measured suggests that many important reactions to generative AI are not strongly driven by differences between early and late adopters.

Next, we examined differential effects by prior contact with elected representatives. We again find very few differences except that individuals with greater prior contact with political representatives are less supportive of individuals using AI chatbots to reach out, as seen in Figure 4. These individuals may have come to value their capacity to produce individualized messages and thus be opposed to lowering the barriers to and cheapening political communication. This provides some modest evidence of an equalizing effect, in that individ-

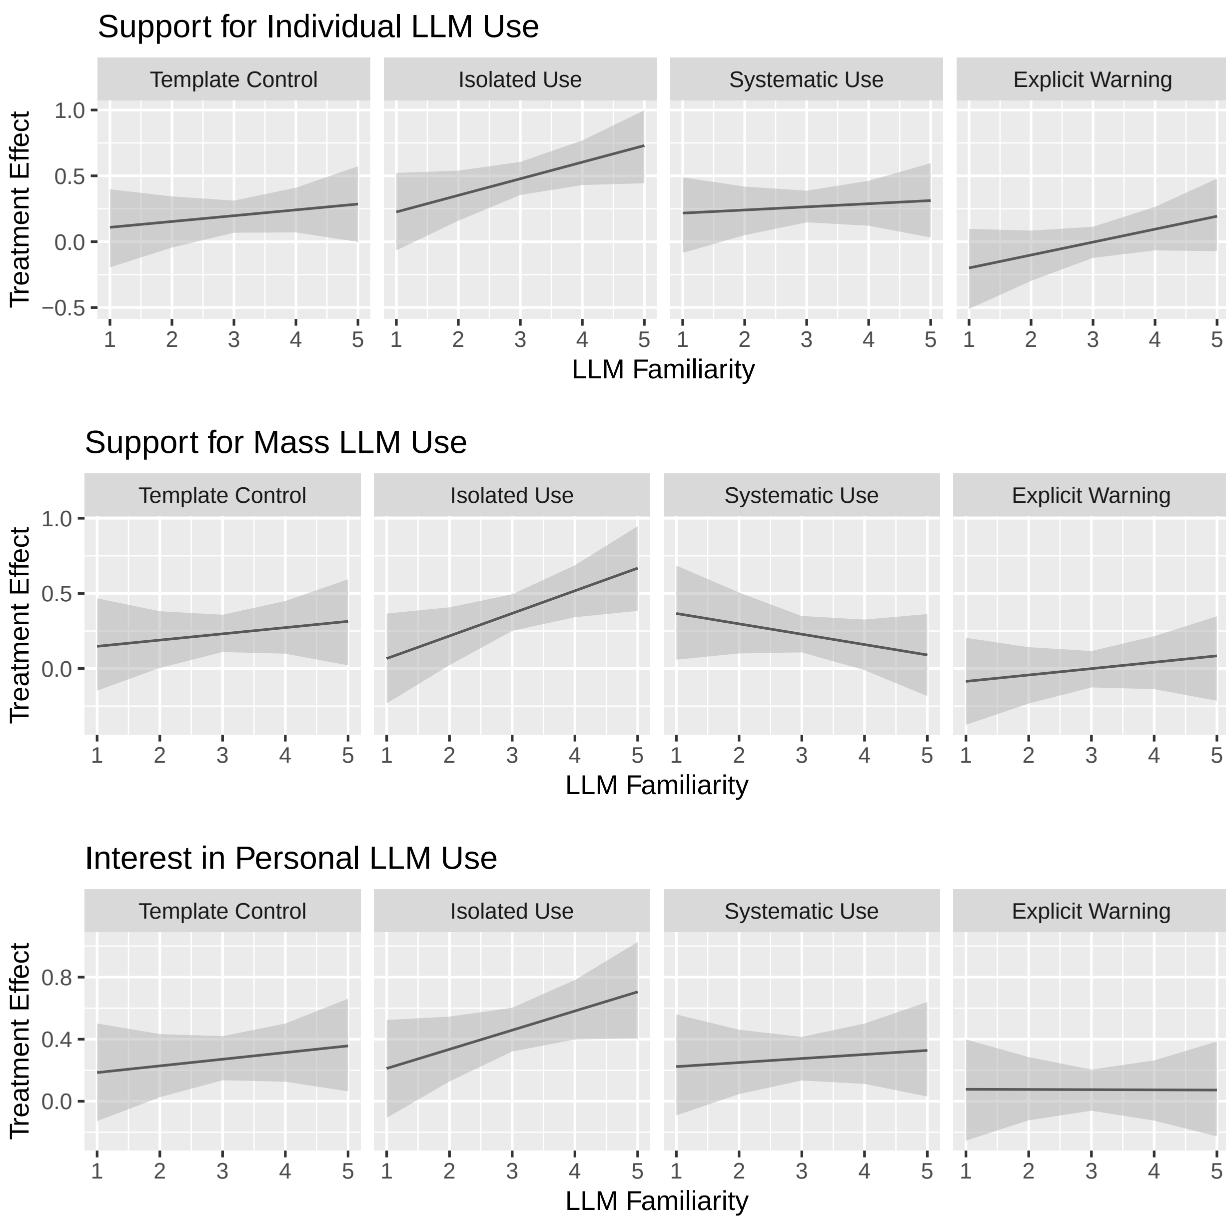

Figure 3: Heterogeneous Effects by LLM Familiarity

uals with greater prior access may feel that their advantages are lessened. However, we fail to see increased interest on the part of individuals with lower incomes or educational levels in using generative AI to support politician outreach, relative to individuals with higher income or educational levels.

Third, we explored heterogeneous effects by perceived benefits versus risks of AI, motivated

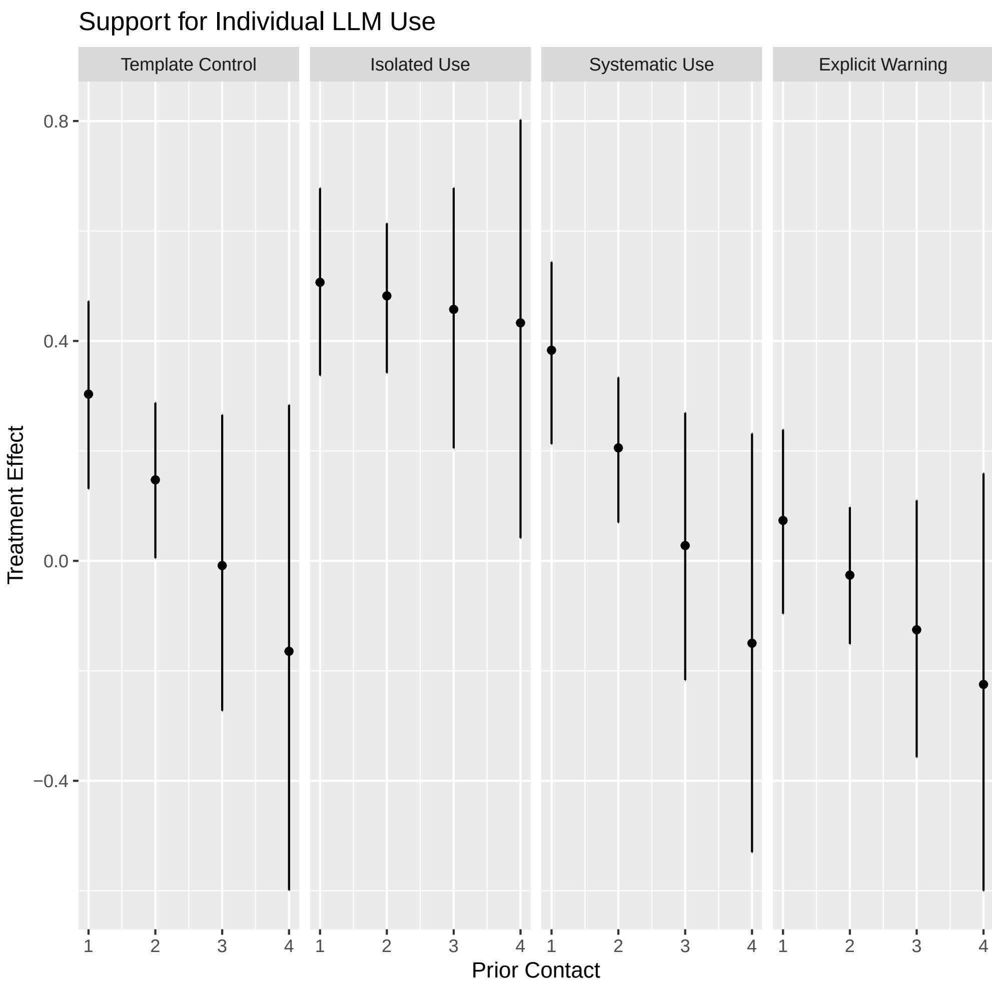

Figure 4: Heterogeneous Effects by Prior Contact with Representatives

by prior work showing that underlying attitudes toward technology like perceived benefits and techno-optimism can shift support for use of AI (O'Shaughnessy et al. 2023). While those who agree that the benefits of AI outweigh the risks are more likely to indicate that it's easier to contact elected officials, to support individual and mass LLM use for political communication, and to be personally willing to use AI chatbots for reaching out to representatives, we find no significant interaction effects with our treatments.

Therefore, while techno-optimism with respective to AI may explain underlying political cleavages and differences in policy preferences, it does not appear to differentially impact how individuals respond to information about political communication in the age of generative AI.

One possibility is that individuals with high levels of techno-optimism were already familiar with LLMs[2] such that the information provided was not new, did not lead to significant updating, and therefore did not produce differences relative to the control group. Assessing this will require further research, particularly over time as public familiarity and receptivity toward generative AI evolve.

Finally, we investigated the impact of trust in government on respondents' reactions to the treatments. Individuals with higher trust in government were indeed significantly more likely to respond to the treatments with increased perceptions that emails are effective in influencing representatives' opinions. Higher trust in government was also associated with reduced support for individual LLM use for political communication across the treatment groups. Lastly, the Systematic Use treatment also led to relatively reduced support for mass LLM use amongst those with higher trust in government. These impacts are depicted in Figure 5. In sum, those with higher trust in government appear to expect use of and effectiveness of constituent outreach using generative AI, and are thus relatively more concerned about the risk of AI chatbots undermining this important democratic channel.

6 Conclusion

As use of generative AI extends into many domains of social and political life, we explore whether AI chatbots for communicating with elected representatives might realize the long-heralded benefits of technology for efficient political participation or might cheapen political communication and amplify concerns over the authenticity of political interactions. We use a survey experiment of 2,891 American adults to randomly vary the presentation of information about how chatbots like ChatGPT or Bard can be used in outreach to politicians. We find that 1) across treatments, individuals support and express interest in using generative AI in political communication, 2) they anticipate that this may make email communication

[2]We find a positive and moderately substantial correlation between LLM familiarity and agreeing that the benefits of AI outweigh the risks ($r = 0.38$).

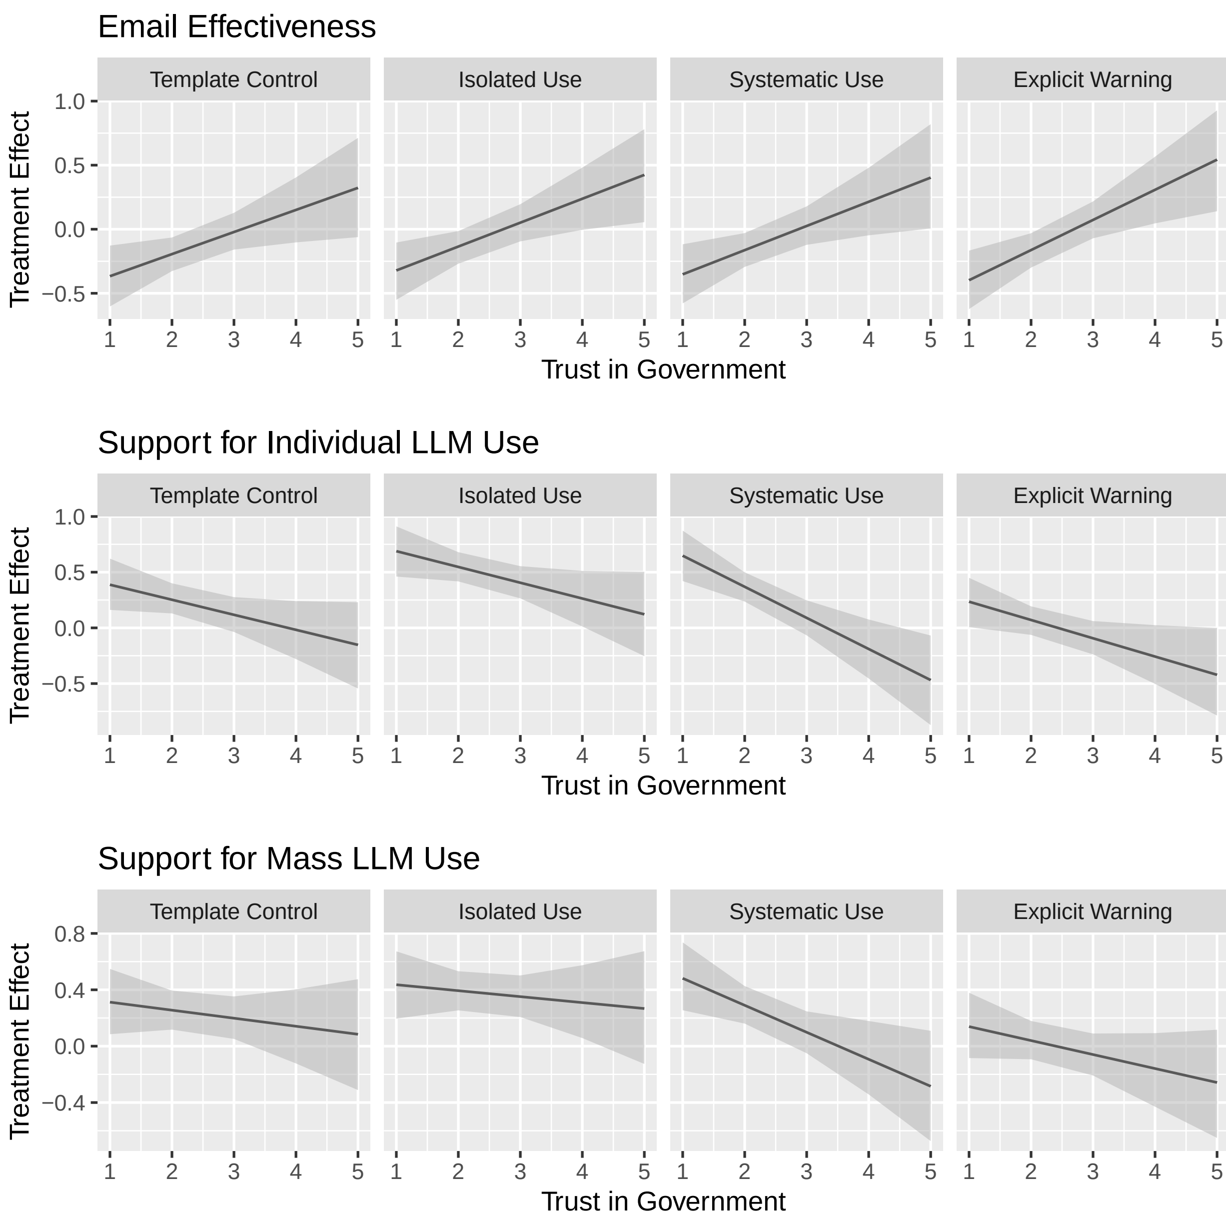

Figure 5: Heterogeneous Effects by Trust in Government

with representatives less effective, and 3) they are unwilling to use alternative, higher-cost contact methods. We do not find benefits in the form of increased perceptions that it is easy to communicate with representatives and we also do not find increased willingness to reach out to representatives. Additionally, we generally find a lack of heterogeneous effects by LLM familiarity and techno-optimism, but those with greater prior contact with representatives and greater trust in government are more opposed to machine-generated correspondence with politicians. The results encourage further investigation into the sources of support for generative AI in political communication and efforts to preserve effective pathways for constituent-representative dialogue.

References

Anderljung, Markus et al. (July 11, 2023). *Frontier AI Regulation: Managing Emerging Risks to Public Safety*. arXiv: 2307.03718[cs]. (Visited on 07/13/2023).

Barari, Soubhik, Christopher Lucas, and Kevin Munger (Jan. 13, 2021). *Political Deepfakes Are As Credible As Other Fake Media And (Sometimes) Real Media*. DOI: 10.31219/osf.io/cdfh3. (Visited on 10/20/2022).

Bimber, Bruce (1999). "The Internet and Citizen Communication With Government: Does the Medium Matter?" In: *Political Communication* 16.4, pp. 409–428. ISSN: 1058-4609, 1091-7675. DOI: 10.1080/105846099198569. (Visited on 08/24/2023).

Boulianne, Shelley (2023). "Participatory Inequality Across Countries: Contacting Public Officials Online and Offline". In: *Social Science Computer Review* 41.4, pp. 1336–1362. ISSN: 0894-4393, 1552-8286. DOI: 10.1177/08944393211071067. (Visited on 08/24/2023).

Burns, Nancy et al. (2018). "What's happened to the gender gap in political participation?: How might we explain it?" In: *100 Years of the Nineteenth Amendment: An Appraisal of Women's Political Activism*. Oxford University Press, pp. 69–104. DOI: 10.1093/oso/9780190265144.003.0004. (Visited on 08/24/2023).

Chesney, Robert and Danielle Citron (2019). "Deepfakes and the New Disinformation War: The Coming Age of Post-Truth Geopolitics". In: *Foreign Affairs* 98, p. 147.

Einstein, Katherine Levine et al. (2022). "Still Muted: The Limited Participatory Democracy of Zoom Public Meetings". In: *Urban Affairs Review*. ISSN: 1078-0874. DOI: 10.1177/10780874211070494.

Gallego, Aina (2014). "Unequal Political Participation Worldwide". In: Cambridge University Press. ISBN: 978-1-139-15172-6 978-1-107-02353-6. DOI: 10.1017/CBO9781139151726. (Visited on 08/24/2023).

Gause, LaGina (2022). "Revealing Issue Salience via Costly Protest: How Legislative Behavior Following Protest Advantages Low-Resource Groups". In: *British Journal of Political Science* 52.1, 259–279. ISSN: 0007-1234, 1469-2112. DOI: 10.1017/S0007123420000423.

Goldstein, Josh A. et al. (Jan. 11, 2023). *Forecasting Potential Misuses of Language Models for Disinformation Campaigns—and How to Reduce Risk*. Brookings Institution. (Visited on 03/20/2023).

Groh, Matthew et al. (Jan. 4, 2022). "Deepfake detection by human crowds, machines, and machine-informed crowds". In: *Proceedings of the National Academy of Sciences* 119.1. Publisher: Proceedings of the National Academy of Sciences, e2110013119. DOI: 10.1073/pnas.2110013119. (Visited on 10/19/2022).

Hemphill, Libby and Andrew J. Roback (2014). "Tweet Acts: How Constituents Lobby Congress via Twitter". In: *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. CSCW '14. New York, NY, USA: Association for Computing Machinery, pp. 1200–1210. ISBN: 978-1-4503-2540-0. DOI: 10.1145/2531602.2531735. (Visited on 08/24/2023).

Hill, Seth J (2022). "A Theory of Intensity, Electoral Competition, and Costly Political Action". In: *The Journal of Politics* 84.1, pp. 291–303.

Kreps, Sarah and Douglas L. Kriner (2023). "The potential impact of emerging technologies on democratic representation: Evidence from a field experiment". In: *New Media Society*, p. 14614448231160526.

Kreps, Sarah, R. Miles McCain, and Miles Brundage (2022). "All the News That's Fit to Fabricate: AI-Generated Text as a Tool of Media Misinformation". In: *Journal of Experimental Political Science* 9.1. Publisher: Cambridge University Press, pp. 104–117. ISSN: 2052-2630, 2052-2649. DOI: 10.1017/XPS.2020.37. (Visited on 04/05/2022).

Lee, Jiyoung and Kim Bissell (Apr. 18, 2023). "User agency–based versus machine agency–based misinformation interventions: The effects of commenting and AI fact-checking labeling on attitudes toward the COVID-19 vaccination". In: *New Media & Society*. Publisher: SAGE Publications, p. 14614448231163228. ISSN: 1461-4448. DOI: 10.1177/14614448231163228. (Visited on 04/24/2023).

Lohmann, Susanne (1993). "A Signaling Model of Informative and Manipulative Political Action". In: *The American Political Science Review* 87.2, 319–333. ISSN: 0003-0554. DOI: 10.2307/2939043.

Marangunić, Nikola and Andrina Granić (2015). "Technology Acceptance Model: A Literature Review from 1986 to 2013". In: *Universal Access in the Information Society* 14.1, pp. 81–95. ISSN: 1615-5297. DOI: 10.1007/s10209-014-0348-1. (Visited on 08/24/2023).

Mustak, Mekhail et al. (Jan. 1, 2023). "Deepfakes: Deceptions, mitigations, and opportunities". In: *Journal of Business Research* 154, p. 113368. ISSN: 0148-2963. DOI: 10.1016/j.jbusres.2022.113368. (Visited on 11/03/2022).

Nisbet, E. C. (2008). "Media Use, Democratic Citizenship, and Communication Gaps in a Developing Democracy". In: *International Journal of Public Opinion Research* 20.4, pp. 454–482. ISSN: 0954-2892, 1471-6909. DOI: 10.1093/ijpor/edn043. (Visited on 08/24/2023).

O'Shaughnessy, Matthew R et al. (Apr. 2023). "What governs attitudes toward artificial intelligence adoption and governance?" In: *Science and Public Policy* 50.2. 10.31219/osf.io/pkeb8, pp. 161–176. ISSN: 0302-3427. DOI: 10.1093/scipol/scac056. (Visited on 10/27/2022).

Pashentsev, Evgeny (2023). "The Malicious Use of Deepfakes Against Psychological Security and Political Stability". In: *The Palgrave Handbook of Malicious Use of AI and Psychological Security*. Ed. by Evgeny Pashentsev. Cham: Springer International Publishing, pp. 47–80. ISBN: 978-3-031-22552-9. DOI: 10.1007/978-3-031-22552-9_3. (Visited on 07/06/2023).

Peer, Eyal et al. (2021). "Data quality of platforms and panels for online behavioral research". en. In: *Behavior Research Methods* 54.4, 1643–1662.

Schiff, Kaylyn Jackson, Daniel S. Schiff, and Natalia Bueno (Aug. 8, 2023). *The Liar's Dividend: The Impact of Deepfakes and Fake News on Trust in Political Discourse*. DOI: 10.31235/osf.io/x43ph. (Visited on 08/10/2023).

Schlozman, Kay Lehman, Henry E. Brady, and Sidney Verba (2018). "Unequal and Unrepresented: Political Inequality and the People's Voice in the New Gilded Age". In:

Unequal and Unrepresented. Princeton University Press. ISBN: 978-1-4008-9036-1. DOI: 10.23943/9781400890361. (Visited on 08/24/2023).

Schlozman, Kay Lehman, Sidney Verba, and Henry E. Brady (2010). "Weapon of the Strong? Participatory Inequality and the Internet". In: *Perspectives on Politics* 8.2, pp. 487–509. ISSN: 1537-5927, 1541-0986. DOI: 10.1017/S1537592710001210. (Visited on 08/24/2023).

Sharma, Isha et al. (Jan. 1, 2023). "Examining the motivations of sharing political deepfake videos: the role of political brand hate and moral consciousness". In: *Internet Research* ahead-of-print (ahead-of-print). ISSN: 1066-2243. DOI: 10.1108/INTR-07-2022-0563. (Visited on 02/23/2023).

Solaiman, Irene et al. (June 12, 2023). *Evaluating the Social Impact of Generative AI Systems in Systems and Society*. arXiv: 2306.05949[cs]. (Visited on 06/15/2023).

Spence, Michael (Aug. 1973). "Job Market Signaling". en. In: *The Quarterly Journal of Economics* 87.3, p. 355. ISSN: 00335533. DOI: 10.2307/1882010. (Visited on 01/06/2022).

Strömberg, David (May 1, 2001). "Mass media and public policy". In: *European Economic Review*. 15th Annual Congress of the European Economic Association 45.4, pp. 652–663. ISSN: 0014-2921. DOI: 10.1016/S0014-2921(01)00106-4. (Visited on 09/24/2020).

Ternovski, John, Joshua Kalla, and Peter Aronow (2022). "The Negative Consequences of Informing Voters about Deepfakes: Evidence from Two Survey Experiments". en. In: *Journal of Online Trust and Safety* 1.22.

Vacarelu, Marius (2023). "Malicious Use of Artificial Intelligence in Political Campaigns: Challenges for International Psychological Security for the Next Decades". In: *The Palgrave Handbook of Malicious Use of AI and Psychological Security*. Ed. by Evgeny Pashentsev. Cham: Springer International Publishing, pp. 203–230. ISBN: 978-3-031-22552-9. DOI: 10.1007/978-3-031-22552-9_8. (Visited on 07/06/2023).

Weidinger, Laura et al. (Dec. 8, 2021). *Ethical and social risks of harm from Language Models*. DOI: 10.48550/arXiv.2112.04359. arXiv: 2112.04359[cs]. (Visited on 01/03/2023).

A Appendix

A.1 Power

We performed power calculations for key comparisons of interest using 200 simulations each and based on constructing a synthetic dataset of 2,891 respondents using random sampling (by treatment group, and with replacement) from the pilot dataset. For each simulation, we used the synthetic dataset to calculate treatment effects through covariate-unadjusted regression, and then we identified power as the proportion of simulations in which we detect a significant effect at the 5% level. Table A1 summarizes the results and suggests that the study is well powered for most comparisons of interest.

Comparison	**Power**
Email Effectiveness, Isolated Use vs. Control	0.965
Email Effectiveness, Explicit Warning vs. Control	0.135
Contact Ease for Familiar Subgroup, Isolated Use vs. Control	1
Contact Ease for Familiar Subgroup, Explicit Warning vs. Control	1
LLMs Support, Isolated Use vs. Control	1
LLMs Support, Explicit Warning vs. Control	1

Table A1: Power for Comparisons of Interest with Sample Size of 2,891

We do note, however, that our pilot study included only four groups rather than five. We did not include a pure control in the pilot, but do so for the main study. Additionally, we note that we modified the wording of many of our outcome questions after fielding the pilot. We believe that this increased clarity, but we also acknowledge that this reduces our confidence in the MDE and power calculations based on the pilot.

A.2 Treatments Compared to Template Control

We also utilized also employed an alternative comparison group in which some respondents were provided with the same template message but absent any mention of LLMs and their potential use for political communication. Yet, we find no significant differences between this template control group and the pure control group, and therefore, the results in Tables

A2 and A3 for comparisons to the template control are highly similar to the results in Tables 2 and 3 for comparisons to the pure control.

This suggests two potentially important findings. First, template messages, relative to no encouragement or support, may be ineffective in shifting individuals' attitudes or intentions toward communicating with elected officials, in contrast to some prior findings. Second, template messages produced by individuals in advocacy or lobbying groups as a more "personal" or "human" alternative to machine-generated content may not be a viable solution to generative AI concerns. That is, these classic strategies may not deter citizens from supporting or using LLMs for political communication, and they may not change perceptions of how representatives will respond to email in this new era.

	Email Attention	Email Effective	Contact Ease	Email Spam	Value Email	Contact Freq	Other Contact	Support (Indiv)	Support (Mass)	Personal Use
(Intercept)	1.831***	2.128***	1.339***	63.837***	0.717*	0.419***	0.458***	1.305***	0.731**	0.251
	(0.260)	(0.267)	(0.298)	(6.649)	(0.280)	(0.119)	(0.130)	(0.256)	(0.272)	(0.285)
Pure Control	0.068	0.073	−0.011	−0.058	−0.090	−0.034	0.015	−0.009	−0.006	−0.077
	(0.065)	(0.060)	(0.073)	(1.334)	(0.062)	(0.035)	(0.029)	(0.064)	(0.061)	(0.068)
Isolated Use	−0.022	0.008	−0.036	−5.099***	0.024	0.009	0.0004	0.486***	0.375***	0.393***
	(0.063)	(0.060)	(0.074)	(1.389)	(0.061)	(0.035)	(0.030)	(0.063)	(0.062)	(0.071)
Systematic Use	−0.087	−0.023	0.001	−9.543***	−0.039	0.069+	0.027	0.260***	0.215***	0.204**
	(0.063)	(0.059)	(0.073)	(1.384)	(0.062)	(0.037)	(0.029)	(0.066)	(0.064)	(0.072)
Explicit Warning	−0.130*	−0.056	0.013	−10.381***	0.019	0.031	0.008	0.195**	0.231***	0.200**
	(0.066)	(0.063)	(0.073)	(1.419)	(0.062)	(0.036)	(0.030)	(0.065)	(0.064)	(0.071)
Num.Obs.	2481	2418	2525	2520	2525	2525	2525	2525	2525	2525
R2	0.041	0.029	0.082	0.057	0.147	0.487	0.057	0.255	0.252	0.295
Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table A2: Treatment Effects Relative to Template Control

Hypothesis	**Treatment**	**Outcome**	**Prediction**	**Coefficient**	**p-Value**
1A	Isolated Use	Ease of contact	+	-0.036	0.623
1B	Isolated Use	Reach out to reps	+	0.009	0.807
1C	Isolated Use	Support for indiv. LLM use	+	0.486	< 0.001
1C	Isolated Use	Support for mass LLM use	+	0.375	< 0.001
2A	Systematic Use	Mass generated emails	+	-9.543	< 0.001
2B	Systematic Use	Other contact methods	+	0.027	0.360
2C	Systematic Use	Support for indiv. LLM use	-	0.260	< 0.001
2C	Systematic Use	Support for mass LLM use	-	0.215	< 0.001
3A	Explicit Warning	Email attention	-	-0.130	0.049
3A	Explicit Warning	Email effectiveness	-	-0.056	0.373
3B	Explicit Warning	Other contact methods	+	0.008	0.779
3C	Explicit Warning	Support for indiv. LLM use	-	0.195	0.003
3C	Explicit Warning	Support for mass LLM use	-	0.231	< 0.001

Table A3: Results of Hypothesis Testing, Template Control Comparison