

BallotBot: Can AI Strengthen Democracy?*

Elliott Ash¹, Sergio Galletta¹, Giacomo Opocher²

¹*ETH Zürich*

²*Università di Bologna*

Abstract

This study explores the potential for AI-powered chatbots to strengthen democracy by boosting political knowledge and engagement through better access to political information. We develop and evaluate BallotBot, an AI chatbot with access to official voter guide information from the November 2024 referendums in California. In a pre-registered three-wave survey experiment in the weeks around election day, participants (California voters) were randomly assigned to use either BallotBot or a traditional digital voter guide to answer questions about ballot initiatives. BallotBot access improved participants' ability to answer in-depth questions, reduced overconfidence, lowered the perceived cost of acquiring information for less-informed participants, and fostered greater engagement with political information. However, it had no clear effect on self-reported turnout or the direction of voting.

Keywords: Voter information, Survey Experiment, Direct Democracy, LLMs

JEL Classification: D72, D83, L82

*We acknowledge valuable comments and suggestions from Maria Bigoni, Enrico Cantoni, Leandro De Magalhaes, Nicola Lacetera, Nicola Mastrococco, Massimo Morelli, David Yanagizawa-Drott, and participants of the AI + Economics Workshop in Zurich, the seminar at Queen Mary University of London, the AI Lab for Human and Social Sciences at the Centre d'Économie de la Sorbonne (CES), the University of Bologna Internal Seminars series, and the Bristol Applied Economics Seminar. We thank Tim Ehrensperger for exceptional research assistance. The study received ethics approval from the Research Ethics Board at the University of Bologna (protocol 0217279); the registration number at the American Economic Association's Registry for Randomized Control Trials is AEARCTR-0014172.

Email addresses: ashe@ethz.ch (Elliott Ash) sergio.galletta@gess.ethz.ch (Sergio Galletta, corresponding author) giacomo.opocher2@unibo.it (Giacomo Opocher).

1. Introduction

In a well-functioning democracy, access to high-quality information is necessary for citizens to make good political decisions. It helps them understand policies, evaluate elected officials, and contribute to political discussions. Without accurate and complete information, or when misinformation spreads, the democratic process can be significantly undermined (Allcott and Gentzkow, 2017). Yet despite its critical importance and the vast availability of information provided by media, many people still find it costly and challenging to become adequately informed. The costs are not just monetary but also involve time and effort – e.g., a Pew Research Center (2024) survey conducted in April 2024 found that 59% of U.S. adults feel worn out by the amount of election coverage. The cost of obtaining and processing information can discourage people from staying informed, leading to decisions based on incomplete or biased information and undermining the electorate’s ability to hold officials accountable and make informed choices (Matějka and Tabellini, 2020).

Given these challenges, emerging technologies such as large language models (LLMs) and AI-powered chatbots offer promising potential solutions for reducing barriers associated with accessing political information. Chatbots equipped with curated knowledge bases may significantly lower the cost of collecting information by providing immediate, concise, and relevant answers to specific user questions. They have the potential to simplify the process of understanding complex issues by presenting information in a structured, user-friendly manner that is customized to the specific voter’s needs, potentially reducing the time and effort required to become informed. Furthermore, these technologies could enhance overall knowledge acquisition, helping users to develop a clearer, more comprehensive understanding of topics and thereby improve democratic decision-making (Noy and Zhang, 2023; Brynjolfsson et al., 2025; Dell’Acqua et al., 2023).

This paper provides the first evidence of how AI-powered chatbots can help address information barriers in the democratic process. We investigate the impact of these technologies on American voters’ political knowledge by developing BallotBot, an interactive chatbot powered by state-of-the-art large language models (LLMs). To examine BallotBot’s impact on voters’ political knowledge, engagement, and perceived costs associated with acquiring political information, we conducted a three-wave survey experiment during California’s November 2024 election cycle, focusing specifically on three citizen-initiated ballot measures.

Participants were randomly assigned to interact with either BallotBot, an AI-powered chatbot designed to provide official information from the California Voter Guide, or the traditional digital version of the guide. In the first wave, we assessed participants' perceived baseline knowledge, political attitudes, and media habits before providing their assigned information source (chatbot or guide). Respondents then answered quiz questions about one of the three selected ballot propositions. We also measured the perceived cost of acquiring information using an incentive-compatible mechanism, allowing us to quantify how much participants were willing to invest in answering an additional question. In the second wave, we re-evaluated participants' political knowledge retention, analyzed their voting intentions, and evaluated the quality of their justifications provided in an open-ended question. In the third wave, we collected information on self-reported turnout and voting behavior, after election day.

First, we find that BallotBot improves voter knowledge of ballot measures. We estimate BallotBot's impact on voter knowledge by separately analyzing responses to basic questions (which can be answered using information provided at the beginning of the voter guide) and in-depth questions (which require respondents to read further into the guide). BallotBot significantly improves participants' ability to correctly answer in-depth questions, achieving an 18% improvement in accuracy compared to the control group, which relied on the traditional voter guide. Additionally, BallotBot reduces the time required to answer these in-depth questions by 10%. However, it has no effect on participants' accuracy when answering basic questions, and if anything, it increased the time needed to answer those basic questions.

Second, using BallotBot improves participants' ability to judge their own performance. While it did not significantly affect their overall confidence in basic or detailed answers, it reduced their tendency to be overconfident. Specifically, participants using BallotBot were about half as likely to be highly confident in answers that were actually incorrect, compared to those using the traditional voter guide. This finding is of particular interest in light of the evidence documenting the negative influence of overconfidence on democratic processes ([Kartal and Tyran, 2022a](#)).

Third, we find that BallotBot can reduce the perceived cost of acquiring additional information about ballot initiatives. During the experiment, participants completed a bidding task based on the BDM mechanism to measure their perceived cost of acquiring additional information ([Becker et al., 1964](#)). While there is no effect of BallotBot on the average participant bids, we find interesting heterogeneity. BallotBot significantly

lowers the perceived cost of information for individuals with lower prior knowledge, but only when they are informed about their initial performance using the bot (which was also randomized). The larger effect for low-knowledge participants is coherent with recent work showing that generative AI tends to disproportionately benefit lower-skilled individuals or those with less initial preparation (Brynjolfsson et al., 2025; Cui et al., 2024). Further, the larger effect for participants receiving performance feedback is coherent with other work showing that feedback and transparency are pivotal for acceptance of algorithms and trust in AI-based recommendations (Ahn et al., 2024).

Fourth, using BallotBot increases curiosity and engagement with ballot initiatives. In a free-form task where participants were given the opportunity to collect information on a dimension of the proposition they had previously identified as important to their voting decision, participants in the BallotBot condition spent 80% longer in collecting that information. When given the option to use BallotBot to learn about additional propositions (rather than the voter guide), respondents were 73% more likely to explore additional propositions beyond the one they were randomized into. These effects indicate that BallotBot can boost motivation to seek political information that is relevant to voting decisions.

Fifth, the effects of BallotBot are not persistent. After the first wave, participants had continued access to either BallotBot or the Voter Guide. In a second wave a week later, participants took a similar quiz without access to the guide or bot. We find no difference between the guide group and the BallotBot group in the average quiz scores. Besides the quiz, we asked participants to write a short statement explaining their voting intention. Based on an AI evaluation of the quality of the reasoning in these statements, we again find no difference between the treatment groups. In a descriptive analysis of the second wave, we assessed differences in outcomes based on the intensity of tool usage in the interim between waves. Participants who accessed BallotBot more often in the interim scored higher on the quiz, demonstrating improved factual understanding. Participants who accessed the voter guide more often in the interim produced higher-quality written motivations for their voting intentions, reflecting better articulation and reasoning. These descriptive results add texture to how AI might influence learning about policies in a referendum setting.

Sixth and finally, we looked at whether BallotBot influenced voting behavior. We find no statistically significant effect on self-reported voter turnout. There were imprecise null effects on the likelihood of voting "yes" on the assigned propositions.

These findings add to the literature on information and democracy, particularly the political consequences of media and social media exposure. Media technologies – including AI assistants – are central to knowledge production, beliefs, and behavior. The most related evidence on this dimension includes [Gentzkow \(2006\)](#), who shows that the introduction of television in the United States led to declines in political knowledge and voter turnout, [Snyder Jr and Strömberg \(2010\)](#) who find that reduced press coverage of U.S. House representatives leads to lower voter knowledge and weaker congressional oversight, and [Schroeder and Stone \(2015\)](#) reporting that Fox News coverage significantly influences public awareness of policy issues, with effects varying based on issue salience and partisan alignment. [Allcott and Gentzkow \(2017\)](#) report on how social media influenced information consumption and knowledge in the context of the 2016 U.S. election. [Zhuravskaya et al. \(2020\)](#) and [Levy \(2021\)](#) show how social media fosters polarization and mobilization, while [González-Bailón et al. \(2023\)](#) reveal substantial ideological segregation in political news exposure on Facebook. Furthermore, [Kartal and Tyran \(2022a\)](#) highlight the special relevance and risks of voter overconfidence, which can amplify the negative political effects of misinformation. We extend this body of evidence by exploring the role of AI chatbots as a new media technology in shaping political knowledge, beliefs, and behaviors, particularly considering how they may influence the costs associated with obtaining and processing political information ([Matějka and Tabellini, 2020](#); [Nyhan, 2020](#)).

Second, the paper adds to the fast-growing literature exploring how LLMs will affect the economy and society. For example, [Noy and Zhang \(2023\)](#) provide experimental evidence with freelance writers that generative AI enhances productivity. Similarly, [Brynjolfsson et al. \(2025\)](#) find that generative AI tools boost productivity in a customer support role, especially for less-experienced workers, while also enhancing customer satisfaction and employee retention. [Dell’Acqua et al. \(2023\)](#) conducted a field experiment in a consulting firm, testing the effects of integrating LLM into typical knowledge-intensive tasks showing that consultants assisted by LLM completed tasks faster with an improved output quality. Moreover, [Cui et al. \(2024\)](#) analyze three field experiments conducted with software developers and find that AI-based coding assistants increased developer productivity, with the largest benefits observed among less-experienced developers. In [Costello et al. \(2024\)](#), personalized dialogues with AI chatbots aiming to debunk conspiracy theories reduced belief in the conspiracy theories. We expand this work into the domain of supporting voters in their civic duties, specifically focusing on how genera-

tive AI can influence voter information acquisition, political engagement, and electoral decision-making.

Finally, our findings relate to the more specific literature on direct democracy and the important role of voter informedness. For example, [Maskin and Tirole \(2004\)](#) argues that representative democracy can outperform direct democracy when politicians use private information to select the best, rather than merely popular, policies. In contrast, [Matsusaka \(2005\)](#) emphasizes that direct democracy allows voter participation in lawmaking, which can enhance governmental performance by incorporating public preferences and aggregating dispersed information. Supporting this view, [Olken \(2010\)](#) use a field experiment in Indonesia to demonstrate that direct democracy mechanisms increase citizen satisfaction, knowledge, and the perceived legitimacy of local public goods decisions. Furthermore, evidence from [Galletta \(2021\)](#) suggests that participation in direct democratic events can make individuals more willing to support expansions of the welfare state. Our results suggest that AI technologies could facilitate direct democracy by helping voters become informed at a lower cost.

The rest of the paper proceeds as follows. Section 2 provides institutional details. Section 3 describes the experimental design, while Section 4 reports the empirical strategy and the results. Section 5 concludes.

2. Background

2.1. Institutional Setting

California offers a distinctive context for analyzing how information technologies influence voter behavior under a direct democracy framework. As one of the nation’s most populous states, it has a long-standing tradition of direct democratic institutions, including ballot propositions, initiatives, and referenda. These mechanisms allow citizens to propose, enact, or repeal legislation through popular vote.

The state constitution empowers voters to propose legislation, approve or reject laws passed by the legislature, and recall elected officials. Ballot propositions, in particular, have become a central feature of California elections, addressing a broad spectrum of issues including taxation, criminal justice, environmental regulations, and social policies.

2.2. California Ballot Measures, November 2024

California's November 2024 election features several significant ballot measures addressing both economic and social issues. Among the most notable are propositions to increase the minimum wage, expand rent control, and raise charges and penalties for specific criminal offenses (see Appendix Table A1).

Proposition 32 - *Raises Minimum Wage* proposes increasing California's hourly wage to \$18 by 2026, with future adjustments tied to inflation. Supporters contend this measure is necessary to address the state's high cost of living and bolster low-income workers, while opponents caution that higher labor costs may burden businesses and limit employment opportunities.

Proposition 33 - *Expands Local Governments' Authority to Enact Rent Control* would repeal the Costa-Hawkins Rental Housing Act, permitting cities and counties to impose rent control on single-family homes and newly built housing units. Proponents view this as a crucial step in mitigating California's housing crisis and safeguarding tenants. Critics warn, however, that it may hinder housing development and investment.

Proposition 36 - *Modifies Felony Charges for Certain Drug and Theft-Related Crimes* increases charges and penalties for offenses linked to organized retail theft and repeat crimes. Proponents argue that stronger deterrents are essential to counter rising crime rates and protect public safety. Opponents express concern that escalating penalties may contribute to higher incarceration rates without addressing underlying causes of criminal behavior.

2.3. Voter Guide

The California Voter Guide serves as an official, comprehensive resource intended to inform voters about statewide ballot measures, candidates, and election procedures. Compiled and distributed by the California Secretary of State, it includes impartial analyses of ballot propositions, summaries prepared by the Attorney General, arguments for and against the measures submitted by proponents and opponents, and rebuttals to these arguments.

In our survey sample, 89% of participants reported receiving the guide. Of these, 54% usually or always read it, while 29% indicated they only read a few pages, and 44% found the guide clear.

The 2024 Voter Guide is 144 pages long, allocating 4–5 pages per proposition. It provides background information, fiscal impact analyses, and concise explanations of pros

and cons, such as those for proposed changes to the minimum wage. A screenshot of an excerpt of the 2024 guide is shown in Figure 3 Panel A. The guide’s primary objective is to promote informed decision-making by presenting objective facts and contrasting viewpoints, thereby enhancing transparency and accessibility in California’s democratic process.

2.4. BallotBot

BallotBot is an AI-powered chatbot developed specifically for this study, offering an interactive, user-friendly way to access information from California’s official Voter Guide on statewide ballot propositions. While it shares the core functionality of standard LLM-based chatbots (e.g., ChatGPT or Claude), BallotBot differs in its use of Retrieval-Augmented Generation (RAG). Rather than relying on general knowledge alone, it supplements the system with proposition-specific information, ensuring answers are derived directly and accurately from the Voter Guide.

Specifically, BallotBot is a hosted instance of OpenAI’s `gpt-4o-mini-2024-07-18` model, served via Microsoft Azure. We created three distinct versions of BallotBot, each focusing on a different proposition. A screenshot of the BallotBot interface is shown in Figure 3 Panel B. A sample conversation with BallotBot is shown in [Appendix A.2](#).

3. Experiment Design

We conducted a three-wave survey experiment to assess the impact of AI-powered chatbots on voters’ political knowledge, engagement, and behavior in the context of California’s 2024 ballot measures. The experiment, pre-registered in the AEA RCT Registry, evaluates the effects of BallotBot, an AI-powered chatbot, relative to the traditional digital version of the official voter guide.

This section describes our experimental design. Figure 1 shows the overall timeline of the experiment, including pre-registration, the survey waves, and the election.

Zooming in on the survey waves, Figure 2 summarizes the study structure, including details of the first and second survey waves, the information sources provided to participants based on their treatment assignments, and the primary outcomes of interest.

3.1. First Wave

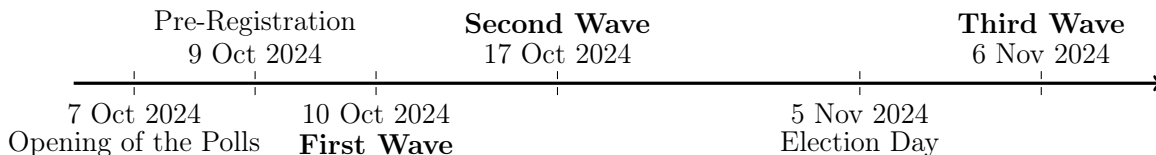
On October 11, 2024, we enrolled 1,463 survey participants via the survey company Prolific. We first collected participants’ baseline demographics, political attitudes,

Table 1: Summary Statistics and Treatment Balance Tests

Variable	All Sample	Control Mean	Treatment Mean	Norm. Difference
<i>A. Demographics</i>				
Gender	0.470 (0.500)	0.479 (0.500)	0.467 (0.499)	0.023 [0.658]
Age	38.360 (13.320)	38.401 (13.232)	38.320 (13.412)	0.006 [0.907]
High Education	0.800 (0.400)	0.784 (0.412)	0.817 (0.387)	-0.083 [0.114]
<i>B. Socio Economic Status</i>				
Employment	0.740 (0.440)	0.736 (0.441)	0.749 (0.434)	-0.029 [0.580]
Democrat	0.540 (0.500)	0.525 (0.500)	0.560 (0.497)	-0.070 [0.182]
Republican	0.170 (0.370)	0.176 (0.381)	0.158 (0.365)	0.048 [0.357]
<i>C. Attitudes Toward A.I.</i>				
Never Used A.I.	0.230 (0.420)	0.239 (0.427)	0.225 (0.418)	0.033 [0.527]
How Useful is A.I.	3.680 (1.480)	3.643 (1.496)	3.725 (1.467)	-0.056 [0.287]
<i>D. Baseline Knowledge</i>				
Know about the Proposition	0.570 (0.490)	0.575 (0.495)	0.570 (0.495)	0.010 [0.850]
Self-Perceived Kn.	0.420 (0.490)	0.421 (0.494)	0.410 (0.492)	0.023 [0.656]
Turnout Pre	0.760 (0.430)	0.758 (0.429)	0.753 (0.432)	0.012 [0.820]
<i>E. Sources of Information</i>				
Social Media	0.640 (0.480)	0.658 (0.475)	0.624 (0.485)	0.070 [0.179]
Traditional Media	0.390 (0.490)	0.379 (0.485)	0.410 (0.492)	-0.063 [0.227]
Official Documents	0.280 (0.450)	0.263 (0.440)	0.290 (0.454)	-0.060 [0.249]
Num. Obs.	1464	731	732	
F-Test of Joint Sig. (P-Val)				0.707
Num. Obs.				1463

Notes: We test for covariate balance across control and treatment groups. Means, standard deviations (in parentheses), and normalized differences (in square brackets) are reported. Panels group variables into demographics, socio-economic status, attitudes toward A.I., baseline knowledge, and sources of information. The F-test reports the joint significance of all covariates.

Figure 1: Timeline of the Experiment



Note: This figure presents the overall timeline of the experiment.

self-reported political knowledge, attitudes toward A.I., and media consumption habits. Summary statistics on the participants are shown in Table 1.

Participants were then randomly assigned to one of the three statewide ballot propositions: Proposition 32, Proposition 33, or Proposition 36. We collected information on their baseline knowledge of the proposition to which they were assigned. Further, we asked participants to pose between one and three questions whose answers they believed would help them reach a confident voting decision.

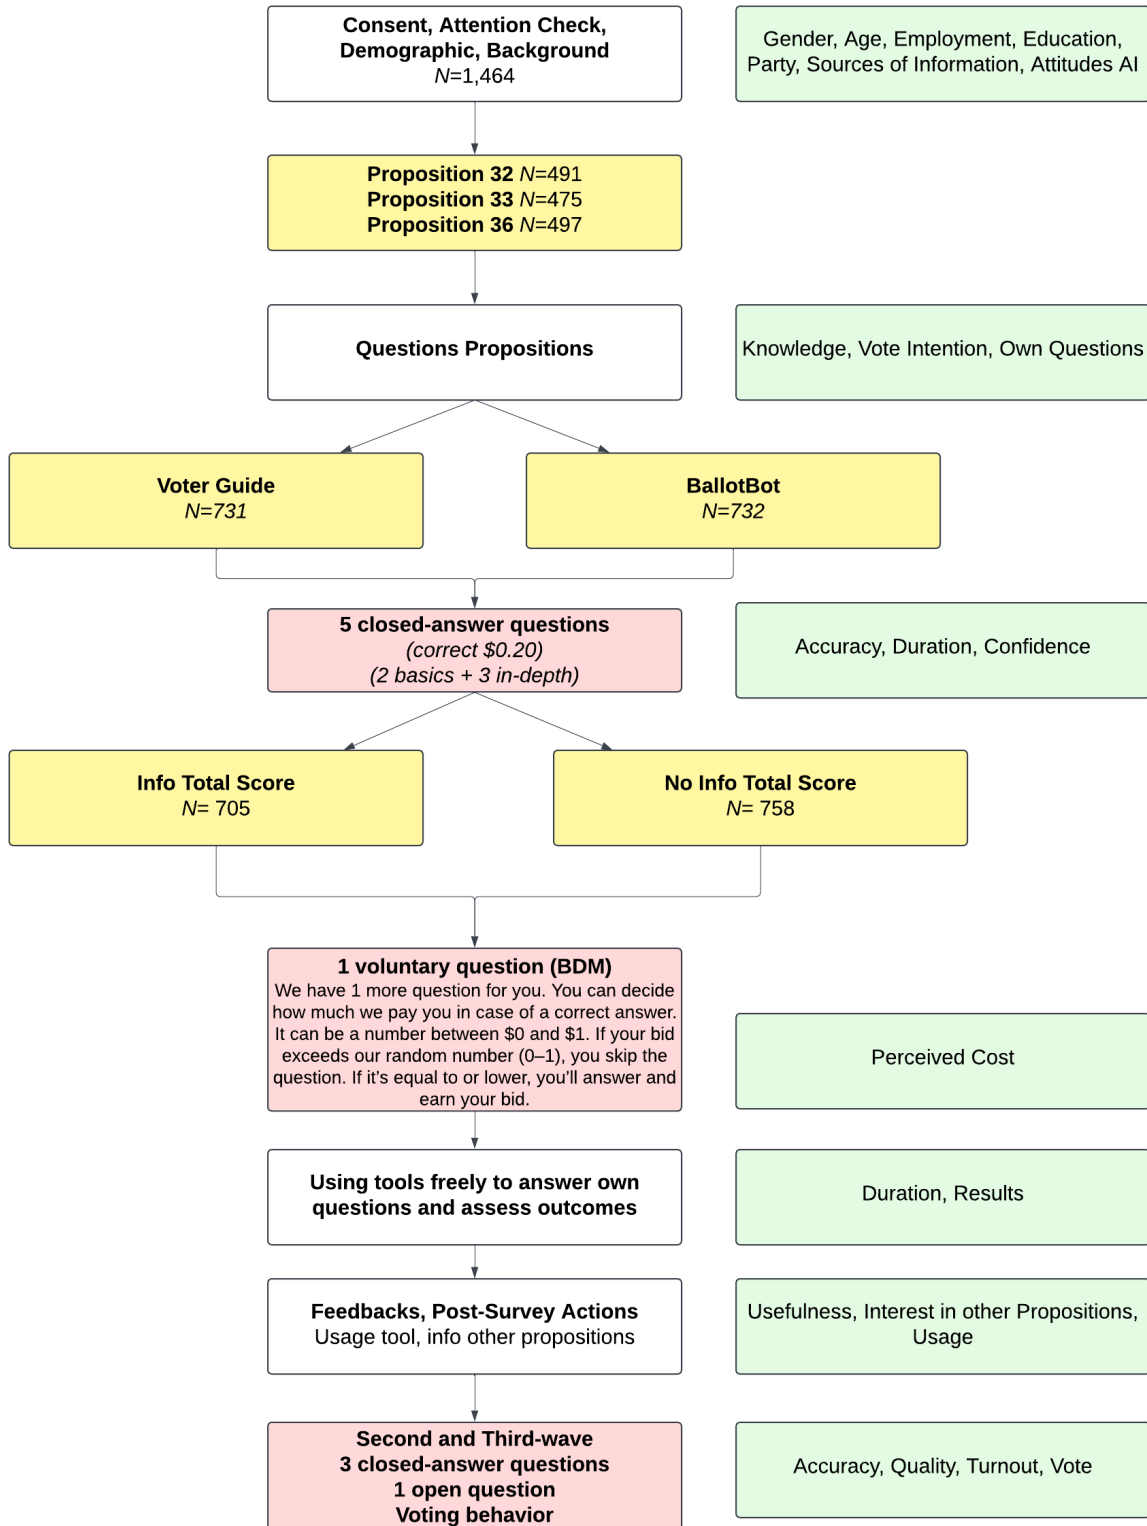
Within each proposition group, participants were further randomly assigned to one of two treatment groups:

- **Voter Guide Group (Control):** Participants in the control group were provided with the official California voter guide related to the proposition converted to HTML format. The guide included detailed information about the ballot propositions, and it is the same as the physical version mailed to voters.
- **BallotBot Group (Treatment):** Participants in the treatment group had access to BallotBot, an AI-powered chatbot designed to answer questions about the assigned ballot proposition, drawing exclusively from the information contained in the voter guide.

The user interface experiences for both BallotBot and the voter guide are illustrated in Figure 3 (see also Appendix Figures A1 and A2).

While participants used the voter guide and chatbot, we could monitor various dimensions of user interactions (see Appendix Table A4). Overall, participants interacted with BallotBot an average of about nine times, submitting queries of about 12 words each. All respondents assigned to BallotBot made at least one interaction. BallotBot responded with an average of 96 words in roughly 3 seconds per query. In Figure A4 we report the full distribution of interactions. In Figure A5 we show that none of the observable dimensions is a strong predictor of the intensity of usage during the survey.

Figure 2: Overview of Experiment Design



Note: This figure outlines the complete structure of the experimental design. Yellow rectangles indicate points of randomization, while the red rectangle marks the stage where the primary outcomes of interest are collected. Green rectangles highlight the variables measured at each stage of the experiment.

Figure 3: Information Sources: Voter Guide & BallotBot

(a) Voter Guide

PROPOSITION **33** EXPANDS LOCAL GOVERNMENTS' AUTHORITY TO ENACT RENT CONTROL ON RESIDENTIAL PROPERTY. INITIATIVE STATUTE.

OFFICIAL TITLE AND SUMMARY PREPARED BY THE ATTORNEY GENERAL

The text of this measure can be found on page 100 and the Secretary of State's website at voterguide.sos.ca.gov.

- Current state law (the Costa-Hawkins Rental Housing Act of 1995) generally prevents cities and counties from limiting the initial rental rate that landlords may charge to new tenants in all types of housing, and from limiting rent increases for existing tenants in (1) residential limiting the ability of cities and counties to maintain, enact, or expand residential rent-control ordinances.

SUMMARY OF LEGISLATIVE ANALYST'S ESTIMATE OF NET STATE AND LOCAL GOVERNMENT FISCAL IMPACT:

(b) BallotBot

BallotBot

Hi, I am your BallotBot!

I'm trained on the data from Proposition 32. Ask me anything about it!

Type a new question...

Notes: These screenshots show the information sources that are then embedded in the survey interface. Panel A: Voter Guide (control group); Panel B: BallotBot (treatment group).

Table 1 reports the balance tests on treatment assignment. There are no significant differences in any of the pre-treatment observables. Further, the joint F-test on these variables is not significant either.

In the first post-treatment task, participants answered five multiple-choice questions about their designated ballot proposition.¹ The questions were randomly selected from a pool and stratified by difficulty level. Two questions were *basic*, with answers located on the first page or two of the guide. Three questions were *in-depth*, with answers found later on in the document in the Analysis section or in the arguments for/against the proposition.²

Throughout this task, the assigned information source (guide or chatbot) was embedded within the survey page window for easy access, as shown in Figure A2 and A1. Participants were encouraged to consult the respective information source to help them with the task. Correct responses earned \$0.20 each.

For each question, after answering participants also indicated their confidence level in their responses on a discrete scale with four levels (1-highly uncertain, 2-quite uncertain, 3-quite certain, 4-highly certain). In addition to recording their answers and confidence levels, we also measured the time taken to answer each question.

Next, half of the sample was randomly assigned to view their score from the previous quiz before proceeding to the next section. Participants were then asked what compensation they would require to answer an additional question. This reservation wage was elicited through an incentive-compatible Becker-DeGroot-Marschak (BDM) bidding task (see Figure A3).³ We interpret the reported bid as a measure of the participant’s perceived cost of answering an additional question.

Next, we reminded participants of the custom questions they had proposed at the beginning of the survey as being relevant to their decision on the ballot initiative. Par-

¹Each question has four options for answers, plus two additional “I don’t know” options: “*I could not find the relevant information*” and “*I could not figure out the answer*”. The questions were designed by the authors based on the content of the official voter guide, independently of the chatbot’s development. Answers do not contain the explicit text of the guide, so that they could not be answered easily through text search. Questions were written to require consultation of the information source; for example, the fraction of Californians living in communities that already have local rent control.

²The question sets, correct answers, and selection details are provided in Appendix Tables A11, A12, and A13.

³Participants were informed that a random number between 0 and 1 would be drawn, and if their bid (in dollar) exceeded this random number, they would proceed to the next section without answering the question. Otherwise, they would answer the additional question and receive the amount they bid as compensation.

ticipants were encouraged to look for the answers using the embedded source, with no incentives provided. We record the time they spend in that window as a measure of engagement, whether they have found the answer, and how much they trust it.

At the end of the survey, we thank participants for their involvement and display the completion link to receive their show-up fee. On the same page, we list three additional buttons: one linking to the information source they used during the survey, and two more linking to the same source configured for the two other unseen ballot initiatives. We record clicking any of these buttons as an additional measure of engagement with the information sources.

Finally, we informed the participants that they could continue using the same source for the next few days as they like.

3.2. Second Wave

Participants were informed during the first wave that there would be a second wave of the survey. Three days after the first wave, we sent participants a message reminding them that they could use the information source (guide or chatbot) they had accessed during the first wave, including a tracked link for easy access. Three days later (six days after the first wave), we invited all participants to complete the follow-up survey. Of the 1,463 participants who completed the first survey, 1169 participated in the follow-up survey, while 294 subjects did not, resulting in an attrition rate of 20.1%. We document this attrition in Table [A2](#).

We asked participants to share their voting intentions and provide a brief written explanation. To assess the quality of these explanations, we prompted a language model to score the text on the quality of reasoning. Next, participants responded to three quiz questions drawn from the same pool used during the first wave, but this time without access to the information source. We randomly select one basic and two in-depth questions. For each question, we recorded their answers, confidence levels, and response times.

3.3. Third Wave

All participants in the first wave were re-contacted to participate in the third wave data collection. Of the 1463 participants who completed the first survey, 913 completed the third wave survey, resulting in an attrition rate of 37.6%. We document attrition in Table [A3](#).

In the third wave, we ask participants to report turnout on all propositions considered, at the general elections, and, in case they voted, what direct they voted.

3.4. Outcomes of Interest

The outcomes of interest are grouped into four broad categories: 1) knowledge-related outcomes, 2) information mapping costs, 3) engagement with political content, and 4) reported voting behavior.

Knowledge-Related Outcomes. In line with established practice (e.g., [Lupia, 2016](#)), we measure political knowledge as the share of correct answers on the multiple-choice questions administered in the first and second waves. This approach captures immediate learning in the first wave and knowledge retention in the second.

Beyond these standard measures of knowledge, we examine participants’ open-ended responses in which they justify their voting positions. These explanations are evaluated using GPT-4 and a pre-registered prompt that rates how well each response justifies the decision, the clarity of reasoning, and the use of relevant evidence or information about the proposition.⁴

Finally, we measure respondents’ confidence using a 4-point scale ranging from 1 (*Highly Uncertain*) to 4 (*Highly Certain*). We define *overconfidence* as demonstrating high certainty in an incorrect answer and *underconfidence* as demonstrating low certainty in a correct answer.

Information Mapping Costs. We examine both the objective and perceived costs of converting information into knowledge. We measure the objective cost by recording the time (in seconds) needed to answer each question. We approximate the perceived cost of answering an additional question using participants’ bids in the BDM task: higher bids indicate a greater perceived effort or cost of seeking out the correct answer.

Engagement with Political Content. We measure engagement during the first wave by examining participants’ willingness to interact with their assigned resource. Specifi-

⁴We prompted the LLM with the following text: “Evaluate the strength and persuasiveness of this response explaining why someone chose to vote ‘yes’ or ‘no,’ decided not to vote, or remained undecided on a specific ballot proposition. Rank the argument on a scale from 1 to 10, and provide a brief explanation for the score, focusing on how well the response justifies the decision, the clarity of reasoning, and the use of relevant evidence or information about the proposition.”. We validated these evaluations by having two distinct human annotators perform the same task for around 10% of the full (90 responses, 30 per proposition). The (positive) correlation between the LLM and annotator 1 was 71%, and with annotator 2 it was 62%. The correlation between annotator 1 and annotator 2 was 77%.

cally, we track the time spent using the resource while participants explore the questions they consider necessary for making a confident voting decision. We also evaluate spillover effects by observing the likelihood of participants engaging with other propositions. This is captured as the probability of clicking a link to the same resource (BallotBot or Voter Guide) for a different proposition—e.g., exploring Propositions 32 and 36 after being assigned Proposition 33—at the end of the survey page.

Additionally, we assess source usage between the first and second survey waves as an extra measure of engagement, recording whether participants accessed the source at least once during this period.

Reported Voting Behavior. In the third wave, we measure participants’ self-reported voting behavior, focusing on two components: turnout (i.e., whether they participated in the election) and their yes/no decision on the assigned proposition.

4. Analysis and Results

This section outlines our specification for data analysis and presents the results.

4.1. Empirical Specification

The empirical analysis in this section relies mostly on a simple comparison of means, which we estimate using the following regression framework. Let Y_i denote an outcome variable for participant i among those discussed in the previous section. Let $D_i \in \{0, 1\}$ be an indicator for being assigned to use BallotBot, and let ε_i be an idiosyncratic error term. We estimate the following equation:

$$Y_i = \alpha + \beta D_i + \mathbf{X}_i' \gamma + \varepsilon_i, \tag{1}$$

where \mathbf{X}_i is a vector of control variables (including demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning) and γ is the corresponding vector of coefficients. The coefficient of interest, β , measures the average treatment effect of being assigned to use BallotBot compared to the standard Voter Guide.

4.2. Results

Effect on answering questions about the ballot initiatives. First, we look at the effect of BallotBot access on answering questions correctly, the response time in

Table 2: Effect of BallotBot on Accuracy, Response Time, and Confidence

Qst. Difficulty	Share Correct		Answering Time		Confidence	
	Basic	In Depth	Basic	In Depth	Basic	In Depth
	(1)	(2)	(3)	(4)	(5)	(6)
BallotBot	-0.009 (0.015)	0.137*** (0.013)	5.532** (1.981)	-6.868** (2.152)	0.006 (0.074)	0.051 (0.034)
Contr. Mean	0.821	0.757	52.577	67.619	3.336	3.232
St. Dev.	0.284	0.280	38.822	47.028	1.438	0.660
R ²	0.030	0.121	0.052	0.072	0.049	0.043
Num. Obs.	1463	1463	1463	1463	1463	1463

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 if participant is in the BallotBot group. The set of outcome variables are: share of correct answers, answering time in seconds, and confidence in the answer on a scale from 1 to 4. Effects are reported separately for basic and in-depth questions, as described in the text. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors in parenthesis: *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

answering, and the self-reported confidence in those answers. Following our pre-analysis plan, we report results separately for answers to basic questions (those located at the beginning of the voter guide) and in-depth questions (which require respondents to read further into the guide).

Table 2 reports the estimates for β from Eq. (1) with the question-answering outcomes on the left-hand side. Columns (1) and (2) show the effect of receiving BallotBot on the share of correct answers to basic and in-depth questions, respectively. Columns (3) and (4) show the effects on the average time needed to answer each question, while Columns (5) and (6) show the effect on self-reported confidence.

Our results show that BallotBot improves accuracy, notably only for in-depth (rather than basic) questions.⁵ In particular, BallotBot users answer 13.7pp more in-depth questions correctly, an 18 percent increase compared to the control group. To put it differently, BallotBot reduces the error rate from 24.3% to 10.6%, effectively cutting the error rate by approximately 50% compared to the control group. Next, while BallotBot users are 5.5 seconds slower to answer basic questions, they are 6.9 seconds faster to

⁵We note that overall accuracy is relatively high already in the control group (a weighted average of 0.78, with 0.82 for basic and 0.75 for in depth questions), reflecting that all participants had direct access to relevant information when answering questions. For comparison, the average accuracy in answering questions during the second wave – when participants did not have direct access to the voter guide – was 12 percentage points lower.

Table 3: Effect of BallotBot on Over/Under-Confidence in Answers

Qst. Difficulty	Over/Under-confidence		Overconfidence		Underconfidence	
	Basic	In Depth	Basic	In Depth	Basic	In Depth
	(1)	(2)	(3)	(4)	(5)	(6)
BallotBot	0.010 (0.017)	-0.121*** (0.020)	0.016 (0.013)	-0.136*** (0.016)	-0.006 (0.012)	0.022 (0.017)
Contr. Mean	0.213	0.334	0.123	0.248	0.090	0.135
St. Dev.	0.308	0.410	0.238	0.376	0.228	0.303
R ²	0.024	0.066	0.034	0.095	0.028	0.022
Num. Obs.	1462	1462	1462	1462	1462	1462

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 for participants in the BallotBot (treatment) group. Outcomes are indicators for being over- or underconfident, being overconfident, and being underconfident. Effects are reported separately for basic and in-depth questions, as described in the text. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors in parenthesis. *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

answer in-depth questions (a 10 percent decrease in time spent per question). Finally, there is no overall effect on self-reported confidence in the answers.⁶

Effect on over- and under-confidence in one’s answers. While there is no overall effect on confidence, we also wondered whether there is an effect on misplaced confidence. That is, does BallotBot differentially affect confidence in correct vs incorrect answers?

Table 3 analyzes this dimension of the data. Columns (1) and (2) report the average effects on the probability of being overconfident (i.e., being sure of a wrong answer) or underconfident (i.e., being unsure of a correct answer), again distinguishing between basic and in-depth questions. Columns (3) and (4) focus on the probability of being overconfident, while Columns (5) and (6) focus on the probability of being underconfident.

We find that using BallotBot reduces the likelihood of inaccurate perceptions of being correct for in-depth questions, by 12.1pp (35%). This effect is driven entirely by a 13.1pp (50%) decrease in the probability of overconfidence, with no significant change in underconfidence. These results are important in light of the discussion by [Kartal and Tyran](#)

⁶In additional analyses (not reported here), we re-estimated our regressions including fixed effects for individual questions, with data aggregated at the question level. Results remain largely unchanged in both magnitude and statistical significance, confirming robustness even after explicitly controlling for specific question characteristics. This outcome was anticipated, given that the randomization should have already balanced out question-specific effects.

(2022a) about the special relevance and risks of voter overconfidence, which can amplify the negative political effects of misinformation among voters.

Effect on perceived cost of answering an additional question. The next question is how BallotBot affects the perceived cost of answering an additional question. As noted above, this perceived cost is measured as the reservation wage (in dollars) required to answer another question using the same information source, elicited through an incentive-compatible BDM mechanism. Following our pre-analysis plan, we compare the perceived cost both by randomized information source (guide or bot) but also by whether the participants got feedback on their share of correct answers among the first five multiple-choice questions. Receiving feedback was cross-randomized.

Table 4 presents the estimates. *Panel A* has results for the group who were informed about their quiz score, while *Panel B* includes results for those who were not informed. In both panels, Column (1) reports the treatment effect for the whole group, Columns (2) and (3) break out the sample by high or low baseline knowledge, and Columns (4) and (5) condition on whether respondents have a college degree.

First, in both groups, we find no effect on the overall average bid. However, in Panel A where respondents got performance feedback, there is significant heterogeneity by baseline knowledge and educational attainment. BallotBot significantly reduces the perceived cost for those with low baseline knowledge ($\hat{\beta}=-\$0.06$, 14%) and/or for high school graduates without college ($\hat{\beta}=-\$0.12$, 26%), with no effect for those with high baseline knowledge and/or with higher education.⁷ Similar to findings on AI applications in other contexts, AI can be seen as a tool to reduce performance inequality in understanding ballot initiatives by lowering the perceived cost of acquiring information for individuals at the lower end of the ability distribution (Noy and Zhang, 2023; Brynjolfsson et al., 2025).

Among respondents not informed about their score (*Panel B*), we detect no significant differences, even for less knowledgeable or less educated groups. BallotBot can reduce the perceived cost of transforming information into knowledge among disadvan-

⁷A potential mechanism for the shift among participants receiving feedback is that BallotBot influences participants' perceived likelihood of answering the BDM question correctly, via the improvement in performance on the earlier questions. We check for this mechanism by running the same regression with fixed effects for the number of previous correct answers. As shown in Appendix Table A10, the results are the same after conditioning on these mediators. Hence, the shift in perceived cost is a net effect beyond changes in beliefs vs increased performance in the previous rounds.

Table 4: Effect of BallotBot on Perceived Cost of Question Answering

	Bid (\$)	Bid (by Prior Kn.)		Bid (by Education)	
	All Sample	High Kn.	Low Kn.	College	No College
	(1)	(2)	(3)	(4)	(5)
<i>A. With Feedback on Previous Performance</i>					
BallotBot	-0.027 (0.018)	0.011 (0.025)	-0.063*** (0.025)	-0.013 (0.020)	-0.119** (0.044)
Contr. Mean	0.448	0.425	0.474	0.448	0.448
St. Dev.	0.232	0.221	0.242	0.236	0.213
R ²	0.048	0.085	0.196	0.048	0.338
Num. Obs.	705	399	306	582	123
<i>B. Without Feedback on Previous Performance</i>					
BallotBot	0.020 (0.016)	0.037 (0.022)	-0.001 (0.028)	0.015 (0.018)	0.034 (0.044)
Contr. Mean	0.449	0.447	0.451	0.445	0.461
St. Dev.	0.215	0.213	0.219	0.207	0.240
R ²	0.052	0.090	0.086	0.055	0.205
Num. Obs.	757	456	301	589	168

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 for participants in the BallotBot (treatment) group. The outcome variable is the bid (in dollars) selected at the BDM task, where higher values indicate a higher perceived effort cost of answering an additional question. Panel A includes respondents that received information on their previous performance, while Panel B includes respondents who did not. Column 1 includes the whole sample; Columns 2 and 3 report heterogeneity by self-reported pre-treatment familiarity with the proposition; Columns 4 and 5 report heterogeneity by education level. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors in parentheses. *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

Table 5: Effect of BallotBot on Engagement with Information

	Custom Questions			
	Time Spent (s)	Found Answer	Clicked Other Prop.	Interim Access
	(1)	(2)	(3)	(4)
BallotBot	19.797*** (2.017)	0.001 (0.023)	0.048*** (0.014)	0.013 (0.017)
Contr. Mean	24.69	.734	.068	.131
St. Dev.	29.433	.441	.253	.337
R ²	0.085	0.008	0.034	0.023
Num. Obs.	1463	1463	1463	1463

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 for participants in the BallotBot (treatment) group. The outcomes considered are the time spent (in seconds) collecting information on one’s own custom question, an indicator for reporting that they found the answer to the question, an indicator for voluntarily collecting information on another proposition at the end of the survey, and an indicator for voluntarily accessing the resource during the week between the two waves. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors in parentheses. *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

tagged individuals, but only when they receive feedback on their quiz performance. One interpretation of this heterogeneity is that without performance feedback, respondents misperceive the quality and effectiveness of BallotBot. In this specific context, providing feedback on actual quiz performance addresses this misperception and helps users evaluate the usefulness of the bot (Ahn et al., 2024).

Effect on voluntary information seeking. Besides these incentivized tasks, we also gave participants several opportunities to voluntarily engage with the information source. Our question is whether BallotBot encourages people to explore ballot initiatives more. Table 7 has the results.

First, we look at the time spent on exploring the subject’s own custom question (Column 1), and whether the question was answered (self-reported, Column 2). We find that using BallotBot increases the time spent by users in looking for answers to their question, by 20 seconds (80%). There is no effect on whether they eventually found an answer, however.

Second, we look at whether, at the end of the survey, the subject clicked on links to information about the other two propositions (with the same type of source, BallotBot or Voter Guide; Column 3). For example, exploring Propositions 32 and 36 after being randomized into Proposition 33. BallotBot has a significant effect, raising the probability

of checking other propositions by 0.05 (73%).

Finally, we look at “Interim Access” a dummy equaling one if the subject voluntarily re-accessed the information source after the first wave but before the second wave. This was tracked by the subject’s individualized link to each information source. As shown in Table 7 Column 4, there is no difference between BallotBot and Voter Guide groups in this outcome.

Effect on persistent knowledge. In Table 6, we report the effect of BallotBot on outcomes in the second wave, about a week after the initial treatment.

First, we look at the share of correct answers in the follow-up quiz (Columns 1–3). In Column 1, we show the treatment effect for the whole sample. We find no difference in the average score between these two groups. This result suggests that the positive effects observed in the first wave do not persist over time. These estimates pool the basic and in-depth questions, but we see a null as well when looking at basic and in-depth separately (see Appendix Table A8).

Columns 2 and 3 show the effect of Interim Access, defined as accessing their resource at least once during the week separating the first and second wave. Accessing the chatbot or voter guide is an endogenous choice and the associated OLS estimates are causal only under relatively strong conditional independence assumptions.⁸ Bearing that in mind, we look at heterogeneity by whether the participant used the information source in between survey waves, and we find that those who used BallotBot in the interim answered more questions correctly in the second wave. Appendix Table A9 provides a number of robustness checks and shows that this coefficient is not that sensitive to various sets of controls, providing some support for the conditional independence assumption and a causal interpretation for this effect.

Next, Columns 4 through 6 show the effects on the quality of reasoning about vote decisions in the second wave. Again, we see no overall effect on the quality of the justification given for voting. However, there is interesting heterogeneity based on whether the subject used their respective information source in between the survey waves (Columns 5 and 6). Those who used the voter guide gave better justifications for their decisions,

⁸We report balance tests of interim access in Appendix Table A6 (for voter guide subjects) and Appendix Table A7 (for BallotBot subjects). While the joint F-test for predicting interim access with observables for voter-guide users is not significant, suggesting limited selection, the F-test for BallotBot users is significant, suggesting endogenous selection into access.

Table 6: Effects on Second-Wave Outcomes

<i>A. Effects on the share of correct answers and reasoning quality.</i>						
	Share Correct			Reasoning Quality		
	All	BallotBot	Guide	All	BallotBot	Guide
	(1)	(2)	(3)	(4)	(5)	(6)
BallotBot	-0.003 (0.017)			-0.170 (0.107)		
Interim Access		0.092*** (0.031)	0.021 (0.037)		0.365 (0.218)	0.832*** (0.216)
Contr. Mean	.662	.649	.658	4.96	4.747	4.842
St. Dev.	.302	.293	.302	1.627	1.605	1.592
R^2	0.013	0.036	0.018	0.038	0.055	0.079
Num. Obs.	1158	586	572	903	454	449

<i>B. Effects on time to answer and over-/underconfidence</i>						
	Time to Answer			Over-/Underconfidence		
	All	BallotBot	Guide	All	BallotBot	Guide
BallotBot	4.864 (3.627)			-0.007 (0.017)		
Interim Access		1.124 (7.222)	10.907 (7.795)		-0.016 (0.037)	-0.053 (0.036)
Contr. Mean	78.829	82.755	77.03	.413	.406	.421
St. Dev.	62.992	62.689	62.142	.298	.293	.298
R^2	0.013	0.015	0.033	0.009	0.018	0.022
Num. Obs.	1158	586	572	1156	585	571

Notes: This table presents the estimated effects on follow-up outcomes. In columns 1 and 4, the treatment variable, BallotBot, is an indicator equaling 1 for participants in the BallotBot (treatment) group. In Columns 2 and 5, the treatment is an indicator equaling one for interim access of BallotBot between survey waves, and the sample is limited to the BallotBot treatment group. In Columns 3 and 6, the treatment is interim access to the voter guide and limited to the voter-guide control group sample. In Panel A Columns 1 through 3, the outcome is share of correct answers. In Panel A Columns 4 through 6, the outcome is LLM-judged reasoning quality on a ten-point scale, as described in the text. In Panel B we use the time spent answering (Columns 1-3) and an indicator for over-/underconfidence (Columns 4-6). All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors are provided in parentheses. Statistical significance is indicated as follows: *** p-value < 0.01, ** p-value < 0.05.

Table 7: Effect of BallotBot on Voting Behavior

	Turnout Proposition	Turnout President	Voted Yes 32	Voted Yes 33	Voted Yes 36
	(1)	(2)	(3)	(4)	(5)
BallotBot	-0.007 (0.023)	-0.023 (0.018)	0.010 (0.056)	0.086 (0.068)	0.090 (0.063)
Contr. Mean	.85	.926	.656	.533	.527
St. Dev.	.357	.262	.477	.501	.502
R ²	0.071	0.087	0.133	0.052	0.084
Num. Obs.	901	901	262	218	240

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 for participants in the BallotBot (treatment) group. The outcomes are indicators for voting at all in the assigned proposition (Column 1), an indicator for voting in the presidential election (Column 2), and indicators for voting “yes” for the assigned proposition (columns 3–5). All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors in parentheses. *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

with no difference for BallotBot users. Again, this could be due to selection, although Appendix Table A9 shows that it is robust to various specifications adjusting for observables. Still, it points to an intriguing possibility that AI chatbots can help voters understand policies enough to answer questions about them, but not deeply enough to provide well-justified arguments for subsequent decisions.

In Panel B, we show the corresponding effects on time to answer (Columns 1-3) and over/underconfidence (Columns 4-6). BallotBot had an effect on these in the first wave. But we find no effect on these outcomes in the second wave, either in terms of BallotBot assignment or Interim Access.

Effect on turnout and voting. The final set of results is on voting behavior. In Table 7, we report the effect of using BallotBot on political participation and the direction of voting. We look at the effect on voting at all on the assigned proposition (Column 1), voting at all in the presidential election (Column 2), and voting yes, separately by the assigned proposition (Columns 3-5).

We find no evidence that using BallotBot increases the probability of voting. As for the likelihood of voting “yes” on specific propositions, the estimates are somewhat imprecise due to the limited sample size. They generally indicate a null effect.

5. Conclusion

Can AI chatbots strengthen democracy? We provide some initial evidence that the answer is *yes*. BallotBot helps lower barriers to understanding ballot initiatives. Participants who used the BallotBot system performed better and faster on in-depth questions and had reduced frequency of over-confidence in incorrect responses. Moreover, less knowledgeable individuals who received performance feedback perceived a lower cost of answering an additional question. At the same time, chatbot availability encouraged broader exploration of ballot initiatives, although it did not change voting decisions or turnout rates.

These results have implications for both researchers and policymakers. First, they suggest that AI-based tools, when firmly grounded in reliable sources, can effectively address information challenges in democratic processes (Gentzkow, 2006; Allcott and Gentzkow, 2017; Matějka and Tabellini, 2020; Kartal and Tyran, 2022b). The value of addressing misinformation is magnified in the context of direct democracy, where voters have to vote directly on potentially complex laws and policies (Tyler, 2020). Second, the findings highlight that the benefits of this technology are heterogeneous, varying depending on the difficulty of the task and individual characteristics of users (Noy and Zhang, 2023; Brynjolfsson et al., 2025; Dell’Acqua et al., 2023). Third, the impact of AI support on reducing perceived information acquisition costs is conditional on respondents receiving clear feedback regarding the quality of the tool. This result suggests that, for AI tools to be effective in practical deployment, policymakers must build trust in their quality (Ahn et al., 2024). Finally, although the benefits appeared to diminish after one week, participants who continued using the chatbot between surveys showed stronger factual recall at follow-up. In contrast, participants who revisited the voter guide provided more thoughtful written justifications, suggesting that conventional resources also reinforce certain dimensions of learning. This points to a potential complementarity between traditional resources and AI-based tools.

We view these findings as an initial step that sets the stage for future research exploring several high-value issues. Future studies should investigate methods for improving the longevity of chatbot-driven learning effects, assess their influence on broader electoral outcomes, and determine whether AI interventions can enhance other forms of civic participation. Additionally, research should address potential unintended consequences, such as the risk of AI tools reinforcing echo chambers or selectively exposing users to

information aligned only with their existing beliefs. As AI tools continue to develop, addressing these questions becomes increasingly important for understanding the broader implications for the health and functioning of modern democracies.

References

- Ahn, Daehwan, Abdullah Almaatouq, Monisha Gulabani, and Kartik Hosanagar**, “Impact of Model Interpretability and Outcome Feedback on Trust in AI,” in “Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems” CHI ’24 2024.
- Allcott, Hunt and Matthew Gentzkow**, “Social Media and Fake News in the 2016 Election,” *Journal of Economic Perspectives*, May 2017, *31* (2), 211–36.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak**, “Measuring utility by a single-response sequential method,” *Behavioral Science*, 1964, *9* (3), 226–232.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond**, “Generative AI at work,” *The Quarterly Journal of Economics*, 2025.
- Costello, Thomas H, Gordon Pennycook, and David G Rand**, “Durably reducing conspiracy beliefs through dialogues with AI,” *Science*, 2024, *385* (6714), eadq1814.
- Cui, Kevin Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz**, “The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers,” *Working Paper*, 2024.
- Dell’Acqua, Fabrizio, Saran Rajendran, Edward III McFowland, Lisa Krayner, Ethan Mollick, François Cadelon, Hila Lifshitz-Assaf, Karim R. Lakhani, and Katherine C. Kellogg**, “Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” Working Paper No. 24-013, Harvard Business School 2023.
- Galletta, Sergio**, “Form of government and voters’ preferences for public spending,” *Journal of Economic Behavior & Organization*, 2021, *186*, 548–561.
- Gentzkow, Matthew**, “Television and Voter Turnout*,” *The Quarterly Journal of Economics*, 08 2006, *121* (3), 931–972.
- González-Bailón, Sandra, David Lazer, Pablo Barberá et al.**, “Asymmetric Ideological Segregation in Exposure to Political News on Facebook,” *Science*, July 2023, *381*.
- Jr, James M Snyder and David Strömberg**, “Press coverage and political accountability,” *Journal of political Economy*, 2010, *118* (2), 355–408.

- Kartal, Melis and Jean-Robert Tyran**, “Fake News, Voter Overconfidence, and the Quality of Democratic Choice,” *American Economic Review*, 2022, 112 (10), 3367–3397.
- and –, “Fake News, Voter Overconfidence, and the Quality of Democratic Choice,” *American Economic Review*, October 2022, 112 (10), 3367–97.
- Levy, Ro’ee**, “Social Media, News Consumption, and Polarization: Evidence from a Field Experiment,” *American Economic Review*, 2021, 111 (3), 831–870.
- Lupia, Arthur**, *Uninformed Why People Seem to Know So Little about Politics and What We Can Do about It*, Oxford University Press, 01 2016.
- Maskin, Eric and Jean Tirole**, “The politician and the judge: Accountability in government,” *American Economic Review*, 2004, 94 (4), 1034–1054.
- Matsusaka, John G.**, “Direct Democracy Works,” *Journal of Economic Perspectives*, 2005, 19 (2), 185–206.
- Matějka, Filip and Guido Tabellini**, “Electoral Competition with Rationally Inattentive Voters,” *Journal of the European Economic Association*, 09 2020, 19 (3), 1899–1935.
- Noy, Shakked and Whitney Zhang**, “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 2023, 381 (6654), 187–192.
- Nyhan, Brendan**, “Facts and Myths about Misperceptions,” *Journal of Economic Perspectives*, 2020, 34 (3), 220–236.
- Olken, Benjamin A.**, “Direct Democracy and Local Public Goods: Evidence from a Field Experiment in Indonesia,” *American Political Science Review*, 2010, 104 (2), 243–267.
- Pew Research Center**, “More than half of Americans are following election news closely – and many are already worn out,” 2024. Accessed: 2024-12-18.
- Schroeder, Elizabeth and Daniel F Stone**, “Fox news and political knowledge,” *Journal of Public Economics*, 2015, 126, 52–63.
- Tyler, Amanda L.**, “Direct democracy, misinformation, and judicial review in the United States,” in Sandrine Baume, Véronique Boillet, and Vincent Martenet, eds., *Misinformation in Referenda*, London: Routledge, 2020, pp. 110–127.
- Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov**, “Political Effects of the Internet and Social Media,” *Annual Review of Economics*, 2020, 12, 415–438.

Appendix A. Supplemental Material

Appendix A.1. Additional Tables and Figures

Table A1: Summary of Propositions

Title	Subject	Result	Yes Votes	No Votes
Proposition 2	Education	Approved	5,740,078 (57%)	4,323,911 (43%)
Proposition 32	Minimum wage	Defeated	5,758,176 (49%)	6,076,253 (51%)
Proposition 33	Housing	Defeated	3,864,542 (39%)	6,155,342 (61%)
Proposition 34	Campaign finance	—	4,995,616 (51%)	4,731,866 (49%)
Proposition 35	Healthcare	Approved	6,674,953 (67%)	3,311,778 (33%)
Proposition 36	Law enforcement	Approved	7,071,086 (70%)	3,008,016 (30%)
Proposition 3	Marriage	Approved	6,193,983 (61%)	3,919,777 (39%)
Proposition 4	Energy	Approved	5,887,354 (58%)	4,243,623 (42%)
Proposition 5	Direct democracy	Defeated	5,190,246 (44%)	6,514,650 (56%)
Proposition 6	Constitutions	Defeated	5,351,965 (46%)	6,233,020 (54%)

Notes: This table summarizes the outcomes of various propositions, including subject areas, results, and vote distributions. Propositions are categorized as "Approved" or "Defeated" based on voting outcomes.

Figure A1: Embedded Information Source - Quiz Task (Voter Guide)

Use the voter guide to answer this question below. If answering takes more than 2 minutes, select "I could not find the relevant information" and go to the next question.

The screenshot shows a voter guide page for Proposition 33. At the top, it reads "PROPOSITION 33 EXPANDS LOCAL GOVERNMENTS' AUTHORITY TO ENACT RENT CONTROL ON RESIDENTIAL PROPERTY. INITIATIVE STATUTE." Below this, it says "OFFICIAL TITLE AND SUMMARY" and "PREPARED BY THE ATTORNEY GENERAL". A note states: "The text of this measure can be found on page 100 and the Secretary of State's website at voterguide.sos.ca.gov." The main text is split into two columns. The left column contains a bullet point: "Current state law (the Costa-Hawkins Rental Housing Act of 1995) generally prevents cities and counties from limiting the initial rental rate that landlords may charge to new tenants in all types of housing, and from limiting rent increases for existing tenants in (1) residential properties that were first occupied after February 1, 1995; (2) single-family". The right column contains the text: "limiting the ability of cities and counties to maintain, enact, or expand residential rent-control ordinances." Below this is a section titled "SUMMARY OF LEGISLATIVE ANALYST'S ESTIMATE OF NET STATE AND LOCAL GOVERNMENT FISCAL IMPACT:" followed by a bullet point: "Reduction in local property tax revenues of at least tens of millions of dollars."

What fraction of Californians live in communities that already have local rent control?

One-quarter

One-half

One-third

Two-thirds

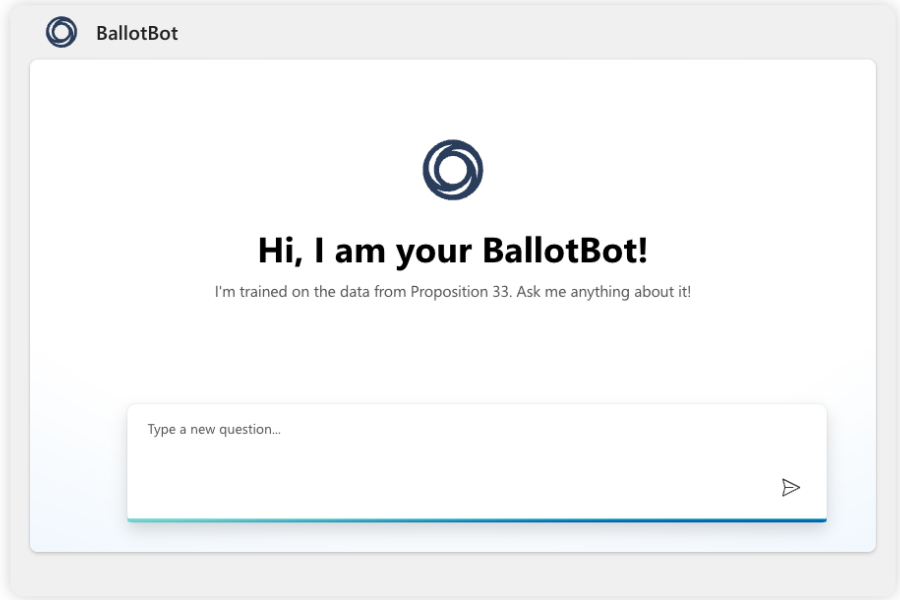
I could not figure out the answer

I could not find the relevant information

Notes: This screenshot represents the window that is displayed in the quiz task to a participant who was randomized into the Proposition 33, Voter Guide group.

Figure A2: Embedded Information Source - Quiz Task (BallotBot)

Use the chatbot to answer this question below. If answering takes more than 2 minutes, select "I could not find the relevant information" and go to the next question.



How much more do Californians typically pay for rent compared to other states?

- About 50% more
- About 10% more
- About 30% more
- About 70% more
- I could not figure out the answer
- I could not find the relevant information

Notes: This screenshot represents the window that is displayed in the quiz task to a participant who was randomized into the Proposition 33, Chatbot group.

Figure A3: BDM Mechanism for Eliciting Perceived Effort Cost

Are you still with us? You are doing **great!**

We have 1 more question for you, but this time **you can decide** how much we **pay you** in case of a **correct answer**. It can be a number between \$0 and \$1.

Be careful, though! We have already drawn a random number between 0 and 1. If your **bid** is **higher** than that **number**, you will **not** take the questions, and you will be directed to the **next step**. If, instead, your **bid** is **lower or equal to** the **number**, you will take the questions and be paid **your bid**.

Let's take an example: say that you bid \$0.20.

- If the random draw gives 0.30, you will get to answer the question and be paid \$0.20 for a correct answer.
- If the random draw gives 0.10, you will not get to answer the question and you will be redirected to the next section.

Select your bid using the slider, and click **Go to the Question!** whenever you are ready.

0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1

Select your Bid



Table A2: Covariates Attrition Test

Variable	Stayers Mean	Leavers Mean	Norm. Difference
Treatment Status	0.506 (0.500)	0.478 (0.500)	0.055 [0.397]
<i>Panel A: Demographics</i>			
Gender	0.477 (0.500)	0.454 (0.499)	0.046 [0.477]
Age	39.192 (13.422)	35.051 (12.374)	0.321 [0.000]
High Education	0.802 (0.399)	0.796 (0.404)	0.014 [0.831]
<i>Panel B: Socio Economic Status</i>			
Employment	0.737 (0.440)	0.762 (0.427)	-0.057 [0.381]
Democrat	0.553 (0.497)	0.498 (0.501)	0.111 [0.090]
Republican	0.160 (0.367)	0.197 (0.398)	-0.096 [0.151]
<i>Panel C: Attitudes Toward A.I.</i>			
Never Used A.I.	0.222 (0.416)	0.271 (0.445)	-0.113 [0.089]
How Useful is A.I.	3.678 (1.480)	3.711 (1.492)	-0.022 [0.731]
<i>Panel D: Baseline Knowledge</i>			
Know about the Proposition	0.577 (0.494)	0.554 (0.498)	0.045 [0.495]
Self-Perceived Kn.	0.418 (0.493)	0.403 (0.491)	0.030 [0.641]
<i>Panel E: Sources of Information</i>			
Social Media	0.630 (0.483)	0.685 (0.465)	-0.114 [0.076]
Traditional Media	0.394 (0.489)	0.397 (0.490)	-0.005 [0.944]
Official Documents	0.263 (0.441)	0.329 (0.471)	-0.143 [0.031]
Num. Obs.	1169	295	
F-Test of Joint Sig. (P-Val)			0.000
Num. Obs.			1463

Notes: We compare covariates between stayers (participants who remained in the study) and leavers (participants who left). Means, standard deviations (in parentheses), and normalized differences (in square brackets) are reported. Panels group variables into demographics, socio-economic status, attitudes toward A.I., baseline knowledge, and sources of information. The F-test checks for the joint significance of all covariates.

Table A3: Covariates Attrition Test

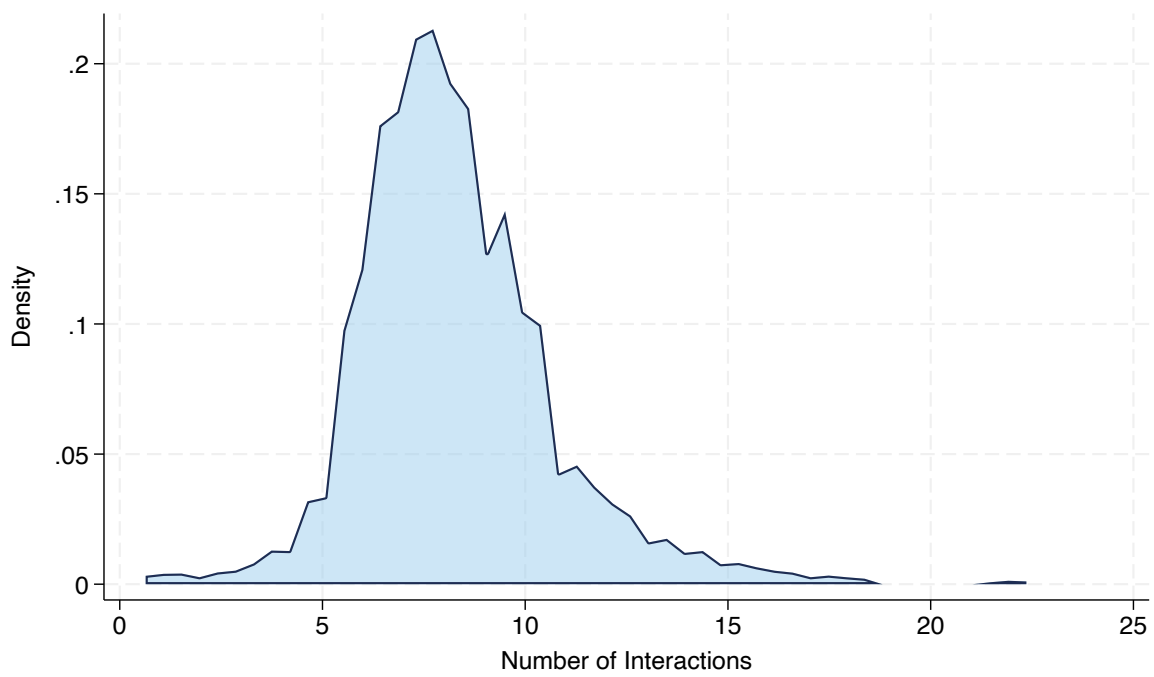
Variable	Stayers Mean	Leavers Mean	Norm. Difference
Treatment Status	0.475 (0.500)	0.516 (0.500)	-0.083 [0.126]
<i>Panel A: Demographics</i>			
Gender	0.464 (0.499)	0.479 (0.500)	-0.030 [0.578]
Age	35.880 (12.910)	39.854 (13.344)	-0.303 [0.000]
High Education	0.780 (0.415)	0.813 (0.390)	-0.081 [0.135]
<i>Panel B: Socio Economic Status</i>			
Employment	0.753 (0.432)	0.736 (0.441)	0.038 [0.478]
Democrat	0.540 (0.499)	0.544 (0.498)	-0.009 [0.871]
Republican	0.178 (0.383)	0.161 (0.368)	0.046 [0.399]
<i>Panel C: Attitudes Toward A.I.</i>			
Never Used A.I.	0.218 (0.413)	0.241 (0.428)	-0.054 [0.314]
How useful is A.I.	3.831 (1.390)	3.596 (1.528)	0.161 [0.003]
<i>Panel D: Baseline Knowledge</i>			
Know about the Proposition	0.567 (0.496)	0.575 (0.495)	-0.016 [0.772]
Self-Perceived Kn.	0.416 (0.493)	0.415 (0.493)	0.003 [0.963]
<i>Panel E: Sources of Information</i>			
Social Media	0.673 (0.470)	0.622 (0.485)	0.106 [0.049]
Traditional Media	0.385 (0.487)	0.400 (0.490)	-0.029 [0.587]
Official Documents	0.267 (0.443)	0.281 (0.450)	-0.032 [0.555]
Num. Obs.	550	913	
F-Test of Joint Sig. (P-Val)			0.000
Num. Obs.			1463

Table A4: Interaction Metrics with BallotBot

	Proposition 32	Proposition 33	Proposition 36
Total Interactions	2,308	2,271	2,226
Instances of Copy-pasting Questions	769	806	801
Instances of Copy-pasting Questions & Answers	14	6	11
Instances of No Answer from BallotBot	268	232	246

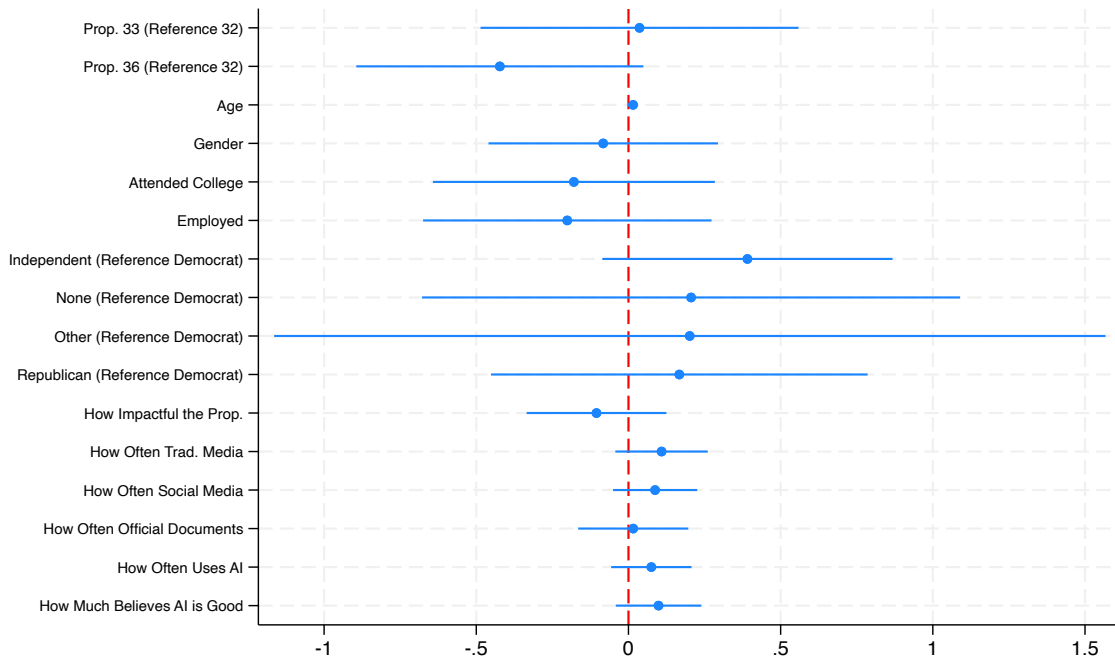
Notes: This table summarizes user interactions with BallotBot for three propositions. "Total Interactions" refers to all recorded responses by the system. "Instances of Copy-pasting Questions" highlights the number of times users directly pasted survey questions into the bot's input field. "Instances of Copy-pasting Questions & Answers" indicates occurrences where users pasted both the exact phrasing of questions and their corresponding answer options. "Instances of No Answer from BallotBot" captures cases where the bot was unable to provide an answer, often due to the absence of relevant information, as flagged by the response: *"The requested information is not available in the retrieved data. Please try another query or topic."*

Figure A4: Distribution of Interactions During the 1st Wave Survey



Notes: The number of interactions is computed as the number of questions that each user made while answering the survey. In [Figure A2](#) we report the user interface displayed on the survey page. In [Appendix A.2](#) we report an example of a conversation that amounted to three interactions.

Figure A5: BallotBot Usage Determinants



Notes: This plot displays the estimated coefficients that relate each dimension to the intensity of usage, measured as the number of interactions is computed as the number of questions that each user made while answering the survey. In Figure A2 we report the user interface displayed on the survey page. In Appendix A.2 we report an example of a conversation that amounted to three interactions.

Table A5: Effect of BallotBot on Accuracy, Log Response Time, and Confidence

Qst. Difficulty	Share Correct		Answering Time		Confidence	
	Basic	In Depth	Basic	In Depth	Basic	In Depth
	(1)	(2)	(3)	(4)	(5)	(6)
BallotBot	-0.009 (0.015)	0.137*** (0.013)	0.158*** (0.033)	-0.039 (0.028)	0.006 (0.074)	0.051 (0.034)
Contr. Mean	.821	.757	3.737	4.025	6.669	3.232
St. Dev.	.284	.28	.682	.612	1.438	.66
R ²	0.030	0.121	0.075	0.073	0.049	0.043
Num. Obs.	1463	1463	1463	1463	1463	1463

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 if participant is in the BallotBot group. The set of outcome variables are: share of correct answers, answering time in log seconds, and confidence in the answer on a scale from 1 to 4. Effects are reported separately for basic and in-depth questions, as described in the text. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors in parenthesis: *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

Table A6: Usage Balance Test - Guide Group

Variable	All Sample	Short Users Mean	Long Users Mean	Norm. Difference
<i>Panel A: Demographics</i>				
Gender	0.480 (0.500)	0.480 (0.500)	0.469 (0.502)	0.023 [0.833]
Age	38.400 (13.230)	37.956 (13.318)	41.344 (12.317)	-0.264 [0.013]
High Education	0.780 (0.410)	0.789 (0.408)	0.750 (0.435)	0.092 [0.409]
<i>Panel B: Socio Economic Status</i>				
Employment	0.740 (0.440)	0.748 (0.434)	0.656 (0.477)	0.201 [0.075]
Democrat	0.530 (0.500)	0.529 (0.500)	0.500 (0.503)	0.058 [0.595]
Republican	0.180 (0.380)	0.178 (0.383)	0.167 (0.375)	0.030 [0.783]
<i>Panel C: Attitudes Toward A.I.</i>				
Never Used A.I.	0.240 (0.430)	0.238 (0.426)	0.250 (0.435)	-0.028 [0.797]
How useful is A.I.	3.640 (1.500)	3.620 (1.493)	3.792 (1.514)	-0.114 [0.300]
<i>Panel D: Baseline Knowledge</i>				
Know about the Proposition	0.570 (0.490)	0.562 (0.497)	0.656 (0.477)	-0.193 [0.073]
Self-Perceived Kn.	0.420 (0.490)	0.430 (0.495)	0.365 (0.484)	0.133 [0.218]
<i>Panel E: Sources of Information</i>				
Social Media	0.660 (0.470)	0.666 (0.472)	0.604 (0.492)	0.129 [0.246]
Traditional Media	0.380 (0.490)	0.373 (0.484)	0.417 (0.496)	-0.089 [0.421]
Official Documents	0.260 (0.440)	0.260 (0.439)	0.281 (0.452)	-0.048 [0.663]
Total Score	4.310 (0.940)	3.921 (1.129)	3.865 (1.184)	0.049 [0.659]
Num. Obs.	1464	635	96	
F-Test of Joint Sig. (P-Val)				0.344
Num. Obs.				731

Table A7: Usage Balance Test - BallotBot Group

Variable	All Sample	Control Mean	Treatment Mean	Norm. Difference
<i>Panel A: Demographics</i>				
Gender	0.470 (0.500)	0.445 (0.497)	0.600 (0.492)	-0.313 [0.003]
Age	38.320 (13.410)	37.652 (13.035)	42.305 (14.932)	-0.332 [0.003]
High Education	0.820 (0.390)	0.820 (0.385)	0.800 (0.402)	0.050 [0.638]
<i>Panel B: Socio Economic Status</i>				
Employment	0.750 (0.430)	0.756 (0.430)	0.705 (0.458)	0.115 [0.284]
Democrat	0.560 (0.500)	0.542 (0.499)	0.667 (0.474)	-0.256 [0.013]
Republican	0.160 (0.370)	0.179 (0.383)	0.038 (0.192)	0.463 [0.000]
<i>Panel C: Attitudes Toward A.I.</i>				
Never Used A.I.	0.230 (0.420)	0.226 (0.419)	0.219 (0.416)	0.018 [0.865]
How useful is A.I.	3.730 (1.470)	3.699 (1.480)	3.886 (1.382)	-0.131 [0.203]
<i>Panel D: Baseline Knowledge</i>				
Know about the Proposition	0.570 (0.500)	0.568 (0.496)	0.581 (0.496)	-0.027 [0.801]
Self-Perceived Kn.	0.410 (0.490)	0.410 (0.492)	0.410 (0.494)	0.001 [0.994]
<i>Panel E: Sources of Information</i>				
Social Media	0.620 (0.480)	0.628 (0.484)	0.600 (0.492)	0.058 [0.583]
Traditional Media	0.410 (0.490)	0.392 (0.489)	0.514 (0.502)	-0.246 [0.021]
Official Documents	0.290 (0.450)	0.284 (0.451)	0.324 (0.470)	-0.087 [0.417]
Total Score	4.310 (0.940)	4.292 (0.954)	4.438 (0.808)	-0.165 [0.094]
Num. Obs.	1464	627	105	
F-Test of Joint Sig. (P-Val)				0.000
Num. Obs.				732

Table A8: Second Wave Quiz Results, by Basic and In-Depth

Qst. Difficulty	Share Correct		Answering Time		Confidence	
	Basic	In Depth	Basic	In Depth	Basic	In Depth
	(1)	(2)	(3)	(4)	(5)	(6)
BallotBot	0.000 (0.027)	-0.005 (0.021)	2.529 (1.748)	0.551 (1.334)	0.008 (0.029)	-0.010 (0.021)
Contr. Mean	.684	.651	26.174	26.328	.414	.412
St. Dev.	.465	.357	27.506	24.147	.493	.363
R ²	0.064	0.052	0.041	0.043	0.033	0.032
Num. Obs.	1158	1158	1158	1158	1158	1158

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 if the participant is in the BallotBot group. The set of outcome variables is: share of correct answers, answering time in seconds, and confidence in the answer on a scale from 1 to 4. All refer to the Second Wave Survey. Effects are reported separately for basic and in-depth questions, as described in the text. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors in parenthesis: *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

Table A9: Effects of Interim Access

<i>A. Effects on share of correct answers</i>				
	1st Wave Share Correct	2nd Wave Share Correct	2nd Wave Share Correct	2nd Wave Share Correct
	(1)	(2)	(3)	(4)
Interim Access	0.013 (0.018)	0.086*** (0.031)	0.091*** (0.031)	0.083*** (0.029)
Contr. Mean	.858	.65	.65	.65
St. Dev.	0.191	0.293	0.293	0.293
R^2	0.042	0.012	0.036	0.110
Controls	Yes	No	Yes	Yes
Previous Performance	No	No	No	Yes
Num. Obs.	732	591	591	591
<i>B. Effects on quality of reasoning</i>				
	-	Reasoning Quality	Reasoning Quality	Reasoning Quality
Interim Access	-	0.762*** (0.210)	0.857*** (0.208)	0.884*** (0.211)
Contr. Mean	-	4.865	4.865	4.865
St. Dev.	-	1.592	1.592	1.592
R^2	-	0.029	0.080	0.096
Controls	-	No	Yes	Yes
Previous Performance	-	No	No	Yes
Num. Obs.	-	468	468	468

Notes: In this table we present robustness checks for the result reported in Table 6. In Panel A, we focus on the BallotBot group and we show the null effect of interim access on the share of correct answers at the first wave quiz (column 1), the positive effect of interim access on the share of correct answers at the follow-up quiz without controls (column 2), with standard controls (column 3), and with standard controls + previous performance (column 4). In Panel B, we focus on the guide group and we present the effect of interim access on the quality of reasoning –as measured as described in Section 3.4– under the same different specifications. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning. Robust standard errors are provided in parentheses. Statistical significance is indicated as follows: *** p-value < 0.01, ** p-value < 0.05.

Table A10: Effect of BallotBot on Perceived Cost with Performance Fixed Effects

	Bid (\$)	Bid (by Prior Kn.)		Bid (by Education)	
	All Sample	High Kn.	Low Kn.	College	No College
	(1)	(2)	(3)	(4)	(5)
<i>A. With Feedback on Previous Performance</i>					
BallotBot	-0.024 (0.018)	0.015 (0.024)	-0.064*** (0.026)	-0.010 (0.020)	-0.115** (0.045)
Contr. Mean	.448	.425	.474	.448	.448
Sd. Dev.	.232	.221	.242	.236	.213
R ²	0.061	0.118	0.205	0.060	0.355
Num. Obs.	705	399	306	582	123
<i>B. Without Feedback on Previous Performance</i>					
BallotBot	0.025 (0.017)	0.041 (0.023)	0.009 (0.029)	0.019 (0.019)	0.037 (0.046)
Contr. Mean	.448	.425	.474	.448	.448
Sd. Dev.	.232	.221	.242	.236	.213
R ²	0.061	0.097	0.105	0.066	0.231
Num. Obs.	757	456	301	589	168

Notes: The main treatment variable, BallotBot, is an indicator equaling 1 for participants in the BallotBot (treatment) group. The outcome variable is the bid (in dollars) selected at the BDM task, where higher values indicate a higher perceived effort cost of answering an additional question. Panel A includes respondents that received information on their previous performance, while Panel B includes respondents who did not. Column 1 includes the whole sample; Columns 2 and 3 report heterogeneity by self-reported pre-treatment familiarity with the proposition; Columns 4 and 5 report heterogeneity by education level. All specifications include as controls demographics, socioeconomic conditions, attitudes toward A.I., news consumption habits, and political leaning, and previous performance Fixed Effects. Robust standard errors in parentheses. *** indicates a p-value < 0.01, ** indicates a p-value < 0.05.

Appendix A.2. Example BallotBot Conversation

An example of a conversation that a user had with BallotBot:

User: "what is prop 32 about"

BallotBot: "Proposition 32 is an initiative statute aimed at raising the minimum wage in California to \$18 per hour by 2026. The proposition is designed to help approximately 2 million Californians who currently earn less than \$18 per hour, many of whom are heads of households and have children.

Supporters argue that this increase is necessary to ensure that full-time workers can afford basic needs and to reduce the burden on taxpayers who subsidize low-wage workers. They believe that raising the minimum wage will stimulate local economies by increasing spending power among workers.

Opponents, however, argue that Proposition 32 could lead to higher costs for consumers, increased prices for goods and services, and potential job losses, particularly affecting small businesses. They also express concerns that it could worsen California's budget deficit and lead to cuts in essential services.

The fiscal effects of Proposition 32 are complex, with potential increases in government costs due to higher wages for public employees, but also possible savings from reduced enrollment in health and human services programs. Overall, the proposition has sparked significant debate regarding its potential impact on workers, businesses, and the state's economy."

User: "what are the controversies concerning this"

BallotBot: "Proposition 32 has sparked significant controversy, with arguments both in favor and against it.

Arguments in Favor:

1. ****Raising Wages****: Supporters argue that Proposition 32 will raise the minimum wage to \$18 per hour, benefiting about 2 million Californians who currently earn less than that amount. They contend that this increase is essential for workers to afford basic needs and reduce reliance on taxpayer-funded aid.
2. ****Economic Boost****: Proponents believe that higher wages will stimulate local

economies as workers will have more spending power, which could lead to job creation and overall economic improvement.

Arguments Against:

1. **Increased Costs**: Opponents argue that raising the minimum wage will lead to higher costs for consumers, particularly affecting small businesses that may struggle to absorb these costs. They warn that this could result in increased prices for goods and services.
2. **Budget Deficit Concerns**: Critics highlight that Proposition 32 could exacerbate California's existing budget deficit, potentially costing state and local governments billions each year. This could lead to higher taxes and cuts to essential services like education and public safety.
3. **Job Losses**: There are concerns that raising the minimum wage too quickly could lead to job losses, particularly among vulnerable populations, as businesses may reduce their workforce or close down due to increased labor costs.

Fiscal Effects:

The fiscal implications of Proposition 32 are complex. While it could lead to higher wages for many workers, it may also result in higher prices and reduced profits for businesses. The overall impact on jobs is expected to be minimal, but the potential for job losses remains a significant concern.

In summary, the controversies surrounding Proposition 32 revolve around its potential economic benefits versus the risks of increased costs, budget deficits, and job losses."

User: "have lawmakers considered the job losses part"

BallotBot: "Yes, lawmakers and analysts have considered the potential job losses associated with Proposition 32.

Concerns About Job Losses:

1. **Impact on Employment**: Opponents of Proposition 32 argue that raising the minimum wage too quickly could lead to job losses, particularly among vulnerable populations, such as young workers and those with less education. They suggest that when businesses face higher labor costs, they may reduce their workforce or even close down, resulting in job losses.
2. **Economic Analysis**: The Legislative Analyst's report indicates that the number of jobs in the state could either increase or decrease as a result of the proposition, but any

change is likely to be minimal—estimated to be smaller than one-quarter of 1%. This suggests that while there may be some job losses, the overall impact on employment levels is expected to be limited.

Summary: While there are concerns about job losses due to increased labor costs, the expected overall impact on employment is considered to be small. Lawmakers and analysts are aware of these potential consequences and have included them in their discussions and analyses of Proposition 32."

Appendix A.3. Knowledge Questions

Table A11: Questions for Proposition 32

Question	Options	Difficulty
If Prop. 32 were to pass, would the change in minimum wage be immediate?	a) Yes, for everybody; b) Yes, for employers with 26 employees or more ; c) No, wages will be updated in January 2025; d) No, except for employees with less than 20 employees; e) I could not figure out the answer; f) I could not find the relevant information	Basic
What would be the fiscal impact of Proposition 32?	a) State and local government costs might increase or decrease, while the tax revenues will likely decrease ; b) State and local government costs will increase, while the tax revenues will likely decrease; c) State and local government costs might decrease, while the tax revenues will likely increase; d) State and local government costs will decrease, while the tax revenues might increase or decrease; e) I could not figure out the answer; f) I could not find the relevant information	Basic
What adjustment happens to the minimum wage in 2027 under Proposition 32?	a) No adjustment is considered; b) It will be adjusted according to the labor market conditions; c) It will be adjusted according to inflation ; d) It will go back to \$16.50; e) I could not figure out the answer; f) I could not find the relevant information	Basic

Continued on next page

Question	Options	Difficulty
What is the current minimum wage in California?	a) \$20 per hour; b) \$13 per hour; c) \$16 per hour ; d) \$15 per hour; e) I could not figure out the answer; f) I could not find the relevant information	Basic
Does Prop. 32 amend also the current laws that establish a higher minimum wage in specified industries?	a) Yes, it raises the minimum also in those, by the same amount; b) No, it doesn't ; c) Yes, but only to a restricted set of industries; d) Yes, but only starting from 2026; e) I could not figure out the answer; f) I could not find the relevant information	Basic
In case Prop. 32 were not to pass, what would likely be the minimum wage in 2026?	a) \$14 per hour; b) \$16.50 per hour ; c) \$17 per hour; d) \$15 per hour; e) I could not figure out the answer; f) I could not find the relevant information	Basic
Does Proposition 32 directly affect the profits of local businesses?	a) Yes, it will likely have a negative effect on profits due to an increase in labor costs ; b) Yes, it will likely have a positive effect due to redistribution of resources; c) No, profits are likely to be unchanged by Proposition 32; d) Yes, but only in the long run; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth

Continued on next page

Question	Options	Difficulty
How does minimum wage affect smaller businesses and local communities, according to Proposition 32?	a) It will hurt smaller businesses by raising costs; b) It will reduce spending in local communities; c) It will benefit local economies by boosting spending and job creation; d) It will force small businesses to close down; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth
By what amount would prices likely increase for consumers as a result of Proposition 32's higher wages?	a) Less than one-half of 1 percent; b) 10 percent; c) 2 percent; d) 5 percent; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth
What is one likely outcome for the number of jobs in the state as a result of Proposition 32?	a) A large increase in jobs; b) A large decrease in jobs; c) A slight increase or decrease, smaller than one-quarter of 1 percent; d) A stop to job growth; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth
What specific rule does Proposition 32 follow when adjusting the minimum wage for inflation after 2027?	a) It adjusts based on California's inflation rate; b) It matches U.S. inflation unless inflation exceeds 3.5 percent; c) It increases only if inflation is above 5 percent; d) It stops adjusting for inflation after 2027; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth

Continued on next page

Question	Options	Difficulty
<p>What impact does Proposition 32 have on low-wage workers' reliance on taxpayer-funded assistance programs, such as Medi-Cal?</p>	<p>a) It increases their reliance on these programs; b) It decreases their need for these programs as their wages rise; c) It has no impact on their reliance on assistance programs; d) It eliminates eligibility for government assistance; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>

Table A12: Questions for Proposition 33

Question	Options	Difficulty
What would Proposition 33 allow cities and counties to control?	a) Rent prices for any type of housing; b) State income taxes; c) Water usage limits; d) Sales tax rates; e) I could not figure out the answer; f) I could not find the relevant information	Basic
How would Proposition 33 impact new renters moving into an apartment?	a) It could limit how much landlords can charge them for rent; b) It would give them tax deductions; c) It would offer them free utilities; d) It would provide free housing; e) I could not figure out the answer; f) I could not find the relevant information	Basic
What is one reason why rents are high in California, as mentioned in the guide?	a) The State does not provide enough housing; b) Overtourism; c) Renters formed a lobby that keeps prices high; d) The State provides too many houses; e) I could not figure out the answer; f) I could not find the relevant information	Basic
How much more do Californians typically pay for rent compared to other states?	a) About 50% more; b) About 10% more; c) About 30% more; d) About 70% more; e) I could not figure out the answer; f) I could not find the relevant information	Basic

Continued on next page

Question	Options	Difficulty
What is one of the key effects of Proposition 33 passing?	a) Expanding rent control in some communities ; b) Increasing income tax rates; c) Reducing corporate influence on elections; d) Eliminating property taxes; e) I could not figure out the answer; f) I could not find the relevant information	Basic
What are the main limits that Costa-Hawkins Rental Housing Act imposes to rent control?	a) It cannot be applied to single-family houses, to houses built after February 1, 1995, new renters ; b) It does not allow rent control in general; c) It cannot be applied to condominiums; d) It can always be applied, but with high maximum rent amounts; e) I could not figure out the answer; f) I could not find the relevant information	Basic
What fraction of Californians live in communities that already have local rent control?	a) One-quarter ; b) One-third; c) One-half; d) Two-thirds; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth

Continued on next page

Question	Options	Difficulty
<p>What impact could Proposition 33 have on local government budgets over time?</p>	<p>a) Property tax revenues may decrease due to a decline in the value of rental properties; b) It could lead to a significant increase in property tax revenues from landlords; c) Local governments would receive direct funding from the state to cover lost revenue; d) It would have no effect on local government budgets; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>
<p>What long-term financial effect could Proposition 33 have on California schools?</p>	<p>a) Schools may lose funding as a result of reduced property tax revenues from rent-controlled properties; b) Schools would receive additional funding from landlords; c) Proposition 33 would require local governments to increase school budgets; d) Proposition 33 would have no impact on school funding; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>
<p>What is one potential consequence of rent control expansion mentioned in the fiscal analysis?</p>	<p>a) Some landlords may sell their properties rather than rent them out; b) Rent prices will dramatically decrease across the board; c) The number of affordable housing units will increase; d) Local governments will see a massive boost in revenue; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>

Continued on next page

Question	Options	Difficulty
Does Proposition 33 affect any already existing local law on rent control?	a) Yes, it replaces any existing law; b) No, existing local laws would still apply ; c) Yes, but only for a restricted type of properties; d) Yes, but only for a restricted set of local entities; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth
According to the argument in favor of Proposition 33, what percent of their income do 55% of Californians pay for rent?	a) Less than 20%; b) 10%; c) More than 50%; d) More than 30% ; e) I could not figure out the answer; f) I could not find the relevant information	In-Depth

Table A13: Questions for Proposition 36

Question	Options	Difficulty
What type of crimes does Proposition 36 increase penalties for?	a) Traffic violations; b) Drug and theft crimes ; c) Environmental crimes; d) Property disputes; e) I could not figure out the answer; f) I could not find the relevant information	Basic
Under Proposition 36, how can some defendants have their felony drug possession charges dismissed?	a) By pleading guilty and completing treatment ; b) By paying a fine; c) By serving six months in prison; d) By doing community service; e) I could not figure out the answer; f) I could not find the relevant information	Basic
How are the crimes referred to in Proposition 36 currently charged?	a) Felonies; b) Misdemeanors ; c) Embezzlements; d) Infractions; e) I could not figure out the answer; f) I could not find the relevant information	Basic
What type of theft crimes can be charged as felonies under Proposition 36?	a) Shoplifting of items under \$500; b) Theft of items worth \$950 or less with two prior theft convictions ; c) First-time shoplifting offenses; d) Minor property damage crimes; e) I could not figure out the answer; f) I could not find the relevant information	Basic

Continued on next page

Question	Options	Difficulty
<p>What happens to individuals who do not complete treatment under Proposition 36's treatment-mandated felony program?</p>	<p>a) They are released without punishment; b) They serve up to three years in state prison; c) They pay a fine instead of serving time; d) They are transferred to community service; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>Basic</p>
<p>How does Proposition 36 address retail theft committed by multiple people?</p>	<p>a) It eliminates punishment for group theft; b) It increases sentences by up to three years for theft involving three or more people; c) It applies no additional penalties for group crimes; d) It reduces penalties for group theft; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>Basic</p>
<p>What should courts warn about to someone convicted of selling drugs if a person dies from the drugs they sold under Proposition 36?</p>	<p>a) That they will be fined; b) They can be charged with murder; c) That they will be given probation; d) That nothing happens; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>
<p>How does Proposition 36 change where some felony sentences (e.g., selling or providing certain drugs such as fentanyl) must be served?</p>	<p>a) It requires them to be served in county jail; b) It requires them to be served in state prison; c) It allows them to be served in home confinement; d) It removes jail time altogether; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>

Continued on next page

Question	Options	Difficulty
<p>What is the impact of Proposition 36 on the jail and community supervision population at the county level?</p>	<p>a) It will likely increase the county population by a few thousand people; b) It will significantly reduce the population; c) It will have no effect on the county population; d) It will reduce the population by transferring offenders to federal prisons; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>
<p>How does Proposition 36 potentially affect state savings from Proposition 47, which focused on mental health and drug treatment?</p>	<p>a) It increases those savings substantially; b) It reduces those savings substantially; c) It redirects funds to education; d) It has no effect on state savings; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>
<p>What specific drug-related felony charges are considered by the treatment-mandated felony provision in Proposition 36?</p>	<p>a) Possession of large amounts of cannabis; b) Possession of fentanyl, heroin, cocaine, or methamphetamine with prior convictions; c) Trafficking prescription drugs, also without prior convictions; d) Possession of fentanyl, heroin, or cocaine; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>
<p>What percentage of the state's General Fund budget is impacted by the increased costs from Proposition 36?</p>	<p>a) Less than one-half of 1 percent; b) 5 percent; c) 10 percent; d) 2 percent; e) I could not figure out the answer; f) I could not find the relevant information</p>	<p>In-Depth</p>