Crowdsourced Adaptive Surveys

Yamil Ricardo Velez[1]

[1]***Department of Political Science, Columbia University, Email: yrv2004@columbia.edu***

December 14, 2023

Abstract

Public opinion surveys are vital for informing democratic decision-making, but responding to rapidly changing information environments and measuring beliefs within niche communities can be challenging for traditional survey methods. This paper introduces a crowdsourced adaptive survey methodology (CSAS) that unites advances in natural language processing and adaptive algorithms to generate question banks that evolve with user input. The CSAS method converts open-ended text provided by participants into Likert-style items, and applies a multi-armed bandit algorithm to determine user-provided questions that should be prioritized in the survey. The method's adaptive nature allows for the exploration of new survey questions, while imposing minimal costs in survey length. Applications in the domains of Latino information environments and issue importance showcase CSAS's ability to identify claims or issues that might otherwise be difficult to track using standard approaches. I conclude by discussing the method's potential for studying topics where participant-generated content might improve our understanding of public opinion.

This is a working paper. Do not cite without permission.

Survey research plays a critical role in informing political decision-making and measuring fluctuations in public opinion (Page and Shapiro, 2010). However, traditional surveys may struggle to identify nascent issues that gain traction in the public consciousness due to the inherent time lag between design and implementation. These challenges are especially acute when examining social media trends or political events that can rapidly rise and fall in popularity. Furthermore, even when surveys manage to capture salient issues, the phrasing of questions and the determination of pertinent policy dimensions are frequently guided by subjective criteria set by the survey researchers.

In this paper, I propose a crowd-sourced adaptive survey methodology (CSAS) that leverages advances in natural language processing and adaptive algorithms to create participant-generated questionnaires that evolve over time. I use open-ended responses from participants to create question banks comprised of potential survey questions, from which questions are prioritized using a multi-armed bandit algorithm. While the survey is in the field, participants contribute to the question bank and rate existing questions, enabling the algorithm to "learn" emerging issues and claims that resonate with the study population.

I employ this methodology across two challenging domains: gauging the prevalence of misinformation within minority communities and evaluating issue importance in the aggregate. In an application examining the prevalence of negative beliefs about American political parties and candidates among Latinos, I find that CSAS recovers beliefs about events that would likely escape the notice of survey researchers. Participant-submitted items are also instructive, in that they focus more on party stereotypes or politically harmful – but true – claims, rather than discrete pieces of misinformation. I then use CSAS to construct an adaptive battery of nationally salient issues. Despite seeding the algorithm with popular items from Gallup surveys, I find that popular issue topics based on open-ended responses depart from the set of "most important issues," reflecting concerns over healthcare, inflation, political accountability, and crime.

Across the two studies, highly-rated items reflect salient news events or socioeconomic conditions, showcasing the method's ability to quickly adapt to changes in the information environment. Moreover, the exploration of survey questions using an adaptive algorithm such as Gaussian Thompson sampling produces issue and claim-level estimates for a variety of items; a feature that distinguishes the method from traditional approaches using fixed item batteries where only a set number of items can be examined. Specifically, as the algorithm attempts to "settle" on the best-performing question, the consideration of questions along the way renders it possible to compute prevalence estimates for an extensive set of claims.

The advantages of CSAS are two-fold. First, it enables survey researchers to capture trends in public opinion in real-time, reflecting the public's evolving beliefs and concerns. Second, it democratizes the survey process by allowing respondents to contribute to instruments. These benefits come at little cost in terms of survey length. Researchers can set the number of "dynamic questions" in advance, and select the appropriate algorithm for determining how questions should be prioritized. For example,

one can rely on a set of tried-and-true items, while allocating one or two survey slots for dynamic questions.

I close the paper by discussing how the CSAS method can be used to study other topics where information environments are evolving and scholarly assumptions about constructs might not match how study populations perceive them. Moreover, I highlight potential guardrails that researchers can implement to ensure a "clean" question bank with respect to quality, redundancy, and toxicity. Given rapidly evolving information environments in the age of social media, it is imperative that survey research is capable of adapting to these changes: the CSAS method is one step in that direction.

1 Dynamic Survey Methodologies: Existing Approaches

Influential texts on survey design stress the importance of a "tailored" approach to recruitment and stimuli (Dillman, Smyth, and Christian, 2014). Adapting questionnaires to respondents can enhance measurement and satisfaction. For instance, branching questions can reduce respondent burden and measurement error by eliminating irrelevant sub-questions (Dillman and Christian, 2005; Krosnick and Berent, 1993). Questions measuring recognition of elected officials and voting in subfederal elections can be tailored using location to produce more relevant question content via "piped in" text (Ansolabehere and Schaffner, 2009). These examples showcase how surveys already possess dynamic elements that respond to user input or data.

1.1 Computerized Adaptive Tests

Scholars have recently developed methods for carrying out computerized adaptive tests (CATs) in public opinion surveys (Montgomery and Cutler, 2013; Montgomery and Rossiter, 2020, 2022). CAT algorithms "respond to individuals' prior answers by choosing subsequent questions that will place them on the latent dimension with maximum precision and a minimum number of questions." (p. 173). CATs are typically employed using latent scales, where the goal is to optimize the number of questions. Montgomery and Cutler show that CATs offer a superior approach to traditional static batteries, and these tools can be easily implemented in survey software such as Qualtrics (Montgomery and Rossiter, 2022).

CATs rely on pre-existing measurement scales (e.g., political knowledge, personality batteries). However, in settings where the objective is to capture emerging issues or changes in the information environment, scholars and practitioners may want to learn about the prevalence of discrete beliefs, some of which cannot be known in advance. Thus, while CATs allow us to enhance precision when estimating latent traits, there are settings where the question bank cannot be fixed in advance and describing the nature of discrete items, rather than estimating positions on latent scales, is the primary objective.

One such setting is research on misinformation and rumors, where survey questions are often developed with assistance from fact-checking organizations (Vosoughi, Roy, and Aral, 2018) and social media databases (Sanderson et al., 2021). Though these sources can be useful, they are limited in their

timeliness and reach. First, fact-checks are lagging indicators, given that news organizations often wait until claims are viral before correcting them (Mena, 2019). Claims may be irrelevant by the time surveys are in the field. Second, identifying public claims through social media databases in real-time is possible, but the rising costs of APIs for platforms such as Twitter have made them less accessible to researchers.[1] Finally, the inclusion of questions can be affected by the institutional resources of fact-checking organizations and their assumptions of the study population. For example, fact-checkers focusing on Latino communities might prioritize immigration-related claims and inadvertently overlook other significant types of misinformation within those communities, such as health-related or economic misinformation.

1.2 Wiki Surveys

The Wiki survey is a collaborative survey methodology that incorporates user input, and thus, may be more flexible than approaches using fixed batteries (Salganik and Levy, 2015). Drawing inspiration from online information aggregation portals such as Wikipedia, Wiki surveys balance three principles: greediness, collaborativeness, and adaptivity. Greediness refers to capturing as much information as respondents are willing to provide, collaborativeness refers to allowing respondents to modify instruments (e.g., proposing new items), and adaptivity refers to optimizing instruments to select the "most useful information." While wiki surveys have shown promise in facilitating collective decision making (e.g., allowing users to vote on policies – both pre-determined and user-generated – that should be considered by local governments), existing applications rely on pairwise comparisons between options provided by survey designers and participants. However, pairwise comparisons may not be useful in settings where options can be accorded the same weight, the decision is not zero-sum, and outcomes can be more accurately measured on an ordinal or continuous scale.

In sum, existing adaptive survey methods offer distinct advantages over traditional surveys with respect to dimensions such as efficiency, but both have limitations when applied to certain question formats or research settings. CATs excel in improving precision for scales with fixed items, but may not be appropriate for settings where the question bank cannot be determined in advance. Wiki surveys, owing to their collaborative and adaptive nature, provide a means for creating surveys that evolve with user input but may not always be suitable, especially when attitudes, behaviors, or beliefs are best measured on an ordinal or continuous scale rather than through pairwise comparisons.

2 The Crowdsourced Adaptive Survey Method

Existing principles of Wiki surveys such as greediness, collaborativeness, adaptivity can guide the construction of methodologies achieving similar objectives. Building on the Wiki survey and other attempts to insert dynamic elements into existing surveys (e.g., CAT approaches), I develop a crowdsourced adaptive survey (CSAS) method that enables question banks to evolve based on user input and does not impose constraints on question formats. I leverage generalized pre-trained transformers

[1] https://www.theverge.com/2023/3/30/23662832/twitter-api-tiers-free-bot-novelty-accounts-basic-enterprice-monthly-price

(GPTs) to convert open-ended text produced by participants into questionnaire-friendly formats (see (Velez and Liu, 2023) for an example) and implement adaptive algorithms (Offer-Westort, Coppock, and Green, 2021) to identify best-performing questions from a question bank. First, each respondent answers an open-ended question about a given topic that is cleaned, summarized, and converted into a structured survey question format. Second, respondents rate their own questions, along with k other questions from a question bank generated by previous participants. Finally, ratings for user-submitted questions and k questions drawn from the question bank are updated using a multi-armed bandit algorithm, adjusting the probabilities of selecting these questions in subsequent surveys.

The three essential features of the proposed method are open-ended questions, a question bank, and a multi-armed bandit algorithm.[2] The open-ended question is used to query participants in a free-form manner about a given topic, issue, or claim. For example, a study examining rumors or misinformation could ask participants to provide a list of claims that present parties or candidates in a negative light. Since these data will be unstructured to some degree, introducing heterogeneity on dimensions such as length, style, and grammar, a response conversion stage is typically necessary. At this stage, researchers can refine the unstructured text to match their objectives. For example, given that Likert scale questions tend to be brief, one can convert the open-ended response into a sentence-long summary using a Generative Pre-Trained Transformer (GPT). GPTs are large-parameter language models that can perform various tasks such as text prediction, summarization, and classification at levels that mirror human performance (Vaswani et al., 2017).

Once open-ended response data are in a usable, structured format, they can be included in a question bank, which is a collection of structured survey questions or statements. Though inclusion of questions can be unrestricted at this stage, such that the question bank includes all responses that have been converted into a structured format, researchers may want to impose additional restrictions to reduce redundant questions and apply filters to ensure that survey questions meet the researcher's objectives. For example, a scholar studying misinformation may want to prioritize the inclusion of verifiable statements or claims versus value judgments. They may also want to minimize exposure to toxic and offensive content that can be produced by antagonistic respondents (Lopez and Hillygus, 2018).

Focusing on redunancy first, if two respondents submit responses about Democratic spending priorities with only minor differences in wording, it may be unnecessary to include both questions in the same question bank. Moreover, multi-armed bandit algorithms can struggle to identify the best-performing arm when arms are equally matched (i.e., "no clear winner"). Assuming near-identical questions are rated similarly, this increases the odds of failing to identify the best-performing item (Offer-Westort, Coppock, and Green, 2021, pp. 832–833). Given that open-ended responses are unlikely to be *exact* duplicates, filtering using more sophisticated text analysis methods such as document embeddings can be helpful. Document embeddings locate texts on a high-dimensional space and can

[2] One could also sample from the question bank using a uniform distribution. However, this entails devoting similar resources to both poorly performing and better performing items.

be used to identify similarities between texts (Rheault and Cochrane, 2020). By applying document embeddings, researchers can quantify the similarity between different questions, even when the wording varies, and retain only questions that surpass a pre-defined threshold of similarity.

Researchers may also want to apply additional filters to ensure that questions meet pre-specified criteria on dimensions such as relevance and toxicity. For example, for a survey of issue importance, a survey researcher may choose to exclude responses referring to the personal characteristics of politicians, and retain only those that refer to policies. Given that this is a classification task, one may opt for a supervised learning model trained on a labeled dataset or a GPT model, among other options. The same holds for identifying and removing toxic responses. Given that a small percentage of respondents resort to "trolling" or toxic behavior, ensuring other participants are not exposed to harmful content is paramount (Lopez and Hillygus, 2018).

The next challenge is in how these questions are presented and selected within the survey. Instead of presenting questions with equal probability, and thus "wasting" responses on questions with low ratings, multi-armed bandit algorithms can be leveraged to identify a set of "best-performing" questions from the pool. Multi-armed bandits have been applied in recent research on adaptive experiments, where the goal is to identify the best-performing intervention from a set of interventions (Offer-Westort, Coppock, and Green, 2021). Rather than assign uniform probabilities to each arm, multi-armed bandit algorithms such as Thompson sampling "learn" the most-effective arm over a set of trials by balancing the trade-off between exploitation (i.e., concentrating on promising arms) and exploration (i.e., experimenting with additional arms to assess their potential). In experimental settings, outcome information is used to determine future treatment allocation, such that arms associated with higher scores on pre-specified metrics are generally prioritized. When measuring the prevalence of a false or contested claim, for example, Thompson sampling can also be employed. Instead of interventions, different survey questions or statements can be considered "arms." By iteratively selecting questions based on their estimated prevalence, Thompson sampling can adaptively allocate more participants to claims that resonate with the sample.

3 An Application to Latino Information Environments

I use the CSAS method to identify rumors, negative political claims, and misinformation reaching Latinos, a group that has received attention among journalists and social scientists due to potential misinformation campaigns targeting the community (Cortina and Rottinghaus, 2022; Velez, Porter, and Wood, 2023). I focus on Latinos for two reasons. First, fact-checking is still a relatively new institution within Latino-oriented media (Velez, Porter, and Wood, 2023, p. 790). Existing organizations might overlook important claims that circulate within the community due to resource constraints and the possibility that best practices for verification have not yet been identified. Second, private, encrypted messaging applications used by Latinos such as WhatsApp and Telegram may hinder the detection of false claims (Lee, Moore, and Hancock, 2023). In contrast to misinformation that is transmitted

through social media such as Instagram, Twitter, and Facebook, these private channels can operate largely beyond the reach of fact-checkers and researchers. These two factors create unique challenges for identifying and addressing misinformation within the Latino community.

Implementing this more "bottom-up" approach to misinformation detection, I fielded a survey of 321 self-identified Latinos using the survey platform, CloudResearch Connect, from July 6-7, 2023. First, participants were asked two open-ended question regarding negative claims they had heard about Republicans and Democrats. These claims were then passed to a fine-tuned generative pre-trained transformer model, OpenAI's *ada*, that classified the text as a "verifiable claim." Fine-tuning was necessary to ensure that questions entering the question bank were falsifiable political claims, rather than value judgments (e.g., politicians are evil). To carry out the fine-tuning step, a mixture of researcher-provided examples and participant-provided examples (N=50) were hand-coded to indicate whether a claim was, in principle, falsifiable. Hand-coded classifications were then used to fine-tune the *ada* model.

Given that similar questions can be introduced by more than one respondent, and to guard against exposing participants to toxic content, I also used a similarity and toxicity filter before adding items to the question bank. For each submitted question, I used OpenAI's embedding API to generated a 1536-dimensional embeddings vector, calculated the cosine similarity between the new item and other items in the question bank by retrieving the five nearest neighbors, and filtered out questions with a similarity score above .90.[3] I also used OpenAI's moderation endpoint to filter out offensive and toxic claims. Claims that passed these filtering steps were added to an item bank that allowed them to be rated by future participants.

After the submission and cleaning step, participants rated their own submissions, along with four items from the question bank and six items capturing conspiracy beliefs and more common misinformation items (e.g., Covid-19 vaccines modify your DNA). All of the questions were presented in a matrix format, with a four-point accuracy scale ranging from "not at all accurate" to "very accurate." Since the first set of participants did not have user-submitted claims to rate, the question bank was seeded with an initial set of claims. Four claims were taken from the front pages of Latino-oriented fact-checking websites (i.e., Telemundo's T-Verifica, Univision's El Detector) to create the initial question bank.

Gaussian Thompson Sampling was then used to determine which questions to present to subsequent participants. Though traditional Thompson Sampling requires binary outcomes due to its use of a Beta distribution (Offer-Westort, Coppock, and Green, 2021), Gaussian Thompson Sampling can be leveraged to accomodate continuous outcomes (Agrawal and Goyal, 2017).[4] Gaussian Thompson sampling was implemented in real-time using an API created by the author. In contrast to previous

[3]In initial tests before data collection, lower similarity thresholds such as .80 were found to produce false positives (e.g., classifying "Biden is a tax cheat" and "Trump is a tax cheat" as sufficiently similar) and higher similarity thresholds such as .95 produced false negatives.

[4]As in other research (Offer-Westort, Coppock, and Green, 2021), a probability floor of .01 was employed to guarantee that every item in the item bank retained a non-zero chance of being presented to participants.

applications of adaptive experiments in political science that have leveraged batch-wise Thompson sampling, probabilities were adjusted at the respondent level.

4 Results

Figure 1 displays mean accuracy estimates for claims receiving more than ten ratings.[5] The highest-rated claims covered information about both Republicans and Democrats. Party stereotypes about Republican positions on moral issues were rated as highly accurate ($\bar{x}$ = 3.54; SE = .07), as were claims that Trump used the presidency to enrich family and friends ($\bar{x}$ = 3.54; SE = .159) and possessed classified documents in his vacation homes ($\bar{x}$ = 3.43; SE = .121). Other highly rated claims focused on Republicans such as Ted Cruz ($\bar{x}$ = 3.40; SE = .049), extreme policy positions (e.g., a claim that Trump plans to deport all undocumented people ($\bar{x}$ = 3.32; SE = .313), and President Biden's son, Hunter Biden ($\bar{x}$ = 3.24; SE = .201). The lowest-rated claims typically involved false statements, allegations, or extreme descriptions of issue positions. Claims that scored especially low on perceived accuracy included "This year, there was a major explosion close to the White House" ($\mu = 1.50$; SE = .15), "There were no wars during the Trump presidency" ($\bar{x} = 1.74$; SE = .17), "Republicans support going back to the 1776 constitution." ($\bar{x} = 1.80$; SE = .33), "Trump will deport **all** immigrants if he wins in 2024" ($\bar{x} = 1.91$; SE = .19), and "The president was involved with his son's bribes" ($\bar{x} = 1.92$; SE = .20).

Whereas highly rated claims mostly reflected actual events or generalizations of party positions, claims with lower accuracy ratings typically involved verifiably false information or oversimplifications. These findings are instructive in that they reveal a level of discernment in the aggregate. Objectively false claims are generally seen as less credible by participants. Instead, higher accuracy ratings are observed when participants judge claims that are widely reported in the news (e.g., Ted Cruz's cancun trip) or that reflect commonly-held perceptions of party positions (e.g., Republicans opposing aboriton and same sex marriage). Though the initial seed claims based on fact-checks were small in number, the analysis revealed a surprising disparity: the most readily believed claims were not identified by the fact-checking organizations, but rather originated from other participants. This reinforces the utility of the method, as it highlights the potential for crowdsourcing to identify and address misinformation that can complement traditional fact-checking efforts.

Moving beyond the aggregate estimates, one can further segment the sample into subgroups. For example, one can identify whether highly rated claims are uniformly perceived as accurate or whether subgroups that may be more prone to seeing misinformation might also differ in their accuracy ratings. Previous research, as well as popular discussions of Latino-targeted misinformation, have highlighted WhatsApp as a potential vector for misinformation. In the survey, participants were asked to rate the credibility of information on the platform on a five-point ordinal scale ranging from "Trust" to "Distrust."

Comparing participants below and above the median of WhatsApp trust, several patterns emerge. For the lowest rated claims, there is a general consensus regarding accuracy. However, for the

[5] Mean estimates are weighted by the inverse probability of selection, as in Offer-Westort, Coppock, and Green (2021).

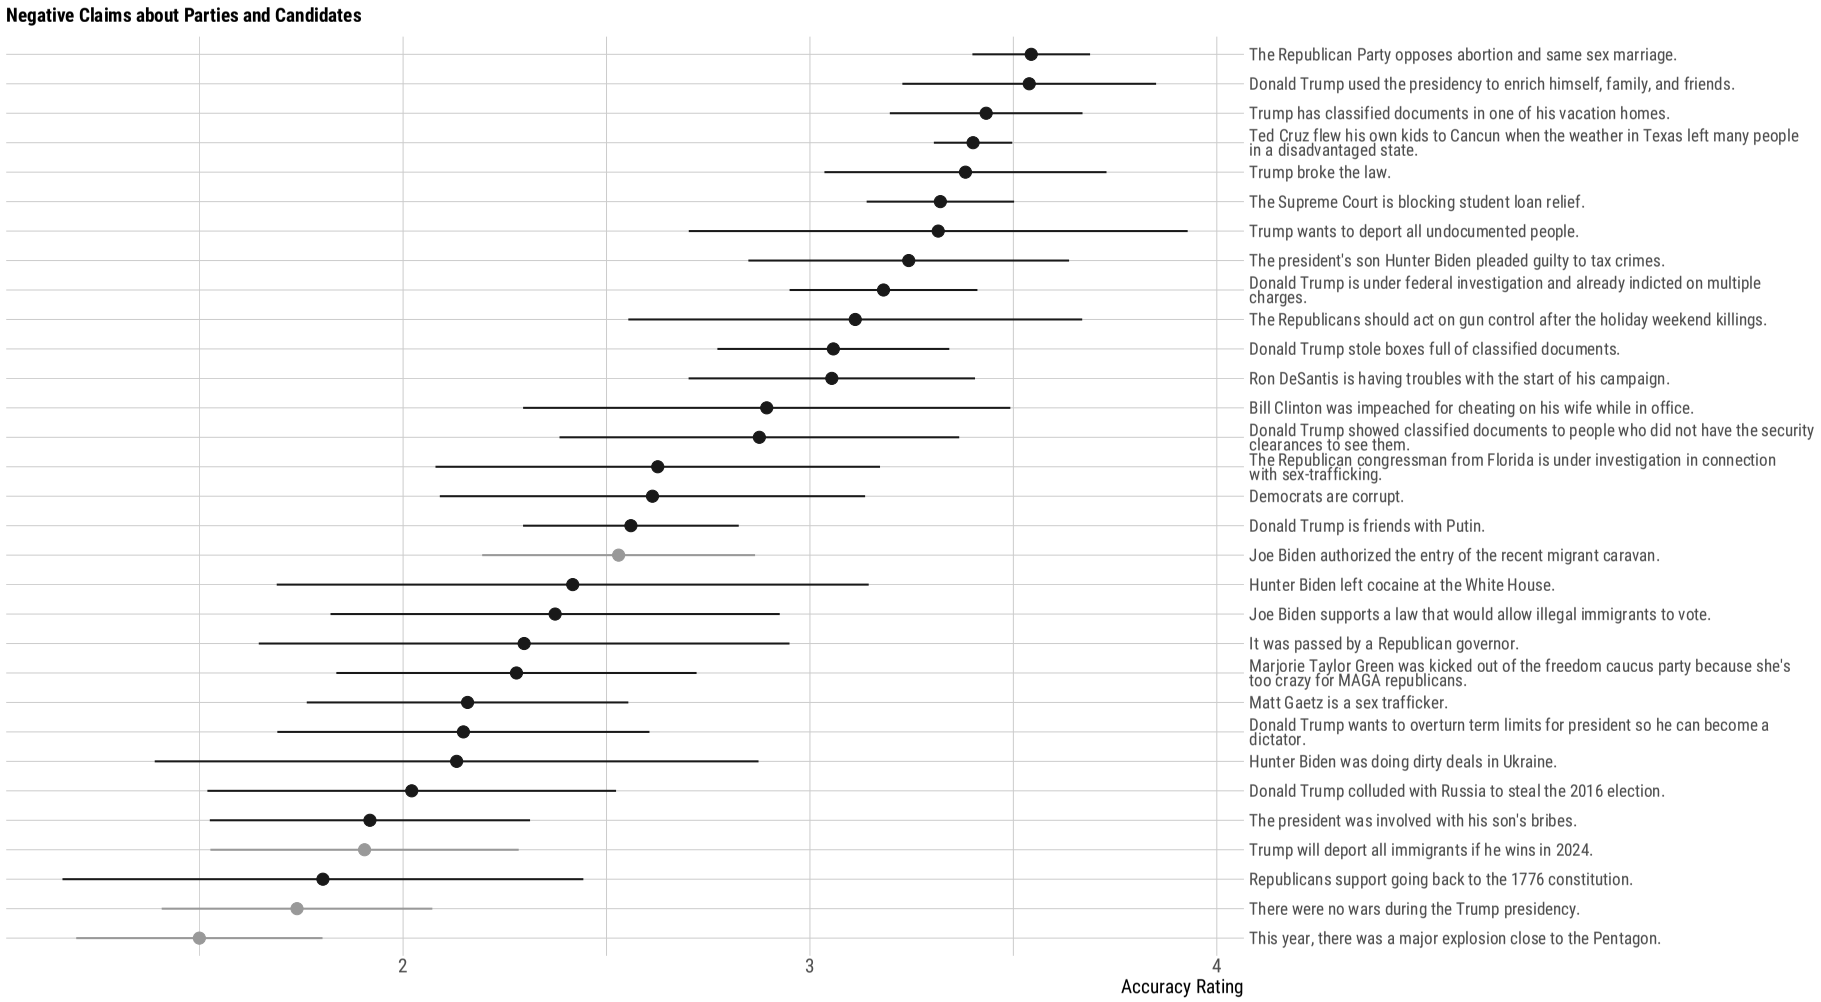

Figure 1: IPW-weighted estimates of survey questions measuring negative political claims with corresponding 95% confidence intervals. Items in gray were initial seed items based on fact-checked claims produced by Latino-focused fact-checking organizations. Black items are participant-generated items.

Figure 2: IPW-weighted estimates of survey questions measuring negative political claims across levels of WhatsApp trust with corresponding 95% confidence intervals. Presented claims were rated by at least 10 participants in each category.

more highly rated claims involving party stereotypes, newsworthy scandals, and salient allegations, differences between the two groups can be seen. For example, those who trust WhatsApp were less likely to have rated true claims about the Supreme Court blocking student loan relief, Ted Cruz visiting Cancun with his family during widespread blackouts in Texas, and Donald Trump being under investigation as accurate relative to other users. These gaps, of course, could reflect other social or demographic causes, but they point to the possibility that observed differences between WhatsApp users and non-users in factual accuracy can be derived from gaps in knowledge of factual political information, rather than a wholesale adoption of misinformation. The relatively flexible nature of the approach, which allows for the inclusion of blatantly false misinformation as well as misperceptions or overgeneralizations allows us to arrive at a more nuanced description of information environments than if we solely measured accuracy ratings for factually verified content.

Our results underscore the complexity of the information environment among Latinos. They engage with a variety of narratives, some of which portray different parties negatively, but also reflect actual events. Understanding this diversity is essential for creating informed strategies to address misinformation and enhance democratic discourse. However, although the CSAS approach can recover useful questions, filtering can be imperfect. Claims like "The Republican party opposes abortion and same sex marriage" departs from more salacious claims about politician behavior that tend to receive attention from fact-checkers. Moreover, there are low-quality items that were added to the question bank despite the filtering. One such example was "It was passed by a Republican governor." Still, this item received low ratings, and thus, was not seen by many participants. This further emphasizes the benefits of using an adaptive algorithm, given that items with lower ratings will be prioritized less.

5 An Application to Issue Importance

Though the CSAS method can be helpful in crowdsourcing questions that resonate within marginalized communities, identifying salient issues is another possible application. That is, rather than fixing a set of issues for participants to rate, as is often done in research on issue attitudes and importance (Ryan and Ehlinger, 2023), one can use the method to provide an issue bank for participants, much like the question bank comprised of claims in the previous study. This can be helpful in estimating support for idiosyncratic issues that may not appear on the national agenda, but still inspire strong reactions among "issue publics" (Elliott, 2020) or serve as issues that could be mobilized in future elections, corresponding to the elusive notion of "latent opinion" described in Key (1961).

Since 1935, Gallup's Most Important Problem question has been used to approximate the issue priorities of the American public. Using an open-ended format, participants are asked "What do you think is the most important problem facing this country today," with responses being classified by hand into a set of categories corresponding to broad issue or policy areas. Though the question has undergone several changes over the years, its consistent use across surveys has enabled scholars to construct time series of a host of issues (McCombs and Zhu, 1995), and has served as the empirical

foundation for numerous studies using national issue salience as an outcome (Dunaway, Branton, and Abrajano, 2010) and independent variable (Ostrom Jr et al., 2018).

Despite its adoption in public opinion research, the measure has been criticized for being an imperfect proxy of national issue salience. As Wlezien (2005) argues, the question asks respondents to provide information on two distinct concepts — importance and "problem status." While some respondents may interpret the question as one where they can offer a personally relevant issue, others may interpret it as an opportunity to highlight a problem affecting the nation as a whole. Focusing on defense spending, Wlezien (2005) finds that increases in MIP estimates regarding foreign policy are predicted by indicators of "problem status" (e.g., dislike of foreign adversaries), but fail to explain defense spending preferences once indicators of "problem status" are included. Though these analyses suggest that revisions of question wording could yield more accurate responses, Jennings and Wlezien (2011) find that a most important issue (MII) question that eliminates the use of "problem" language performs similarly over time, and thus, "an effective measure of issue importance remains elusive" (p. 546).

More recently, Ryan and Ehlinger (2023) make a case for moving beyond fixed questions about a slate of national issues and hand-coded classifications based on open-ended questions. Like the MIP and MII, Ryan and Ehlinger (2023) use open-ended questions to elicit issues from participants. However, in contrast to these approaches, they convert unstructured text by hand into rating scales, and present these scales to participants in a follow-up wave. This approach provides a richer amount of information about the *degree* to which an elicited issue may hold importance to an individual. Recent studies in this vein have recovered high levels of stability (Velez and Liu, 2023) and sizable causal effects for this "core issue" in conjoint settings (Ryan and Ehlinger, 2023; Velez, 2023), suggesting that such a measure might serve as a candidate for the elusive concept of issue importance.

Applying the CSAS method in this context is straightforward. Participants report personally relevant issues using open-ended responses; open-ended responses are transformed into topics for rating and included in a question bank; and participants rate the different issue topics with Gaussian Thompson sampling used to optimize question selection. In contrast to the misinformation setting, where it may be more difficult to select seed questions because of a lack of high-quality data on Latinos, the top-performing MIP items produced by Gallup can be used as a baseline for comparison. Specifically, we can seed the question bank with a set of popular issues and assess whether the crowd-sourced issue topics receive higher importance ratings. The issues were the following: 'Immigration', 'Economy', 'Race Relations', 'Poverty', 'Crime', 'Ethics, Moral, and Family Decline', 'Unifying the country', and 'Inflation.'

From September 11 until September 13, 2023, I collected data from a national quota sample balanced on age, race, and gender using CloudResearch Connect (N=820). The question bank was seeded with eight popular issue areas taken from Gallup. In this study, a more capable LLM, OpenAI's GPT-4, was used to convert unstructured text into issue topics and filter out redundant issues in

one shot. Recent studies have shown that in the task of classifying open-ended text to identify "most important issues," the efficacy of Large Language Models (LLMs) like GPT-4 is on par with classification algorithms trained on thousands of examples, achieving performance levels marginally below that of human evaluators (Mellon et al., 2022). User-submitted items spanned salient issues such as abortion and climate change, and less salient issues such as universal basic income and deficit spending. Each participant rated their own issue, along with eight others that were determined using Gaussian Thompson sampling, as in the first study.

5.1 Results

Figure 3 displays mean estimates for issue topics receiving 50 or more ratings. As shown in the figure, the highest-rated issues were focused on the economy and health care, with issues such as "Cost of Living," "Healthcare Affordability," "Healthcare Costs," "Economic Stability," and "Universal Healthcare." Issues rated lower on importance include more social and culture issues such as immigration and voting rights. The MIP issue topics of "Race Relations" and "Ethics, Moral, and Family Decline" appeared among the lower issue ratings, with topics related to immigration (i.e., "Illegal Immigration," "Border Security") being given lower ratings. In the list of highly rated topics, we see issues that would likely not appear in traditional issue importance barries such as "Mental Health Access," "Privacy Protections," and "Candidate Transparency." Moreover, the frequent mention of various economic and healthcare dimensions is instructive, revealing a trend where socioeconomic issues like healthcare and cost of living are taking center stage.

As in the previous study, I present issue importance ratings across relevant subgroups. Figure 4 presents mean estimates of issue importance among Republicans and Democrats. Estimates are ordered by the absolute size of the partisan gap. As seen on the bottom of the figure, there some areas of issue consensus across the two parties, namely those related to healthcare costs, political polarization, and economic issues (i.e., cost of living, economic stability). However, by and large, issue importance ratings exhibit considerable gaps across partisans. For example, Democrats generally rate inequality, progressive policies (e.g., universal healthcare), gun control, and social issues (e.g., women's rights, race relations) as more important than Republicans, whereas Republicans rate issues involving morals, government spending, and immigration higher than Democrats. While partisan differences in issue positions are well known and expected, the dynamic nature of adaptive algorithms allows us to delve into a broader range of "issue gaps" than traditional surveys with fixed sets of questions.

To what extent does our method capture personally meaningful issue positions? As mentioned earlier, recent studies suggest that issues elicited in an open-ended format score high on dimensions such as attitude strength, certainty, and reliability over time. Still, we can use the data from the existing study to verify whether this is the case. First, we find that on a 1-5 scale ranging from "Not at all important" to "Very important," the mean importance score for the personal issue is 4.7 (s = .72). These ratings are generally higher than importance ratings for other issues in the question bank ($\bar{x}$ =

3.43; s = 1.24). Moreover, the mean ranking for the "personal issue" is 1.43 in the list of nine rated issues.

Overall, the findings highlight the ability for CSAS to identify idiosyncratic issues that may be deeply important to segments of the population. For example, universal healthcare and other health-related issue topics received some of the highest importance ratings. Though healthcare policy is not currently a salient issue in American politics, the CSAS method reveals that it remains a deeply important concern for individuals in our sample. The findings offer a valuable glimpse into less salient issue priorities and suggest the potential for the CSAS method to complement traditional measures of public opinion.

6 Concerns and Caveats

6.1 Is the CSAS method compatible with traditional survey design?

Despite the limitations of "static surveys" in identifying changing information environments or measuring responses within niche communities, the two approaches are not at odds. Researchers can decide the number of adaptive questions, and include these questions in standard batteries. For example, in the Latino survey, participants rated a pre-existing set of false claims and conspiracies, along with an adaptive set, in a question matrix. Before using this method, scholars should determine whether the marginal benefit of having a designated slot for exploratory questions is worth the survey time and cost. A distinct advantage of multi-armed bandits is that several items can be explored despite having a smaller set of "dynamic survey slots." The two approaches can also work in tandem when there are multiple phases of data collection. An initial wave (or pilot) could use CSAS to develop a fixed battery of questions for future waves, functioning much like pilot studies that gather open-ended data to inform scale construction (Weller, 1998). With CSAS, however, future surveys can be designed not only with open-ended content in hand, but question ratings and posterior probabilities that a given question is the "best-performing question." This approach may be optimal if researchers prefer to split their research process into exploratory and confirmatory stages, as is recommended in Egami et al. (2018).

6.2 Prompting participants

Leaving aside potential challenges with prompting large language models, prompting participants is also an important step in the process. For example, in the case of misinformation, asking about false claims directly would likely fail to elicit meaningful items, given that participants may not know the veracity of statements or may be inclined to disregard validity when thinking about congenial pieces of information. Instead, asking about claims that possess common features of misinformation, such as negativity, appear to be more fruitful (Carrasco-Farré, 2022). In extending this approach to construct other kinds of question banks, researchers should carefully consider whether the questions they create may diverge significantly from how respondents conceptualize them. Though best practices involving closed-ended questions are well-developed, the use of LLMs and other text analysis methods to extract

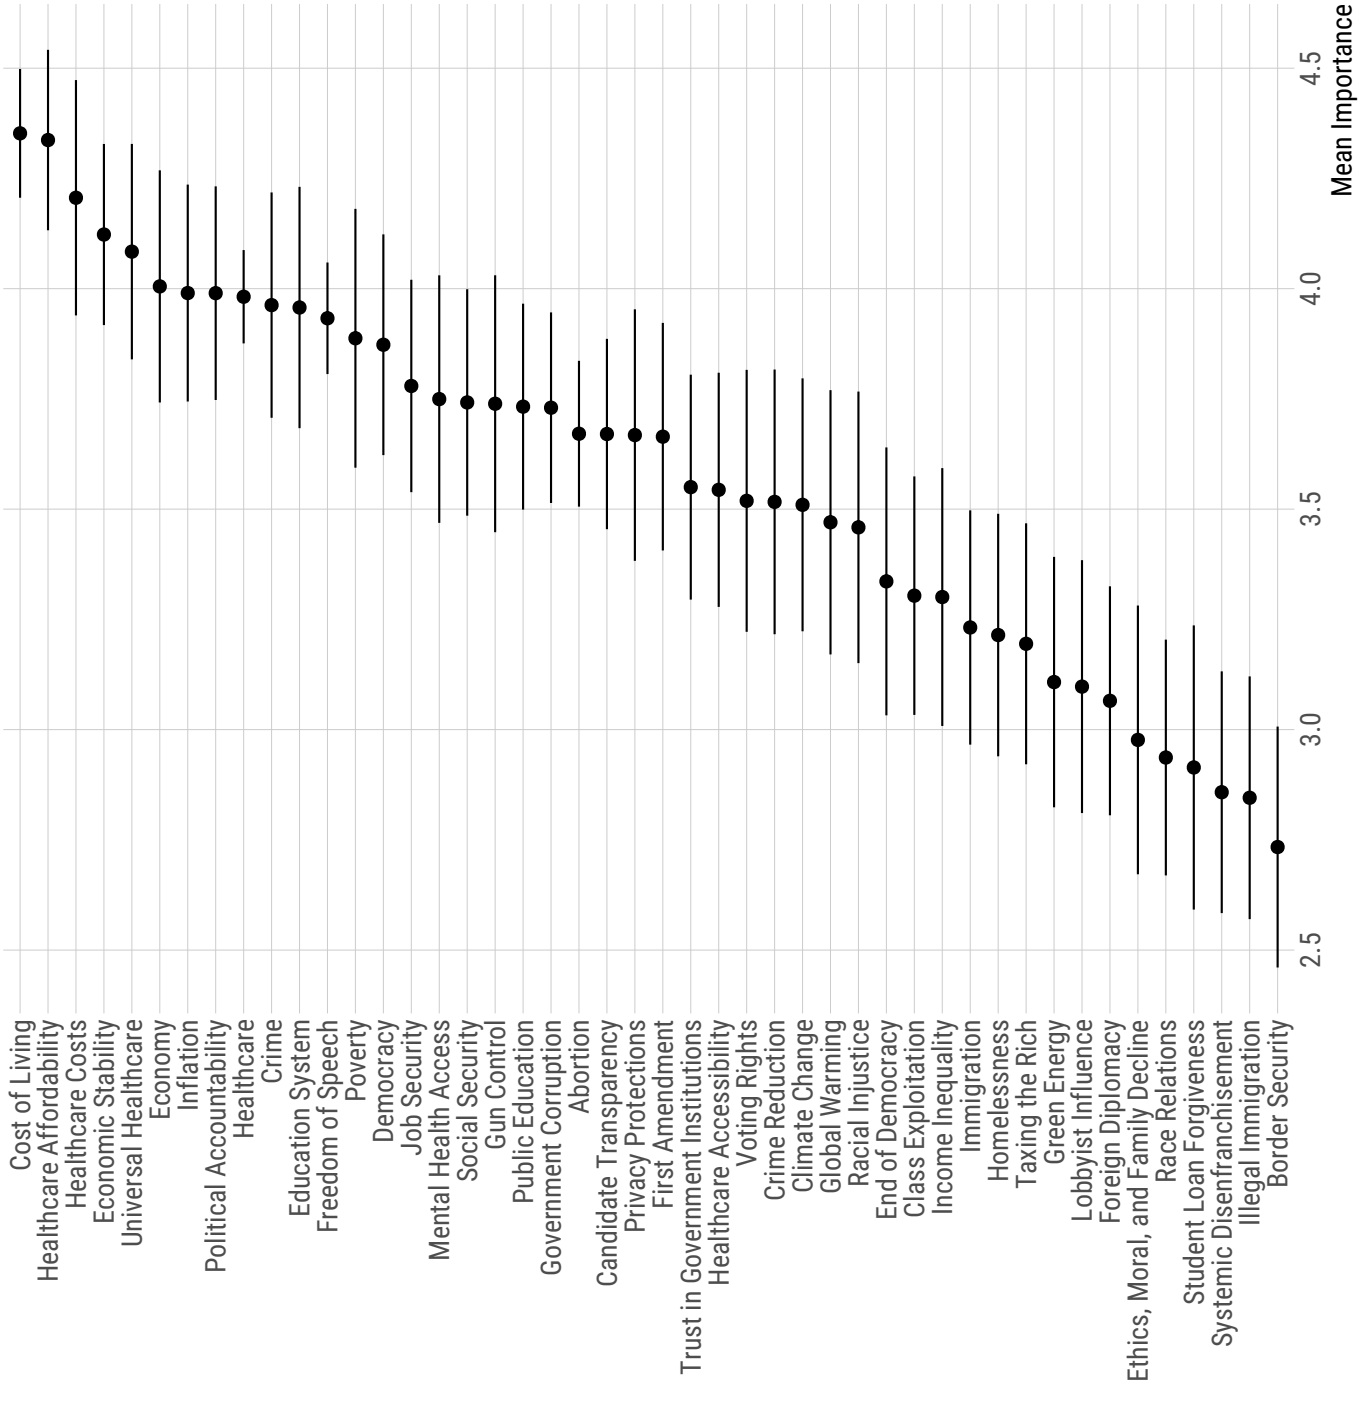

Figure 3: IPW-weighted estimates of survey questions measuring issue importance with corresponding 95% confidence intervals.

https://doi.org/10.33774/apsa-2023-r2jzm **ORCID:** https://orcid.org/0000-0002-3243-0975

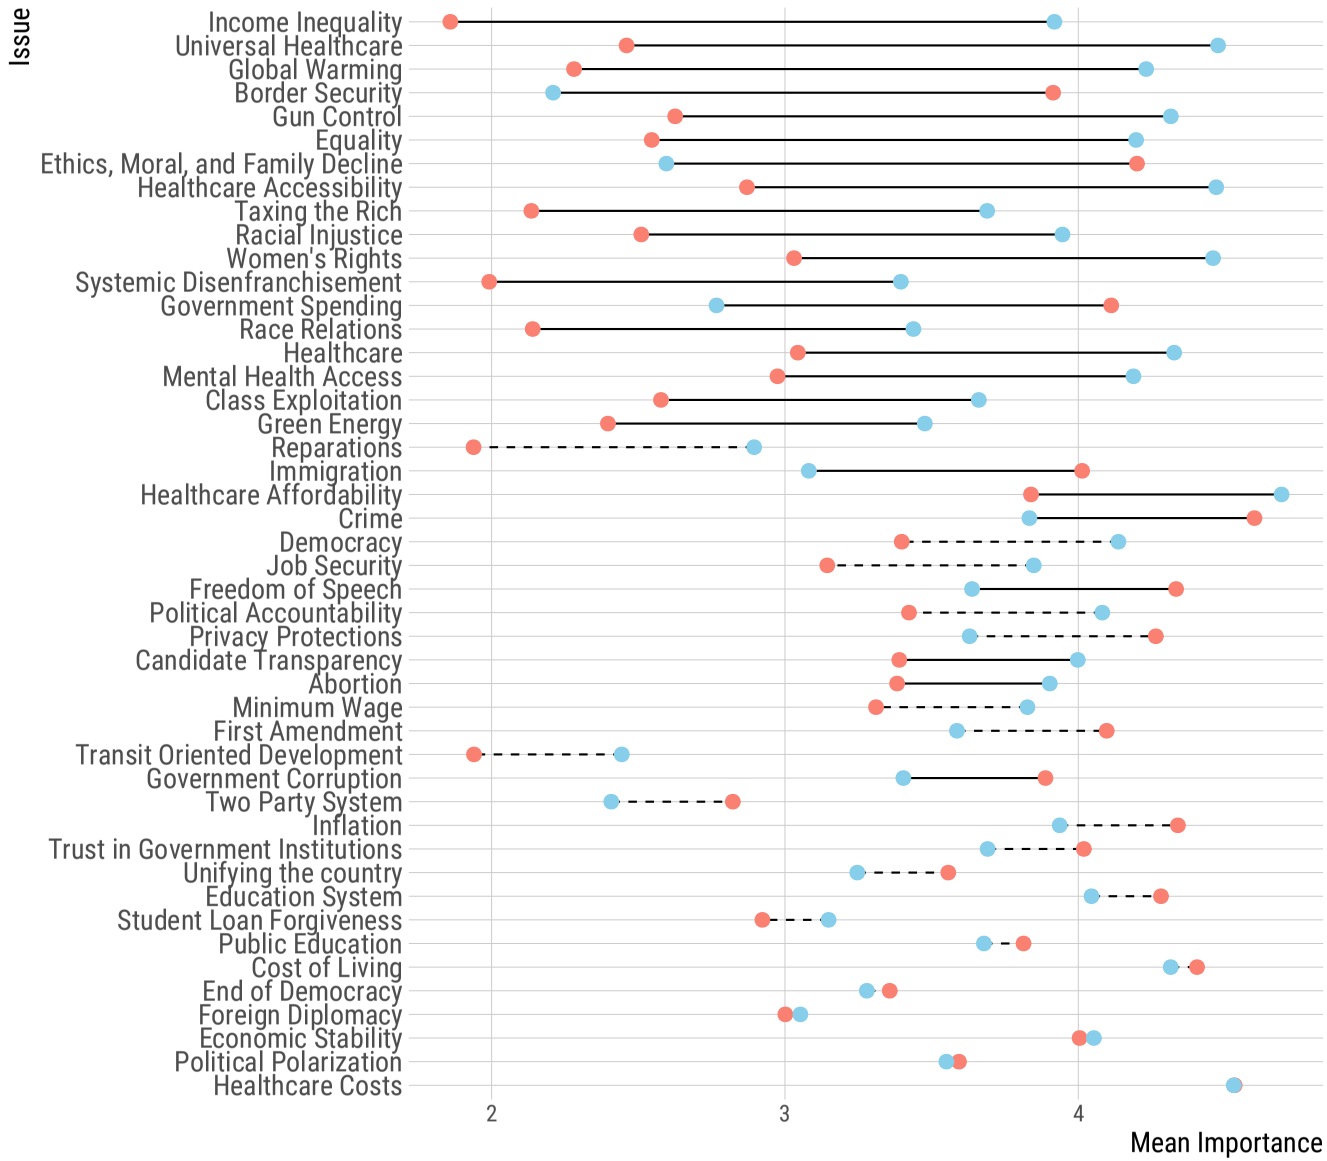

Figure 4: IPW-weighted estimates of survey questions measuring issue importance across Republicans and Democrats with corresponding 95% confidence intervals. Statistically significant (insignificant) differences at conventional levels of significance (α = .05; two-tailed) are depicted using solid (dashed) lines.

meaningful information from text suggests that more thought and care may need to go into identifying optimal open-ended formats. Here, lessons from scholars using semi-structured and unstructured interview modes could prove useful (Leech, 2002).

6.3 Late arrivals

Given that later submissions to the question bank are not considered by as many respondents, there is a risk that the "true" best-performing question will not be identified. In experimental settings where the determination of optimal treatment allocation may be vital, such as the testing of pharmaceutical interventions, this risk would represent a significant limitation for the approach. However, it is worth highlighting the exploratory nature of the proposed method, and how it performs against the alternative of pre-determined fixed question batteries. In the absence of an adaptive survey design, the selected questions are – in principle – assumed to be the best-performing set. If one compares a fixed battery to a dynamic battery comprised of the same items, the dynamic battery, with its ability to adapt and explore new questions, can potentially uncover additional questions that a static design may ignore. Therefore, the risk of failing to identify useful survey items may be even higher with a static design. Indeed, though "seed questions" included in the Latino survey were drawn from Latino-oriented fact-checking sources (i.e., Telemundo's T-Verifica, Univision's El Detector) and drawn from Gallup in the issue importance study, user-submitted questions were generally rated as more accurate and important. This further highlights the inductive advantages of the method. One can include questions that they possess strong priors about, and allow those questions to be tested alongside new submissions. This creates an environment where established, theoretically grounded items are examined in tandem with potentially novel items. Items introduced toward the end of data collection might also be incorporated into future waves, leveraging the information gleaned from existing ratings.

6.4 Toxicity and Quality of Responses

The CSAS method thrives on user-generated questions, but this can also be a source of problematic content. Toxic and low-quality content, such as hate speech and gibberish, can pollute the question bank and compromise the validity of the approach. To address these concerns, several mitigation strategies can be implemented. Using moderation filters to remove potentially harmful content before it is presented to users can be useful, as can "quality filters." For example, Velez and Liu (2023) use GPT-3 to classify open-ended responses into low and high quality on the fly, and find that such a classification is able to distinguish between inattentive and attentive participants, on average. Ensuring that open-ended responses are legible, not harmful, and provide relevant information is crucial to maintaining the integrity of the CSAS method. This can be achieved through a combination of automated filtering and human oversight.[6]

[6]Upon publication, scripts for implementing toxicity, relevance, and quality filters will be publicly provided.

7 Conclusion

This paper introduced the Crowdsourced Adaptive Survey (CSAS) method, an approach to developing evolving survey question banks. Leveraging the capabilities of large language models, open-ended text was converted into structured survey questions, while adaptive algorithms identified best-performing questions from the question bank. The method's efficacy was demonstrated through an application to Latino information environments, uncovering both factual and partisan narratives. Moreover, the method was applied to issue importance, highlighting the ability for the method to recover niche issue positions that would otherwise be missed by traditional survey methods.

In a study of negative beliefs about parties and candidates held by Latinos, I found that the most highly rated claims were either popular stereotypes of the parties, true claims, or widely reported allegations. In contrast, claims that were rated lower on accuracy were often objectively false and reflected more blatant misinformation. An additional analysis of Latino subgroups defined by WhatsApp usage revealed similar accuracy ratings across a variety of claims. The most consistent difference between the two groups emerged for true, but politically incongenial, claims, such that participants who trusted WhatsApp less were more likely to rate these claims as accurate relative to those who trusted WhatsApp for information. These descriptive patterns suggest that information gaps across WhatsApp trust may be more pertinent with respect to more mainstream claims, rather than outright misinformation. However, research is needed to assess whether this holds more broadly.

Moving from the identification of claims within marginalized communities to the topic of issue importance, the CSAS method was used to develop a question bank comprised of issues provided by participants. While economic factors, healthcare, crime, and education were among the most highly-rated issues, more niche issues such as "candidate transparency" and "privacy protections" also received high ratings from the sample. A follow-up analysis comparing issue ratings across partisan subgroups revealed substantial differences in importance ratings, but also some areas of agreement on topics such as healthcare costs, public education funding, and inflation. A distinct advantage of the adaptive design is that the exploration of new items can yield a larger number of items than one would typically measure in a traditional static survey. By dynamically adjusting the question bank, CSAS can delve deeper into emerging topics and unexpected areas of agreement across diverse groups. Given heterogeneity in ratings across individuals, applying more sophisticated adaptive designs such as contextual adaptive experiments could improve estimates of politically contentious claims or issues (Offer-Westort, Rosenzweig, and Athey, 2022).

Despite the advantages of the CSAS method, scholars should be aware of potential limitations or drawbacks of using LLMs. Simpler classification tasks, like those demonstrated in the work on important issue identification (Mellon et al., 2022), may not require the additional investment of fine-tuning pre-trained LLMs, as their existing capabilities may suffice. However, capturing more abstract concepts such as misinformation may require additional steps such as fine-tuning and few-shot learning (i.e., including positive and negative examples in an LLM prompt). For more sensitive

topics, implementing filters is also strongly recommended, given that participants can submit toxic content that may be seen by other participants.

Future research employing CSAS has the potential to explore a variety of topics where participant-generated content is particularly valuable. For instance, in candidate choice, voters not only consider issue positions, personal attributes, and party affiliation, but also other, less tangible factors. CSAS can be useful in uncovering these additional elements that influence voting decisions. Furthermore, when identifying norms, core beliefs, or key sources of identity within hard-to-reach communities, CSAS could provide significant insights, preventing the misalignment of researcher-defined concepts with respondents' actual perceptions. CSAS could also be used to develop measurement scales that more accurately reflect the diverse considerations of participants. This kind of research, which prioritizes the perspectives and realities of study populations, not only bridges the gap between researchers and respondents but could also produce significant advances in our understanding of public opinion and political behavior, more broadly.

References

Agrawal, S. and N. Goyal (2017). "Near-optimal regret bounds for thompson sampling". In: *Journal of the ACM (JACM)* 64.5, pp. 1–24.

Ansolabehere, S. and B. Schaffner (2009). "Guide to the 2008 cooperative congressional election survey". In: *Harvard University, draft of February* 9.

Carrasco-Farré, C. (2022). "The fingerprints of misinformation: how deceptive content differs from reliable sources in terms of cognitive effort and appeal to emotions". In: *Humanities and Social Sciences Communications* 9.1, pp. 1–18.

Cortina, J. and B. Rottinghaus (2022). "Conspiratorial thinking in the Latino community on the 2020 election". In: *Research & Politics* 9.1, p. 20531680221083535.

Dillman, D. A. and L. M. Christian (2005). "Survey mode as a source of instability in responses across surveys". In: *Field methods* 17.1, pp. 30–52.

Dillman, D. A., J. D. Smyth, and L. M. Christian (2014). *Internet, phone, mail, and mixed-mode surveys: The tailored design method*. John Wiley & Sons.

Dunaway, J., R. P. Branton, and M. A. Abrajano (2010). "Agenda setting, public opinion, and the issue of immigration reform". In: *Social Science Quarterly* 91.2, pp. 359–378.

Egami, N. et al. (2018). "How to make causal inferences using texts". In: *arXiv preprint arXiv:1802.02163*.

Elliott, K. J. (2020). "Democracy's Pin Factory: Issue Specialization, the Division of Cognitive Labor, and Epistemic Performance". In: *American Journal of Political Science* 64.2, pp. 385–397.

Jennings, W. and C. Wlezien (2011). "Distinguishing between most important problems and issues?" In: *Public Opinion Quarterly* 75.3, pp. 545–555.

Key, V. O. (1961). "Public opinion and the decay of democracy". In: *The Virginia Quarterly Review* 37.4, pp. 481–494.

Krosnick, J. A. and M. K. Berent (1993). "Comparisons of party identification and policy preferences: The impact of survey question format". In: *American Journal of Political Science*, pp. 941–964.

Lee, A. Y., R. C. Moore, and J. T. Hancock (2023). "Designing misinformation interventions for all: Perspectives from AAPI, Black, Latino, and Native American community leaders on misinformation educational efforts". In: *Harvard Kennedy School Misinformation Review*.

Leech, B. L. (2002). "Interview methods in political science". In: *PS-WASHINGTON-* 35.4, pp. 663–664.

Lopez, J. and D. S. Hillygus (Mar. 2018). "Why So Serious?: Survey Trolls and Misinformation". In: Available at SSRN: https://ssrn.com/abstract=3131087 or http://dx.doi.org/10.2139/ssrn.3131087.

McCombs, M. and J.-H. Zhu (1995). "Capacity, diversity, and volatility of the public agenda: Trends from 1954 to 1994". In: *Public Opinion Quarterly* 59.4, pp. 495–525.

Mellon, J. et al. (2022). "Does GPT-3 know what the Most Important Issue is? Using Large Language Models to Code Open-Text Social Survey Responses At Scale". In: *Using Large Language Models to Code Open-Text Social Survey Responses At Scale (December 22, 2022)*.

Mena, P. (2019). "Principles and boundaries of fact-checking: Journalists' perceptions". In: *Journalism practice* 13.6, pp. 657–672.

Montgomery, J. M. and J. Cutler (2013). "Computerized adaptive testing for public opinion surveys". In: *Political Analysis* 21.2, pp. 172–192.

Montgomery, J. M. and E. L. Rossiter (2020). "So many questions, so little time: Integrating adaptive inventories into public opinion research". In: *Journal of Survey Statistics and Methodology* 8.4, pp. 667–690.

— (2022). *Adaptive Inventories: A Practical Guide for Applied Researchers*. Cambridge University Press.

Offer-Westort, M., A. Coppock, and D. P. Green (2021). "Adaptive experimental design: Prospects and applications in political science". In: *American Journal of Political Science* 65.4, pp. 826–844.

Offer-Westort, M., L. R. Rosenzweig, and S. Athey (2022). "Battling the CoronavirusInfodemic'Among Social Media Users in Africa". In: *arXiv preprint arXiv:2212.13638*.

Ostrom Jr, C. W. et al. (2018). "Polls and elections: Terror, war, and the economy in George W. Bush's approval ratings: The importance of salience in presidential approval". In: *Presidential Studies Quarterly* 48.2, pp. 318–341.

Page, B. I. and R. Y. Shapiro (2010). *The rational public: Fifty years of trends in Americans' policy preferences*. University of Chicago Press.

Rheault, L. and C. Cochrane (2020). "Word embeddings for the analysis of ideological placement in parliamentary corpora". In: *Political Analysis* 28.1, pp. 112–133.

Ryan, T. J. and J. A. Ehlinger (2023). "Issue Publics: How Electoral Constituencies Hide in Plain Sight". In: *Elements in Political Psychology*.

Salganik, M. J. and K. E. Levy (2015). "Wiki surveys: Open and quantifiable social data collection". In: *PloS one* 10.5, e0123483.

Sanderson, Z. et al. (2021). "Twitter flagged Donald Trump's tweets with election misinformation: They continued to spread both on and off the platform". In: *Harvard Kennedy School Misinformation Review*.

Vaswani, A. et al. (2017). "Attention is all you need". In: *Advances in neural information processing systems* 30.

Velez, Y. and P. Liu (2023). "Confronting Core Issues: A Critical Test of Attitude Polarization". In.

Velez, Y. R., E. Porter, and T. J. Wood (2023). "Latino-targeted misinformation and the power of factual corrections". In: *The Journal of Politics* 85.2, pp. 789–794.

Velez, Y. R. (2023). "Trade-offs in Latino Politics: Exploring the Role of Deeply-Held Issue Positions using a Dynamic Tailored Conjoint Method". In: *Aletheia*. URL: https://osf.io/g5rwk/.

Vosoughi, S., D. Roy, and S. Aral (2018). "The spread of true and false news online". In: *science* 359.6380, pp. 1146–1151.

Weller, S. C. (1998). "Structured interviewing and questionnaire construction". In: *Handbook of methods in cultural anthropology*, pp. 365–409.

Wlezien, C. (2005). "On the salience of political issues: The problem with 'most important problem'". In: *Electoral studies* 24.4, pp. 555–579.