Received: 25 January 2023 | Accepted: 17 July 2023

DOI: 10.1002/aaai.12103

ARTICLE

Civicbase: An open-source platform for deploying Quadratic Voting for Survey Research

Madeline E. Bassetti[1] | **Gustavo Dias**[1,2] | **Daniel L. Chen**[1] | **Alan Mortoni**[1] | **Ritesh Das**[1]

[1]Data and Evidence for Justice Reform, The World Bank, Washington, District of Columbia, USA

[2]The University of Chicago, Chicago, Illinois, USA

Correspondence
Madeline E. Bassetti, Development Impact Evaluation (DIME) group, The World Bank, Washington, DC, USA.
Email: madelineb@protonmail.com

Abstract

Civic engagement is increasingly becoming digital. The ubiquity of computing increases our technologically mediated interactions. Governments have instated various digitization efforts to harness these new facets of virtual life. What remains to be seen is if citizen political opinion, which can inform the inception and effectiveness of public policy, is being accurately captured. Civicbase is an open-source online platform that supports the application of Quadratic Voting Survey for Research (QVSR), a novel survey method. In this paper, we explore QVSR as an effective method for eliciting policy preferences, optimal survey design for prediction, Civicbase's functionalities and technology stack, and Personal AI, an emerging domain, and its relevance to modeling individual political preferences.

INTRODUCTION

Survey research has traditionally been used to uncover preferences. When it comes to policy preferences, in particular, the goal of the survey method is to understand the preferences of not only individuals, but also groups, which informs decision-makers of policies that best reflect the preferences of a community. The survey methods available in traditional platforms do an adequate job when it comes to informing surveyors of whether respondents agree or disagree with a particular policy. This is the case of Likert Scale or Conjoint Analysis, surveys widely used in empirical re-search. These techniques, however, fall short when it comes to measuring the preference intensity of survey respondents. This becomes an issue for behavior prediction models since the policy preferences held more intensely are the ones more likely to decisively influence a person's action. If we are interested in predicting the behavior of individuals, the intensity of preferences is a point of particular emphasis. Until now, measuring preference intensity has been a difficult task for field researchers (Cavaillťe, Chen, and Van Der Straeten 2022).

Civicbase fulfills this demand by offering researchers traditional survey methods plus Quadratic Voting for Survey Research (QVSR). QVSR, the current focus of the Civicbase project, is a new survey method created in response to E. Glen Weyl and Eric Posner's book Radical Markets: Uprooting Capitalism and Democracy for a Just Society (Posner 2018). In the book Weyl and Posner describe Quadratic Voting as a system, which enables people to cast votes that reflect the intensity of their preferences, not just the direction.

In the following section, we will detail the survey methodologies available on the platform. In the third section of this paper, we discuss the optimization of survey methods and recent evidence that suggests QVSR is the

Madeline E. Bassetti and Gustavo Dias contributed equally to this work.

survey design with the most predictive power relative to Likert and Conjoint. In the fourth section, we detail the features and capabilities of Civicbase. In the fifth section, we comment on the future of prediction models and the potential application of personal AI models in collective decision-making and politics.

SURVEY METHODOLOGIES

Researchers can use Civicbase to create surveys with three different methodologies: QVSR, Likert, and Conjoint analysis. In the next pages, we will review each of these methods in more detail.

Likert Scale

The Likert Scale, developed by Rensis Likert, is a psychometric survey instrument widely used in survey research today (Likert 1932). The most common version asks respondents to rate on a scale of 1–5 (or 7) ordered responses ranging from "strongly agree" to "strongly disagree." The Likert Scale does not place any constraints on the elicitation of preferences. This allows survey participants to take extreme positions on the scale across many questions. It has been demonstrated that a characteristic W distribution with clustering at the extremes prevails in large studies that investigate responses to important public policy issues (Quarfoot et al. 2017; Caviallťe, Chen, and van Der Straeten 2019). Respondents tend to shout their views, either very strongly in favor or against (Posner 2018). This is problematic for researchers when trying to justify respondent's preferences as "true" and only shows the direction of an individual's preferences, not their intensity (Caviallťe, Chen, and van Der Straeten 2019).

For this reason, Likert Scale responses have already been described as cheap talk. Cheap talk creates an "abundance problem" due to the lack of restrictions on voting across potentially competing policy issues. In other words, the survey response provides limited information on the relative priorities of the respondents when faced with a forced choice (e.g., in case of lockdowns during a pandemic, public health at the expense of economic activity or vice versa). A second concern is the "bunching problem": some respondents may face conflicting motives when answering survey questions, caring both about reporting their "sincere" view on an issue, but also willing to appear as loyal to their preferred group. In a highly polarized world, the distortion introduced by partisan motives may cause survey answers to cluster on both extremities, with respondents tending to "shout" their views. Thus respondents with the same preference orientation but varying preference intensities end up getting clustered together. A simple and parsimonious decision-theoretic model of survey answers, which includes both the sincerity motive and other competing motives (Caviallťe, Chen, and Van Der Straeten 2022), illustrates this challenge. Respondents are likely to systematically misreport their views when they are measured using Likert items. This results in similar responses from those who care about the issue and those who do not care as intensely and are merely paying lip service to a group norm (Akerlof and Kranton 2000).

In order to improve measures of preference intensity, Likert survey items are often followed by a question asking how important a given issue is to the respondent. Responses are collected using a similar categorical scale ranging from "not at all important" (1) to "very important" (5) (Miller and Peterson 2004; Howe and Krosnick 2017). Though this gives us some insights on their preference intensity, it does not explicitly address the above-mentioned abundance and bunching problems that arise from cheap talk.

Conjoint analysis

Conjoint analysis is another popular survey technique for determining how individuals make decisions when faced with a trade-off. First proposed by Green and Rao (Green and Rao 1971), Conjoint analysis can take up three forms: ranking, rating, or choice-based conjoint. Civicbase features the choice-based version, in which participants are presented with a random pairwise combination of options (i.e. products, policies) with varying attributes. For each new pair, the respondent must select the option they prefer. After an arbitrary number of selections is made, the survey ends and researchers can infer the respondent's relative preferences. This survey method is especially useful for evaluating marginal changes in what would otherwise be very similar options. With the sample's relative preferences collected, surveyors can create a rank of the options for each respondent, respecting the assumptions of completeness and transitivity. Each position in the rank is assigned points that are consistent with the preference ordering (options higher in the rank receive a higher value than lower-ranked options). But other than this limitation, points are assigned arbitrarily. Next, the information in the ranks of each participant is aggregated by adding all the points assigned to each option by each participant.

Conjoint analysis is a useful instrument for surveyors in capturing a participant's order of preferences. Like Likert, however, Conjoint fails when it comes to capturing individuals' preference intensity, concealing a better estimate

of the utility that an individual derives from each option. It is also a less realistic survey instrument as people do not necessarily consider all the specific trade-offs chosen by the surveyors whenever they need to decide between the available options.

Quadratic voting for survey research

A recent innovation in survey technology, QVSR brings the power of Quadratic Voting (QV) to the context of survey research. Before we discuss QVSR, a review of QV would be helpful. QV is a collective decision-making mechanism that can capture the intensity of people's preferences using a budget constraint, real or artificial (Lalley and Weyl 2018). This allows voters to "buy" votes and spend them on policy proposals. As a voter expresses more intensity for a policy, this becomes more expensive, represented using a quadratic cost function. "What's important is not so much the total cost of each number of votes, but that the marginal cost of casting the next vote grows proportionally to the number of votes cast" (Posner 2018). This design will incentivize voters to think critically and relatively about survey items with respect to preference intensity.

Crucially, QV addresses not just the direction of one's preferences but their strength. It is posited as a remedy to the "Tyranny of Majority," popularized in John Stuart Mill's "On Liberty" (Mill 1859). On the surface, the 1 person, 1 vote (1P1V) framework seems equitable as it allows voters to express the same unit of influence across the population on issues important to them. However, it does not allow voters to signal the intensity of their preferences or increase influence on certain issues that are relatively more important to them (Quarfoot et al. 2017). These shortcomings of 1P1V can lead to tyranny of the majority outcomes (Posner 2018). That is, situations in which policy decisions are suboptimal because the majority of the population is less informed or has less stake in this decision than a small minority. QV addresses this by allowing individuals to express truer preference intensity by casting more votes for particular issues and less for others. In theory, this allows a minority group to exert a stronger influence over the outcome of an issue that may materially or personally affect them more.

Furthermore, by following the same design as QV, QVSR can also address bunching problems in opinion polls. Because the price for each vote is quadratic, it becomes increasingly costly to acquire additional votes to express support or opposition to the same issue. The quadratic pricing makes it costly to express intense preferences by voting repeatedly for the same issue (Lalley and Weyl 2018).

OPTIMAL SURVEY DESIGN FOR PREDICTION

As we have argued in the previous section, there are many reasons why, in theory, one would have better results if one chose to conduct a survey using the QVSR design. But would capturing preference intensity actually have a significant impact on the performance of political behavior prediction models? If we compared responses to surveys with the actual behavior of the respondents, would QVSR results reflect their behavior more accurately than the results of other available survey designs?

This is the question that Cavaillé, Chen, Das, and Van Der Straeten set out to answer in "Willingness to Say? Optimal Survey Design for Prediction" (mimeo). The researchers selected three different survey methods to compare: Likert, Likert items with a follow-up issue importance question ("Likert+" for short), and QVSR. A nationally representative sample of roughly 4000 individuals were randomly assigned to take the same survey varying only the technology used to measure their policy preferences. The survey had respondents answer whether they favored or opposed the following 10 policies:

1. Giving same-sex couples the legal right to adopt a child.
2. Laws making it more difficult for people to buy a gun.
3. Building a wall in the US Border with Mexico.
4. Requiring employers to offer paid leave to parents of new children.
5. Preferential hiring and promotion of blacks to address past discrimination.
6. Requiring employers to pay women and men the same amount for the same work.
7. Raising the minimum wage to 15 dollars an hour over the next 6 years.
8. A nationwide ban on abortion with only very limited exceptions.
9. A spending cap that prevents the federal government from spending more than it takes.
10. The government-regulating business to protect the environment.

After collecting the policy preferences of each respondent, the researchers introduced participants to three behavioral tasks. Each of these tasks tried to measure a particular political behavior and was designed so that specific survey questions were expected to correlate with the behavior measured by the task. For instance, the first assignment tried to measure how much a participant is willing to engage in political donations and it was restricted to the topics of immigration and gun control.

For this experiment, answers to survey questions 2 and 3 should correlate with the participants' behavior. Participants were told that they had been automatically entered into a lottery with a prize of 100 dollars for randomly selected respondents. Next, they were prompted to imagine that they had won the lottery and asked whether they wanted to donate part of their lottery money to an advocacy group. They were told that they had a choice between four advocacy groups working on the two issue areas: immigration and gun control. For each issue, organizations fell on different sides of the political divide: for and against immigration, as well as for and against gun control. Respondents could choose to donate to one of the four advocacy groups, or to not donate at all. Whatever they did not donate, they could keep.

The second task measured how respondents behaved in a situation of power. Participants were placed in a dictator game in which they were faced with a clear trade-off: reward or punish another participant, identified as an independent, who agrees with them on one issue but disagrees with them on another. Respondents were asked how they would behave in three dictator games involving a Republican, a Democrat, and an Independent. Next, the game followed a setup very similar to that of the donation scenario. Respondents had the option to donate anywhere between 0 and 100 dollars of some lottery. After they made their decisions, participants were informed of the Independent respondent's preferences on two issues, against the border wall and against gun control. Then, respondents were asked if they wanted to change the amount they had previously decided to donate to this individual. This specific mix of preferences (antigun control and pro-immigration) was chosen to correlate with survey questions 2 and 3 again. This choice was also interesting because it usually involved a clear trade-off for the participants. For example, if you were a Democrat, choosing to punish the Independent for donating to a pro-gun organization meant also punishing her for donating to a pro-immigration organization.

A third and final task was designed to capture "the willingness to spend time and effort promoting a political cause one agrees with," and it was restricted to the topics of abortion rights and minimum wage. In this assignment, answers to survey questions 7 and 8 were expected to correlate with the participants' behavior. Respondents were given the option to write a letter about one of two policy proposals being discussed in Congress at the time. One proposal sought to restrict abortion rights and the other sought to increase the minimum wage.

From the outcomes of the three behavioral tasks, the demographic features of participants, their survey responses, and their treatment assignment, the researchers trained different machine learning models to predict subjects' behavior. Their goal was to assign individuals to survey methodologies solely based on the covariates in order to be able to predict the behavioral outcomes more accurately. The machine learning methods used for this were XGBoost, random forest, logistic regression, decision tree, support vector classifier, ADA boost, K-Nearest Neighbor, Gaussian Naive Bayes and Elastic Net. To identify the conditional expected treatment impacts of each survey methodology, the study resorted to a two-step process. First, they chose the machine learning method which had the least mean squared error. In this case, the method was XGBoost. From this optimal machine learning model, they calculated the Brier score. Next, the Brier score was used as the outcome variable of a policy learning algorithm (Athey and Wager 2017). Given a doubly robust estimator of the causal effect of assigning everyone to treatment, the policy learning algorithm uses the covariates, and the treatment assignment information to map the covariates to a particular treatment arm that resulted in the least Brier score. Since the policy tree assigns the treatment arm based on the maximum predicted treatment impact, the chosen treatment arm for the covariates is the one that yields the best prediction of behavior given the survey responses and the covariates. And in fact, QVSR was chosen more often as the treatment arm that results in the least Brier score when compared to Likert and Likert+. This evidence suggests that QVSR could indeed be a survey design, which yields data best predictive of individual behavior such as donations, voter turnout, and letter writing compared to the alternatives.

As part of their study, the researchers also compared Likert, Likert+, and QVSR with a fourth mechanism, a variation of QVSR, which they named QVSR with monetary numeraire ("QVSRN" for short). QVSRN works the same way as QVSR, except that it allows respondents to convert their credits into a currency that they can use in future surveys. That is, it enables respondents not only to express their variation in preference intensity within the topics of a survey, but also their relative preferences between different surveys. When it comes to predicting political behavior such as the ones measured by the behavioral tasks, the study found that QVSRN performs even better than QVSR. These findings are still being analyzed, but they already prompt the question: what other variations on these survey designs could bring us closer to an optimal survey design for prediction? As we have stated in the previous section, Civicbase can only host surveys with Likert, Conjoint, and QVSR methods. But there is already a list of promising new variations on QVSR that can be tested and added to the platform such as QVSRN, and QVSR with linear, cubic, or exponential cost functions. In the

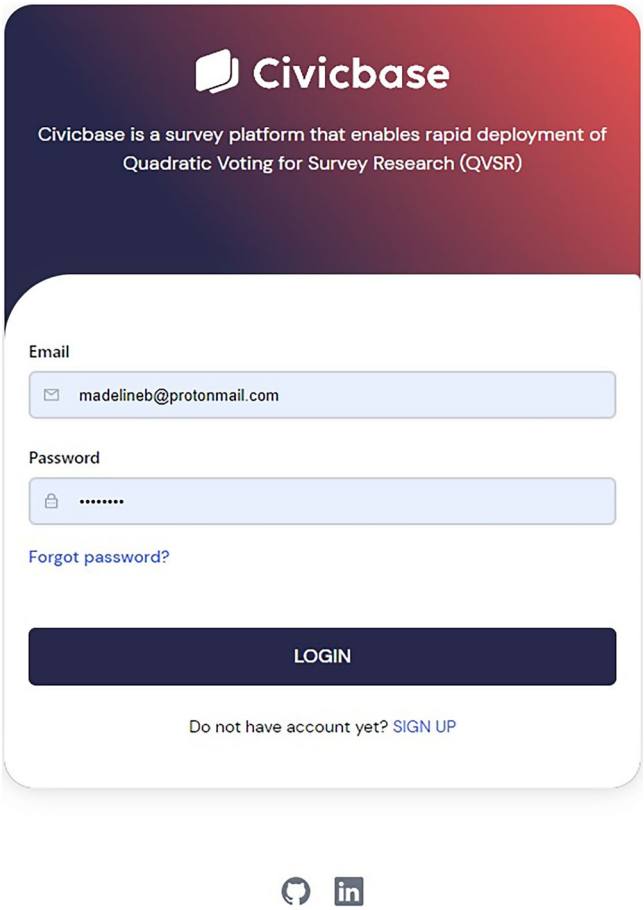

FIGURE 1 Landing page with login screen for Civicbase.

next section, we will dive into the current functionalities of Civicbase.

CIVICBASE

Civicbase offers the deployment of QVSR surveys together with traditional survey methods. Below is an overview of the functionalities of the software.

Landing page, registration, and login

Civicbase allows for account creation via a registration form on the landing page (see Figure 1 on page 8.) The user must fill in a form with their name, preferred email, and password. Then, an email communication will be sent to the address provided by the user. The message will request the user to confirm that this is the correct email and redirect them to the Civicbase survey creation dashboard. After signing up and making the first login, all future login can be done in Civicbase's landing page by only providing the correct email and password associated with the account. Civicbase does not currently charge for account creation and hosting, therefore users are able to create and deploy surveys in a seamless manner.

Survey creation

Once logged into Civicbase, researchers can create surveys using Quadratic Voting for Survey Research (QVSR), Conjoint, and Likert survey methodologies (see Figure 2 on page.) Surveys can be customized using field inputs in the creation process. A drop-down menu allows the surveyor to choose between the different survey methodologies. A customization section lets the surveyor add field inputs for respondents to provide feedback at the end of the survey. Additionally, it is possible to configure surveys so that respondents can take them more than once, require participants to identify themselves and radomize questions presented to respondants.

Surveys created with Civicbase can also have the option of an audio recorder embedded into the survey page. This helps researchers doing field work to record the exact interactions they had with survey respondents. Another design feature of the app that addresses a typical complication of field research is GPS tracking functionality. With it, researchers can capture the GPS coordinates of the surveyor.

If a survey created with Civicbase uses the QVSR methodology, then surveyors can choose between four suggested vocabulary options for respondents to convey their opinion about a survey question. The options are "Disagree/Agree," "Favor/Oppose," "Approve/Reject," and "Aye/Nay." If none of these accurately capture the survey's intention, then there also exists a fifth, customizable option. Surveyors can write the terms they believe to best suit the context of the questions or statements. The platform also gives surveyors three suggested vocabulary options to refer to the voting credits: "tokens," "credits," or "coins." Beyond these, there is also a fourth, customizable vocabulary option. Surveyors can establish their own vocabulary for this component of QVSR.

After it is created, the survey will then be saved to the researcher's dashboard in pilot mode. Surveys can then be edited and further customized based on the researcher's needs. Civicbase works offline and surveyors are able to record and store audio, capture GPS coordinates, load surveys, and capture survey inputs from respondents. This allows researchers not to be constrained by internet accessibility or connection quality when choosing their survey sample.

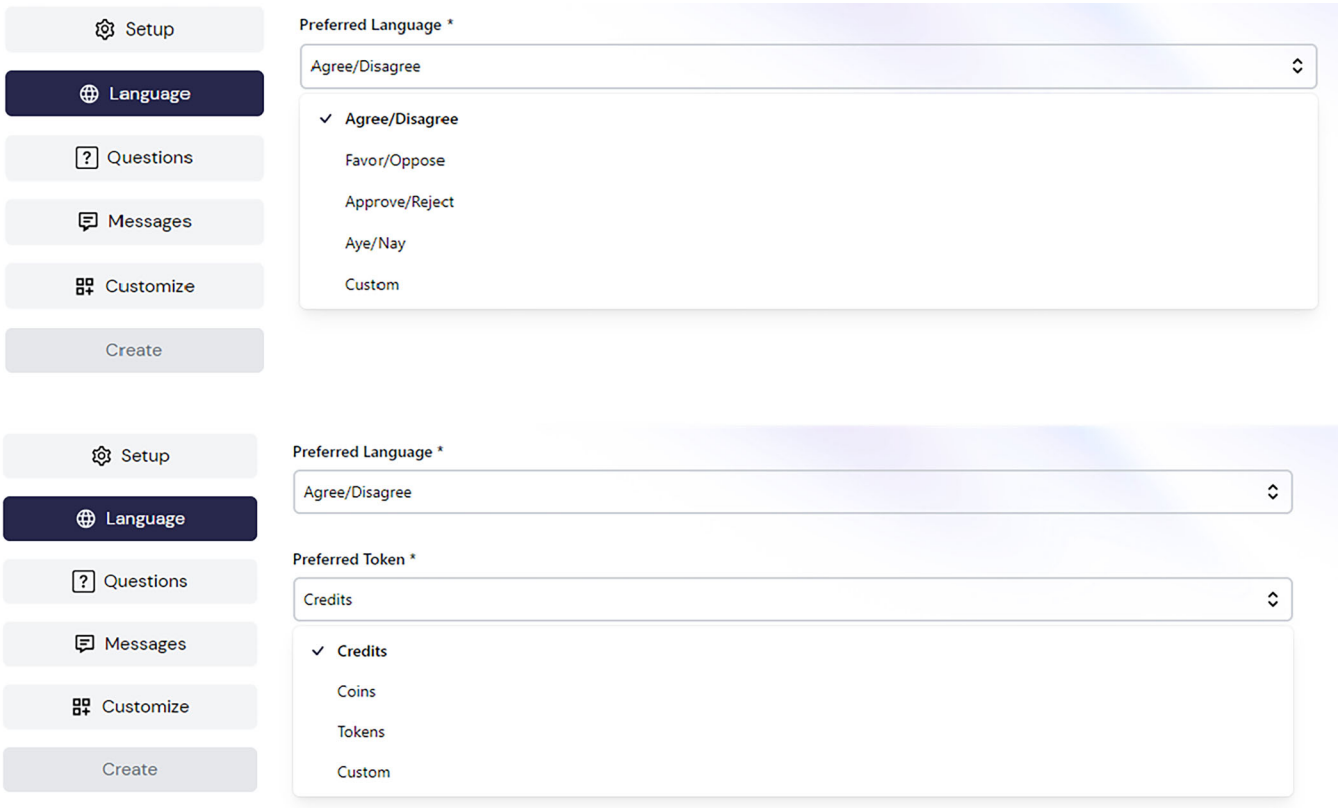

FIGURE 2 Survey creation flow for Civicbase.

Pilot and published modes

Once a survey is created and drafted, it will be saved in the pilot state by default. In the pilot state, a survey URL is generated that can be shared with pilot survey participants. This allows researchers to effectively test their survey before publishing it for a live experiment. In the pilot state, the survey can still be edited, even when the survey link has been distributed and participation is active. This allows researchers to tinker with their survey design, which can minimize errors with live deployment. Pilot surveys collect data on respondents and this can be viewed within the Analytics page. Published mode represents live survey deployment. A new answer set is generated with published surveys that distinguishes it from Pilot mode. Published surveys collect data on survey respondents as well and this data can be viewed on the Analytics page.

Survey respondent view

The survey respondent view displays the survey interface with which respondents interact and record their answers. The Conjoint and Likert pages show a clean representation of each method. QVSR displays a credit bar at the top of the page which moves dynamically based on the respondent's allocations. Below the credit bar, we see the question/statement set. The most important interactive element of the respondent page is the credit allocation buttons. This design facilitates the QVSR survey method. It allows the survey participant to agree or disagree within a scale of 0–7 for each question, while observing the credit bank reduce or replenish quadratically as the user allocates stronger preferences for certain statements or questions. A radial circle design is used to graphically communicate the intensity of preferences with 0–7 bi-directional movement based on agreement or disagreement. The colors green (agreement) and red (disagreement) are used to visually reinforce these allocations. Furthermore, total credits allocated for each question is displayed to the respondent so that they can see how many credits they have allocated. This is especially useful if the respondent decides to tinker with their allocation to ensure that they have allocated all of their budget. There is current consideration to improve the QVSR experience, by displaying the marginal cost of each added vote. Finally, the respondent is able to submit their responses and is taken to a confirmation page, that is customizable in its text.

Analytics

The Analytics page allows researchers to better understand how their survey respondents are interacting with their experiment. It provides two interactive tables, "Results" and "Insights." The Results table shows the Question ID or survey item, aggregation of responses for "Agree" and "Disagree" for a particular question. This table assists the researcher in understanding how respondents are voting in real time as the survey is being conducted and at the end of the survey when analyzing the data.

The Insights table displays an overview of the survey's performance. It shows the total number of respondents, total number of times the survey was accessed, a conversion rate (total number of respondents divided by the total access) and it shows the state of the survey, pilot or published. All data in the analytics tables can be downloaded as a CSV file for further analysis. The CSV file contains more data points useful for analysis, including the time the survey was started and completed, total number of credits left over, how questions were randomized, and additional fields that can be collected through the Civicbase API, if interfacing with another website.

We intend to expand the functionality of this feature with data visualization, such as pie charts, histograms, and barplots containing the survey information. (Figure 3)

Technology stack

Civicbase uses a set of modern libraries and frameworks, these include: Tailwind CSS https://tailwindcss.com/; React JS https://reactjs.org/; Draft JS https://draftjs.org/; firebase for hosting; firebase functions to write/read from database; firebase analytics to get a minor insights from our users.

Charts and plots in the data visualizations page were created using Nivo, a package of dataviz built on top of the d3 and React libraries. An example of visualization we provide is a regional infographic built with visitors IP addresses. It is also possible to view the average time the respondents are taking per question. A graph with the time difference per answer is created. This gives researchers more insight into how much time respondents are spending on questions. All of these visualizations update in real time, allowing the researchers to view how modifications to their survey may alter these data.

Civicbase is now an open-source Github repository. In the flow chart on page 14, we can observe a high-level abstraction of Civicbase's architecture. At the top, we have the main technology behind our application. In the middle portion of the chart, we have the front end and how the pages are structured. Finally, at the bottom, we have the firebase functions. Functions in orange are the ones that require authentication.

ETHICAL CONSIDERATIONS AND LIMITATIONS

Survey research and preference elicitation, especially in regards to citizenry voting can crossover into the domain of sensitive data. Special attention should be paid to ensuring the security of the data, and its use. Civicbase is open-source, therefore its codebase can be continuously and rigorously checked for accuracy and integrity. However, because open-source code can be forked and implemented by anyone, within the grounds of a licence, we cannot fully control how actors may use this codebase down the line. Furthermore, ubiquitous computing makes is easier and lest costly for governments and private institutions seek to understand their citizens, employees, and consumers, QVSR would be difficult to implement in an analog nature, and is therefore suitable for the screen. As we posit it elicits truer preferences and sheds light on relative preferences, these insights could be potentially abused and exploited by authoritarian actors. It is intended by working with credible institutions to roll out Civicbase, the benefits of understanding truer preferences in democratic polities can assist authorities in strengthening civic feedback loops. It is hoped that this lead by example and research contribution approach can set better precedents of use for Civicbase and QVSR in real-world contexts.

COLLABORATIVE DEVELOPMENT

Civicbase is a collaborative project between the Data and Evidence for Justice Reform (DE JURE) program at the World Bank and volunteers from around the world. Currently, DE JURE is engaged with projects at various stages in India, Estonia, Australia, and Mozambique. These project scopes include comparative performance in the field of QVSR and Likert scale, application in municipal dashboards, predicting voter turnout through QVSR elicited political opinions, and willingness to pay for public goods. Civicbase can be applied in many municipal and public policy contexts. We encourage opportunities for research partnerships and collaboration. (Figure 4)

PERSONAL AI

Civic participation can be described as the extended involvement of individuals in a collective political

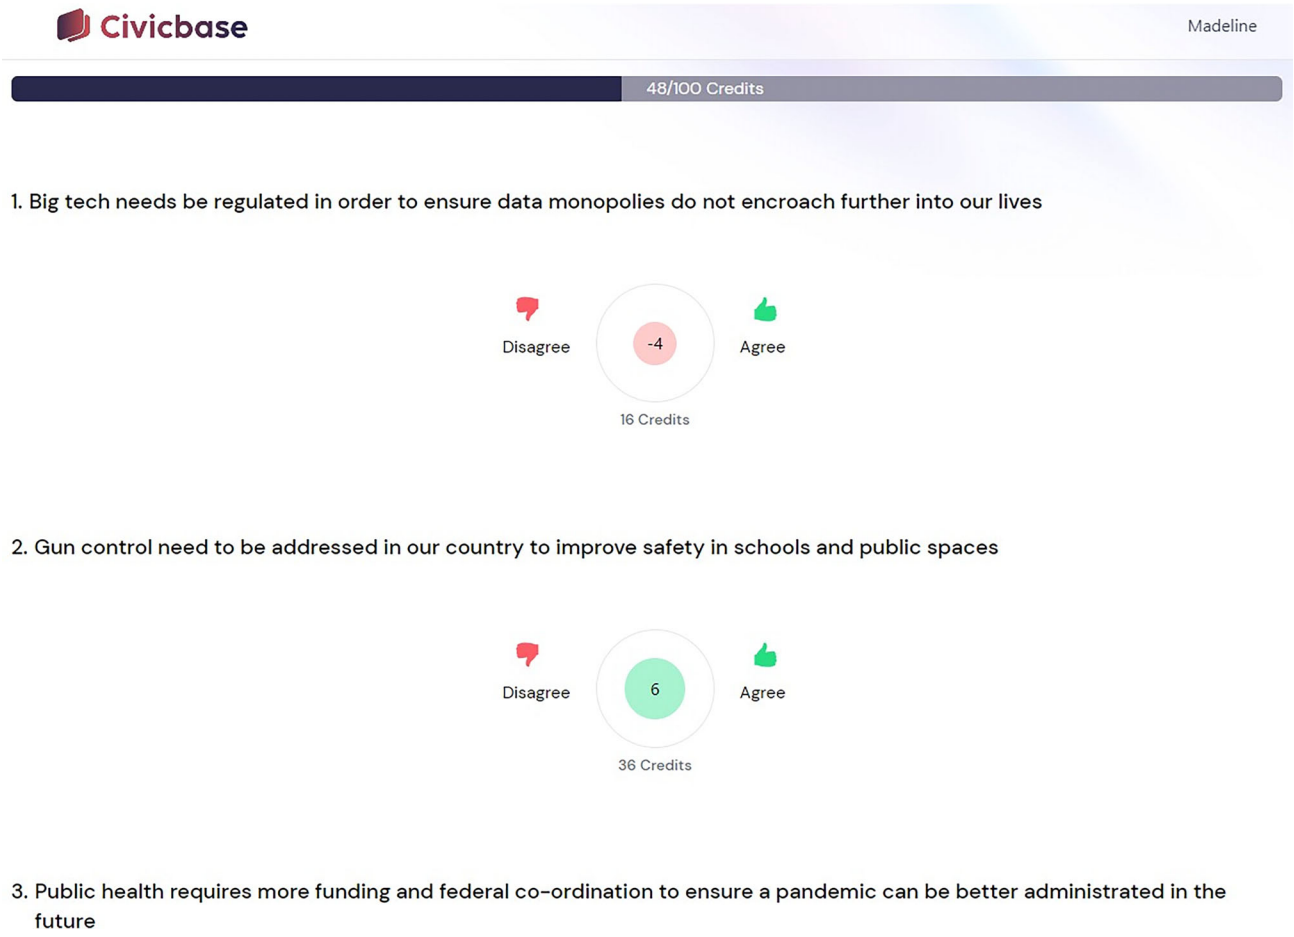

FIGURE 3 QVSR interface for survey response collection. QVSR, Quadratic Voting Survey for Research.

decision-making process ("Front Matter" 2019). Two prominent driving factors for citizen engagement include a rights-based view of citizenship, which includes voting, demonstrating, campaigning, joining political parties, and a "dutiful citizenship" motive, which involves being informed by news, and forming organizations ("Front Matter" 2019). We posit that while these motivations exist, they often compete with various concerns that encompass daily life for many citizens, including work commitments, family ties, and even household chores. Additionally, there is a cognitive load on citizens to stay across the many public policy issues that exist at a local, state, and federal level within their country.

The term "Personal AI" encompasses any data-driven model associated with an individual, and which continuously learns from that individual's knowledge and experiences. Theoretically, a personal AI would have a person's unique personality and would be able to interact with communities as a digital agent of that person. One area that researchers believe Personal AI and machine learning models will have a significant impact on is the development of automated decision-making systems. Governments and federal agencies are increasingly using data-driven models for improving decision-making (Engstrom and Ho 2020). In the context of civic engagement, Personal AI could be used to model an individual citizen's public policy preferences, and thus create a digital agent for each citizen in a community. We posit that QVSR could be an effective survey method for uncovering truer preferences, and therefore could be useful in uncovering public policy preferences. From these preferences, an AI model could theoretically predict a citizen's political preferences, and even make political decisions on behalf of its user. Careful ethical consideration will need to be applied, particularly in an acknowledgment of the "Ironies of Automation" in which it is propounded that as automation increases within a set of functions, it deskills the human agent over time as they no longer apply their learned expertise (Bainbridge 1983). We can imagine a scenario where a Personal AI who modeled a citizen's preferences accurately and votes on public policy on behalf of their citizen, may over time, create political apathy within the citizen.

A participatory framework for algorithmic governance called WeBuildAi has been developed and tested in the field (Lee et al. 2019). This framework allows people to participate in designing algorithmic governance policies,

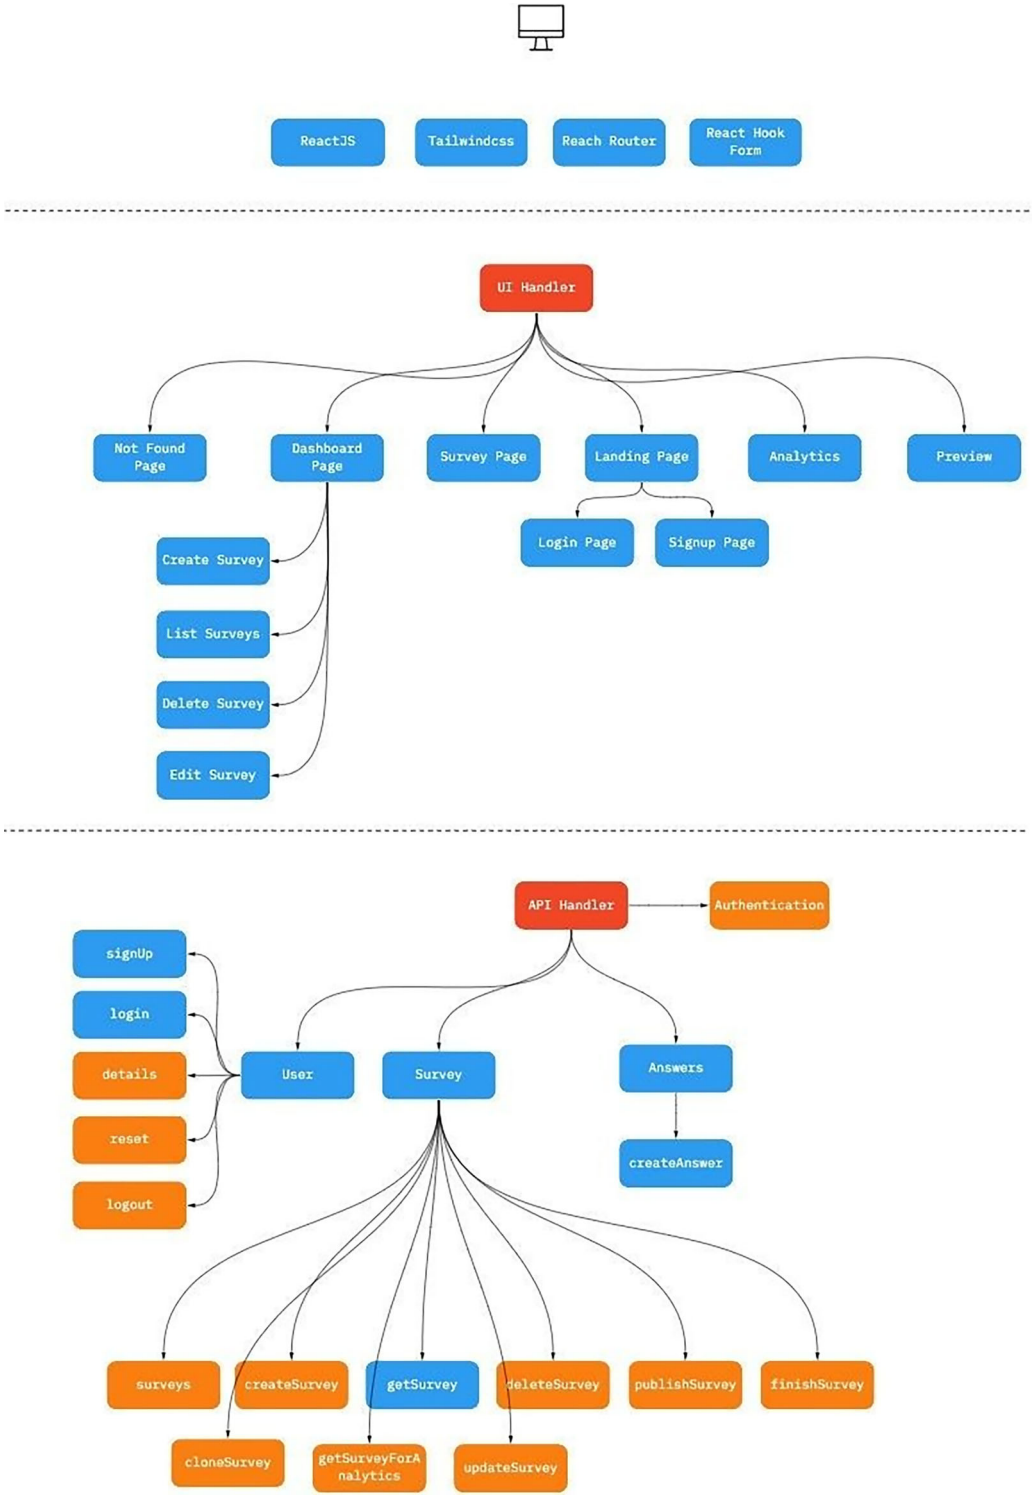

FIGURE 4 A high-level abstraction of Civicbase's software architecture.

with a key aspect being the facilitation of individuals in creating computational models that reflect their beliefs. These models then vote on the individual's behalf. It uses a three-stepped approach, feature selection, model building, and then model selection. Feature selection involves consultation with individuals about what features should be used by the algorithm to make decisions. Pairwise options were used to elicit responses at this stage. The next step is model building in which the researchers used a machine learning model and an explicit rule model to reflect an individual's decision criteria. The final step is model selection, where once people build their models using the model-building

methods, a visualization of the model and its example decisions are presented to the individual so they can better understand each model and select which one best reflects them (Lee et al. 2019). We would like to replicate the WeBuildAI experiment in a civic engagement context with a sample of citizens and with further investigation, explore ways QVSR may apply to elicit responses. Additionally, we also wish to explore how Civicbase could become a public repo for capturing government policy and political decision-making and subsequent citizen preferences on these decisions. This could serve as a foundation for a rich training dataset for personal AI agents.

CONCLUSION

The Civicbase platform aims to allow researchers to further explore the benefits of QV for preference elicitation. Application of a budget constraint over Likert's Scale offers an alternative survey tool for researchers. It holds promise when the goal of the study is to understand truer and relative preferences, control for "shouting" responses, or polarization. QVSR will require further stress testing in the wild, and it is hoped that this contribution will enable that. The long-term vision of Civicbase is a place where people can create communities for eliciting preferences using novel methods like QVSR, this could take the form of municipal dashboards, citizen portals, worker collectives, and domain hobbyists. Civicbase could also be particularly helpful for the development of personal AI projects focused on political representation (Suter et al. 2022). It will allow surveyors to improve the way they collect and aggregate preferences, and perhaps improve the accuracy and sensitivity of voting behavior predictive models. We intend to continue to grow the community of researchers and developers who use Civicbase by creating an open-source ecosystem (OSS) and will endeavor towards realizing collaborative potential.

CONFLICT OF INTEREST STATEMENT

The authors declare that there is no conflict.

ORCID

Madeline E. Bassetti https://orcid.org/0000-0003-1926-0657
Gustavo Dias https://orcid.org/0000-0002-7218-4537

REFERENCES

Akerlof, G. A., and R. E. Kranton. 2000. "Economics and Identity*." *The Quarterly Journal of Economics* 115(3): 715–53. https://doi.org/10.1162/003355300554881

Athey, S., and S. Wager. 2017. Policy learning with observational data. https://doi.org/10.48550/ARXIV.1702.02896

Bainbridge, L. 1983. "Ironies of automation." *Automatica* 19(6): 775–9. https://doi.org/10.1016/0005-1098(83)90046-8

Cavaillé, C., D. L. Chen, and K. Van Der Straeten. 2022. "Who Cares? Measuring Preference Intensity in a Polarized Environment." TSE Working paper no. 22-1297, Toulouse School of Economics (TSE). https://ideas.repec.org/p/tse/wpaper/126566.html

Cavaillé, C., D. L. Chen, and K. van Der Straeten. 2019. "A Decision-Theoretic Approach to Understanding Survey Response: Likert vs. Quadratic Voting for Attitudinal Research." *The University of Chicago Law Review* 87: 22–43.

Engstrom, D. F., and D. E. Ho. 2020. "Algorithmic Accountability in the Administrative State." http://hdl.handle.net/20.500.13051/8311

Front matter. 2019. In *The Playful Citizen: Civic Engagement in a Mediatized Culture*, 1–4. Amsterdam: Amsterdam University Press. Retrieved January 24, 2023, from http://www.jstor.org/stable/j.ctvcmxpds.1

Green, P. E., and V. R. Rao. 1971. "Conjoint Measurement for Quantifying Judgmental Data." *Journal of Marketing Research* 8(3): 355–63. Retrieved January 24, 2023, from http://www.jstor.org/stable/3149575

Howe, L. C., and J. A. Krosnick. 2017. "Attitude strength." *Annual Review Of Psychology* 68(1): 327–51.

Lalley, S. P., and E. G. Weyl. 2018. "Quadratic Voting: How Mechanism Design Can Radicalize Democracy." *AEA Papers and Proceedings* 108: 33–7.

Lee, M. K., D. Kusbit, A. Kahng, J. T. Kim, X. Yuan, A. Chan, D. See, R. Noothigattu, S. Lee, A. Psomas, and A. D. Procaccia. 2019. "Webuildai: Participatory Framework for Algorithmic Governance." *Proceedings of the ACM on Human Computer Interaction* 3: 1–35. https://doi.org/10.1145/3359283

Likert, R. 1932. "A Technique for the Measurement of Attitudes." *Archives of Psychology* 22(140): 55. https://psycnet.apa.org/record/1933-01885-001

Mill, J. S. 1859. *On Liberty*. London: J.W. Parker & Son London.

Miller, J. M., and D. A. M. Peterson. 2004. "Theoretical and empirical implications of attitude strength." *The Journal of Politics* 66(3): 847–867.

Posner, E. A. 2018. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Edited by Eric Posner and Glen Weyl. Princeton: Princeton University Press.

Quarfoot, D., D. von Kohorn, K. Slavin, R. Sutherland, D. Goldstein, and E. Konar. 2017. "Quadratic Voting in the Wild: Real People, Real Votes." *Public Choice* 172(1-2): 283–303.

Suter, V., M. Meckel, M. Shahrezaye, and L. Steinacker. 2022. "AI Suffrage: A Four-Country Survey on the Acceptance of an Automated Voting System." In *Hawaii International Conference on System Sciences*. https://doi.org/10.24251/HICSS.2022.290

How to cite this article: Bassetti, Madeline E., Gustavo Dias, Daniel L. Chen, Alan Mortoni, and Ritesh Das. 2023. "Civicbase: An open-source platform for deploying Quadratic Voting for Survey Research." *AI Magazine* : 1–11. https://doi.org/10.1002/aaai.12103

AUTHOR BIOGRAPHIES

Madeline Bassetti is a User Experience Designer and Researcher. She works as a UX Consultant for the Development Impact Evalutation (DIME) department at the World Bank. Madeline holds a Bachelor of Business (Economics) from the Queensland University of Technology. She is interested in open-source software and public goods.

Gustavo Dias is an undergraduate at The University of Chicago. Gustavo is currently pursuing his bachelor's degree in Economics. His topics of interest are market design, voting mechanisms, and econometrics.

Daniel L. Chen is the Director of Research at the Centre National de la Recherche Scientifique (CNRS), and Professor at the Toulouse School of Economics. Previously, he was the Lead Principal Investigator with DE JURE (Data and Evidence for Justice Reform) at the World Bank. Chen received his BA and MS from Harvard University in Applied Mathematics and Economics; Economics PhD from MIT; and JD from Harvard Law School. He has attained prominence through the development of open source tools to study human behavior and through large-scale empirical studies—Data Science, Artificial Intelligence, and Machine Learning—on the relationship between law, social norms and the enforcement of legal norms, and on judicial systems.

Alan Mortoni is an experienced Front-end Developer skilled in React, Javascript, and Firebase. He has held positions at Flight Centre and Virgin Australia and holds a Bachelor of Computer Science from the Universidade Catolica de Brasília. He is a creative problem solver and dedicated to staying current with the latest web development technologies and trends.

Ritesh Das is a PhD student at the Robert H. Smith School of Business, University of Maryland. Prior to that he was working as a Consultant with the World Bank's DE JURE group in the domain of justice and development. He has prior experience in the consulting domain and cofounded two firms—one on rural infrastructure consulting and the other on monitoring and evaluation of government policies. He has completed 5-year integrated M.Sc. in Physics from IIT Kharagpur and Master of Applied Science in Data, Economics, and Development Policy from MIT.