

Heterogeneous participation and allocation skews: when is choice “worth it”?

NIKHIL GARG
Cornell Tech

A core ethos of the EconCS community is that people have complex private preferences and information of which the central planner is unaware, but which an appropriately designed mechanism can uncover to improve collective decisionmaking. This ethos underlies the community’s largest deployed success stories, from stable matching systems to participatory budgeting. I ask: is this choice and information aggregation “worth it”? In particular, I discuss how such systems induce *heterogeneous participation*: those already relatively advantaged are, empirically, more able to pay time costs and navigate administrative burdens imposed by the mechanisms. I draw on three case studies, including my own work – complex democratic mechanisms, resident crowdsourcing, and school matching. I end with lessons for practice and research, challenging the community to help reduce participation heterogeneity and design and deploy mechanisms that meet a “best of both worlds” north star: *use preferences and information from those who choose to participate, but provide a “sufficient” quality of service to those who do not.*

1. INTRODUCTION

A deserved point of pride for the EconCS community is the integration into everyday life the systems we have long studied, an integration often done in collaboration with researchers. In New York City, I recently voted in a participatory budgeting election and used ranked choice voting for a mayoral primary election; my neighbors submit preferences to stable matching processes that assign their children to 3-k (for three year olds), pre-k, kindergarten, middle school, and high school; and the city has embraced crowdsourcing: whenever we encounter problems as mundane as potholes or as serious as suspected lead in our water, we can submit a 311 report or request a testing kit. Each of these systems represents a triumph of an underlying community ethos: that the people have complex preferences and information of which the government is unaware, but which an appropriately designed mechanism can uncover to improve collective decisionmaking.

This article’s purpose is to raise a simple, perhaps surprising, question: *is this choice and information aggregation “worth it”?* Just as democratic decisionmaking generally privileges those who (can) vote, these systems skew public resource allocation and decisionmaking in favor of those who (can) participate. And, as I will describe, substantial empirical evidence has established that participation in these mechanisms correlates with existing axes of privilege. Thus, we must ask whether the gain in information aggregation is worth the cost—or have we, in the guise of preference optimization, deployed ways to allocate scarce public resources to those best positioned to take advantage? As I will argue, this question is central to the legitimacy – and perceived legitimacy – of our systems.

My thesis is analogous to, and motivated by, those recently advanced in poli-

Author’s address: ngarg@cornell.edu

cymaking, public interest technology, and behavioral economics. In their seminal book, “Administrative Burdens: Policymaking by Other Means,” Herd and Moynihan [2019] argue that the information requirements to access rights such as voting and Medicare – often imposed in the name of safety, fraud detection, and choice – in practice cause people to not receive what they are entitled to. In their respective books, Schank and McGuinness [2021] and Pahlka [2023] argue that poor technology design – something as innocuous as long forms – contributes to this loss, even when well-intentioned. In “Scarcity: Why Having Too Little Means So Much,” Mullainathan and Shafir [2013] explain how poverty begets poverty, because it inhibits long-term planning in favor of urgent needs. All then argue that system designers must design with this phenomenon – the time cost of participation – in mind. Analogously, I argue that effective outcomes in the face of heterogeneous participation must be a primary design goal for our field, if we want our information aggregation mechanisms to be “worth it.” In other words, we should either deem “equal” participation as a necessary precondition to using choice to allocate scarce resources, or ensure that our mechanisms are robust despite heterogeneity.¹

In this article, I first detail three case studies central to our community: complex voting mechanisms, resident crowdsourcing, and school matching. In each, I overview the promise and on-the-ground realities of how these systems affect collective decisionmaking. I highlight recent research, including my own, that has sought to understand and close the gap caused by heterogeneous participation. I then summarize shared patterns from the three case studies, including potential solutions and design principles. Finally, I overview practical and research directions on the use of choice to allocate scarce public resources. I challenge us to meet a “best of both worlds” north star: *use preferences and information from those who participate, but provide a “sufficient” quality of service to those who do not.* Simultaneously, we should help develop approaches to support balanced participation.

2. CASE STUDIES

2.1 Complex democratic mechanisms

“Equal” voting rights and participation is central to democracy. Of course, equal participation is difficult to achieve; in the United States, eligible voters who are young, lower-income, racial and ethnic minorities, or have less formal education are less likely to vote [Hartig et al. 2023]. These patterns are also present in two democratic innovations advanced in the community: participatory budgeting and deliberative democracy (“citizen assemblies”).

In participatory budgeting, voters select which community projects to fund, from libraries in schools, to gym renovations, to park beautification. The Stanford Participatory Budgeting Platform has helped run over 150 elections, each of which may allocate millions of dollars [Gelauff and Goel 2024b]. New York City, Cambridge, Paris, Porto Alegre, Budapest, Helsinki, and many other cities globally all run participatory budgeting elections. At their best, these elections promise to increase

¹I use the words, “equal,” “heterogeneous,” and “representative” informally. What exactly constitutes equal, or equal enough, depends on context and may be subjective. See Chasalow and Levy [2021] for a history and analysis of “representativeness” as a “foundational yet slippery concept.”

civic engagement and ensure that project funding decisions are made by the people, instead of elected representatives or administrators. In deliberative democracy mechanisms, people (“panelists”) are selected to deliberate, potentially over several days, over a prescribed set of issues; they are polled before and after regarding their beliefs and sometimes are tasked to make recommendations; at their best, such processes gather a diverse set of people to make decisions in accordance with what a “public sample would think if it had better conditions and information with which to explore and define the issues” [Fishkin 1991]. Such processes have been used in over 25 countries, including to make constitutional amendments in Mongolia [Lee 2024], and EconCS researchers are involved in both building online deliberation platforms and in selecting panelists [Fishkin et al. 2019; Flanigan et al. 2021].

Given time costs and the use of unfamiliar methods, ensuring representative participation in these processes is a continuous, challenging task, on which researchers have rightfully focused. Participatory budgeting is often conducted online and open to all residents, including children and non-citizens; however, turnout rates are sometimes low, including at or below 5% of eligible voters, and there may be unequal participation rates by race, ethnicity, education, immigration status, and home ownership [Zepic et al. 2017; Stewart et al. 2014; Hayduk et al. 2017]. Others report that participatory budgeting increases civic engagement by otherwise disadvantaged groups [Johnson et al. 2023], and there are mixed findings on its distributive effects [Shybalkina and Bifulco 2019; Stewart et al. 2014].

Motivated by unequal participation in participatory budgeting, Gelauff et al. [2020] and Shen et al. [2021] study targeted advertising for demographically balanced participation. Gelauff and Goel [2024a] advocate for the design of “civic feedback processes that are robust against disparities in the representation of demographic and opinion minorities,” including reweighting techniques “for more equitable voice among demographic minorities which were underrepresented in the process;” such reweighting could especially be appropriate for processes which are consultative for policymakers as opposed to binding.

Analogously, motivated by unequal volunteer and dropout rates in deliberative democracy, an important line of work shows how balanced panels can be selected (“sortition”). Both individual fairness (volunteers should have sufficiently high selection probabilities, even if from overrepresented groups) and overall representative balance (on both observed and unobserved covariates) are important [Benadè et al. 2019; Flanigan et al. 2020; Flanigan et al. 2021; Flanigan et al. 2021; Ebadian et al. 2022; Flanigan et al. 2023; Baharav and Flanigan 2024; Flanigan et al. 2024; Caragiannis et al. 2024; Ebadian and Micha 2025; Assos et al. 2025]. Their algorithms have been deployed at scale to support panel selection [Flanigan et al. 2021].

What lessons does this literature provide? (1) Representative participation is seen as a central design goal by researchers and practitioners, including in the EconCS community—it is well accepted that an unrepresentative process is not legitimate, though there is empirical equivocation on real-world participation disparities and its effect on resource allocation. (2) Representative participation (and overall rates) is nevertheless an ongoing challenge. In NYC participatory budgeting, fewer than 100,000 people vote city-wide,² far less than even other local elections such as for

²The exact number of people eligible is unclear, as it depends on the current number of residents

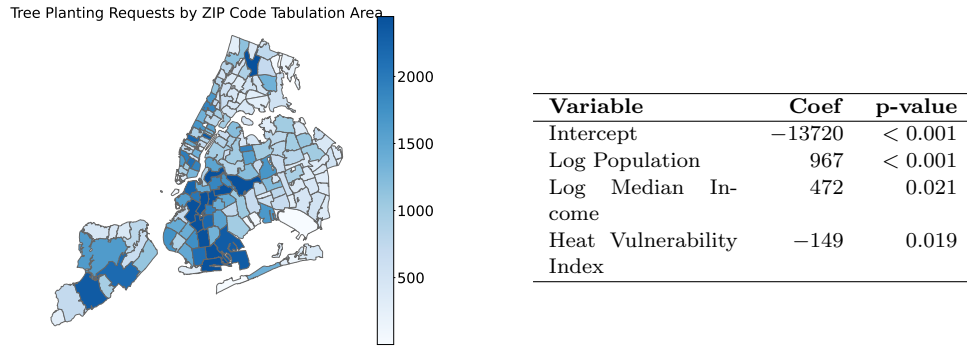


Fig. 1: In NYC, the number of tree planting requests by ZIP Code Tabulation Area in 2015-2024. Controlling for population, requests correlate positively with neighborhood income and *negatively* with a heat vulnerability index, a proxy for the need for shade. NYC no longer takes requests to plant trees, and instead will develop a schedule that prioritizes the most heat vulnerable areas.

city council—and some have questioned its use for decisionmaking [Golliher 2025]. It thus remains unclear how to mitigate the effects of heterogeneous participation.

2.2 Resident crowdsourcing

Another type of system deployed at scale to influence public resource allocation is *resident crowdsourcing*: people make service requests, such as through “311 systems” in the United States. NYC receives over 3 million requests a year—for incidents ranging from fallen trees on powerlines, to potholes, flooding, rodents, and to request new tree planting—and similar systems are in place in hundreds of cities globally. This is an important avenue for the government to learn about problems – supplementing and informing less frequent active inspections – and there is a large government bureaucracy to respond to requests.

However, substantial research, including my own, has established that participation is heterogeneous, even *conditional on ground truth conditions*. For example, in Liu et al. [2024], we show how to use *duplicate* reports about the same incident to estimate reporting delays; in Agostini et al. [2024], we use *spatial correlation* to probabilistically identify unreported incidents; in Balachandar et al. [2025], we combine regularly scheduled government inspections with crowdsourced reports; and in Franchi et al. [2025], we identify true flooding prevalence using a vision-language model on dashcams street imagery. In all cases, despite the diverse identification strategies, we find that (a) crowdsourced reporting data can be informative about ground truth conditions, e.g., that more hazardous conditions are reported at far higher rates [Liu et al. 2024]; but also (b) reporting is correlated with socioeconomic characteristics, also substantially: e.g., in Liu et al. [2024], we find that higher income, population density, voter participation, fraction of people with college degrees, and fraction of the population that is white all correlate with higher reporting rates. These patterns induce heterogeneous delays in incidents being

over 11 years old in the city council districts that participated. There are almost 8 million residents over 10 years old in NYC, implying a less than 2% turnout rate.

addressed, potentially leading to inequitable government service.

What should we do given this heterogeneous participation? I do *not* believe that these results imply we should not crowdsource this information; rather, we should ensure efficient resource allocation despite it, as in the work to balance citizen assemblies. For example, in a followup project, we seek to optimize inspection resources to efficiently and equitably set service level agreements [Liu and Garg 2024]; it would be conceptually simple to account for heterogeneous reporting delays.

I believe that the design of such modifications is urgent, before practitioners decide that information aggregation is not worth the resulting allocation skews. NYC no longer allows the public to request new tree planting locations; instead, the Department of Parks and Recreation will plant “street trees on a cyclical basis and prioritiz[e] the most heat-vulnerable neighborhoods first” [New York City 2025b]. A simple analysis using public data [New York City 2024; 2025a] helps explain why: as shown in fig. 1, planting requests historically correlated positively with median neighborhood income and *negatively* with heat vulnerability, one measure of “need.” Optimizing solely for stated resident preferences would lead to inefficient allocation, when the government has some expertise. Such pullback may also occur in other settings, if the mechanisms are not viewed as legitimate.

2.3 School matching

Finally, consider school matching. In many urban environments, students are assigned to public schools through the deferred acceptance algorithm [Abdulkadiroğlu et al. 2005]. The algorithm inputs applicant preferences (via ranked lists of schools) and school priorities (with factors such as geography, academic performance, diversity, and lottery numbers). The promise is twofold: (a) these systems provide the opportunity to access desired schools, even if they are not in the student’s neighborhood; (b) when slots in high-value schools are scarce, they are allocated not solely due to geography but also accounting for student preferences, academic performance, and random chance—thus, these allocation systems can be more effective and equitable than those that simply reflect geographic segregation.

In practice, applying effectively can be time consuming for families: in NYC, there are over 800 high school programs to choose from, each with varying locations, classes and sports teams offered, and school quality metrics. Families who can afford it often pay for admissions consultant services. A long line of research has empirically shown that information access, awareness, and the time-consuming process – not just preferences – affects application behavior, both in NYC high school admissions [Corradini 2024; Corradini and Idoux 2025] and elsewhere [Larroucau et al. 2024; Tomkins et al. 2023; Arteaga et al. 2022; Ajayi and Sidibe 2020].

The “administrative burdens” [Herd and Moynihan 2019] of applying lead to participation heterogeneity and outcome inequity. For example, Cohodes et al. [2022] documents the large fraction of students who apply to non-competitive, “nonoptimal” schools first in their rankings. In Peng et al. [2025], we show that such behavior leads to substantial “undermatching”: students not matching to as high-quality programs as they could have (that are no further geographically than their actual match), because they did not apply. In particular, this gap between where students matched and where they could have matched is almost twice as large for the most competitive Black and Hispanic students as it is for the corresponding

Asian and white ones, when quality is measured by program performance, value add, selectivity, school graduation rate, or college enrollment rate. We then show that simple application behaviors explain a large portion of this undermatching. Including with surveys, Corradini and Idoux [2025] show that differential awareness of schools rated as high-quality and racial homophily preferences explain such gaps, as opposed to preferences over other characteristics like quality.

Substantial work has further gone into developing and evaluating *informational interventions* to close the participation gaps [Corcoran et al. 2018; Arteaga et al. 2022; Cohodes et al. 2022; Corradini 2024; Larroucau et al. 2024]—for example, by providing students lists of high-quality programs close to their neighborhood. These interventions have changed behavior, when used.

Despite this focus, much work remains to be done. Informational gaps and heterogeneous participation persist, as documented by recent studies [Corradini and Idoux 2025; Peng et al. 2025]. One constant challenge, highlighted by Cohodes et al. [2022], is that informational interventions only work to the extent that they are *used*, i.e., they put the burden on participants—just as targetted advertising for participatory budgeting ultimately requires people to respond to the ads. Thus, it remains open how to deploy (a) interventions with high takeup rates and (b) new mechanisms robust to heterogeneous participation.

2.4 Common Themes and Implications

The above examples all follow a similar pattern: a mechanism allocates scarce public resources or makes joint decisions; a core mechanism component is to input preferences or information from participants; when the mechanism is deployed, participation is heterogeneous, despite it being monetarily “free.” Such heterogeneity both makes the mechanism less effective and potentially skews allocation and decisionmaking against those already disadvantaged. While substantial work has been done to measure and reduce these disparities, they persist.

Related concerns potentially apply in other settings in which preferences are elicited from participants who may have heterogeneous capabilities: in refugee matching, refugees may be asked for preferences over host countries [Jones and Teytelboym 2017]; in food bank allocation, large food banks (but not small ones) have dedicated staff to interface with the mechanism [Prendergast 2017]; in kidney exchange, preference elicitation from doctors regarding compatible kidneys is a practical challenge [Ashlagi and Roth 2021].

What should we do, given this fact pattern? In any given setting, the options are to (a) defend the status quo, by establishing that the mechanism is nevertheless effective, or at least preferable over any feasible counterfactual mechanism; (b) aim to reduce participation heterogeneity, as behavior is far from fixed; (c) reform the mechanism, so that it is robust; and/or (d) replace it entirely, likely to one that minimally uses the people’s preferences and information. How should we choose which option(s) to pursue? Different applications have and should take different paths, and the paths are complementary.³ In the face of heterogeneous participation, we

³There are key differences between the applications. Voting leads to a collective decision, and individuals who do not participate nevertheless benefit if they agree with those who do. In stable matching, allocations are individual and more arguably ‘zero-sum.’ Crowdsourcing lies in-between,

defend standard democracy *and* invest substantially in voter turnout efforts; this is also the path taken so far for participatory budgeting and deliberative democracy. In other cases, we’ve seen either reform or a retreat from participatory mechanisms.

I posit that reform should aim for the following “best of both worlds” north star: *use preferences and information from those who choose to participate, but provide a “sufficient” quality of service to those who do not.* In other words, we should attempt to retain the benefits of public participation while mitigating the resulting resource allocation skews. With this goal, we would still have “power users” who benefit from their invested time; however, non-participation would lead to a reasonable, default allocation. Of course, “sufficient” and “reasonable” are subjective, and themselves policy choices; when allocating scarce resources, these defaults may come at some cost to the “power users.” Policymakers and the public are in the best position to choose the context-dependent operating point on the inevitable tradeoff between information aggregation and allocation skews. My position is that this choice should be explicit, as opposed to the too-common status quo of maximally supporting aggregation at the cost of allocation skews.

3. APPROACHES FOR PRACTICE

In the remainder of this article, I highlight potential directions for practice and research, in service of this goal. These paths are informed by the above literature and Herd and Moynihan [2019] in particular. They lay out three reforms to respond to the “Medicare Maze,” in which the elderly must annually learn about complex options to choose a health care plan, leading to worse health outcomes and increased costs: (1) reduce choice by simplifying options; (2) expand outreach and human assistance in navigating the choices; (3) use administrative data and information technology to provide personalized defaults or recommendations. These options have their analogues for participatory mechanisms.

3.1 Reduce participation heterogeneity

The simplest response to participation heterogeneity is to try to reduce it. In participatory budgeting, this is done via targeted advertising; in deliberative democracy, this is done more directly by modifying selection probabilities.⁴ This approach is also a key tool to reduce disparities in the takeup of other entitlements, like SNAP benefits in the United States; Koenecke et al. [2023] show public support for targeted advertising that improves allocation equity. However, as continued disparities prove, turnout efforts are not a panacea in the presence of structural barriers to participation, such as those discussed by Mullainathan and Shafir [2013]. Approaches that more directly tackle structural barriers, such as those that provide childcare and video conferencing technology for deliberative democracy, may be necessary.

Another, more systematic approach to reducing participation *heterogeneity* is to use preferences *within* areas with relatively homogeneous participation. For exam-

as allocations (e.g., pothole fixes) are geographically localized, but everyone may benefit from the information shared by participants. The exposition has ignored these differences, as they are not crucial to my core thesis. However, they may be relevant in considering paths forward.

⁴This is only possible because deliberative democracy purposely is designed to select a subset of the people, with the goal of making that subset representative.

ple, NYC runs participatory budgeting separately for each city council district, with a set budget per district; *if* districts are drawn such that participation is similar within each district, then heterogeneity across districts would not skew allocations.⁵ Analogously, in the tree planting context, the following approach could incorporate geographic balance, need as determined by the agency, and resident requests: make neighborhood-level scheduling and quantity decisions according to agency expertise; then, within each neighborhood, allow requests to inform precise planting locations, alongside expertise. Appropriately designed, such an approach could be “best of both worlds” and combine elicited preferences with expert decisionmaking.

3.2 Provide personalized defaults or recommendations

Turnout campaigns may be effective when participation is (meaningfully approximated as) binary: in voting mechanisms and resident crowdsourcing, the most important outcomes are how regularly someone votes or submits requests. In school matching, on the other hand, *whether* people submit ranked lists is not the only concern, as doing so is required to enroll a child in public school. Rather, submitting *informed* ranked lists is a challenge, as it requires awareness of program quality and admissions probabilities; only some may have access to expensive consulting services or advice from social networks to help them navigate these decisions. In such systems, practitioners, alongside researchers, may be able to provide recommendations or even default options to users. Then, applicants can – just as in the status quo – provide preferences if they are dissatisfied with the recommendations or defaults; others can choose to follow the recommendations. Of course, as with targetted advertising, one challenge with recommendations is takeup [Cohodes et al. 2022], and so stronger user interfaces or “nudges” are important.

In many cases, as in school matching, there already *is* a default option, e.g., a manual administrative placement if an applicant does not match with any school. One approach is for these defaults to be more systemically planned, to provide better allocations to those who do (can) not participate meaningfully. Recommendations and defaults are also related to – and ‘lighter-touch’ than – another approach developed in school matching: limiting options, potentially in a data-driven manner: Shi [2015] develops short choice menus for each family in Boston, citing “too many options” as contributing to long commute times, unpredictability, loss of neighborhood cohesion, and a research burden on families; Allman et al. [2022] develop small zones in San Francisco, in support of school diversity.

The use of personalized defaults and recommendations, powered by modern machine learning methods, may also be effective in other contexts. Ashlagi and Roth [2021] advocate for a related approach in the context of preference elicitation difficulties for kidney exchanges: “it may be useful to develop machine learning models to predict positive crossmatches and ... to understand the trade-offs involved with waiting (while on dialysis) for a better match.”

⁵It is not clear that NYC’s districts meet this criteria. My district spans relatively wealthy areas in the Upper West Side, to Columbia University, to lower-income areas in West Harlem. However, granular turnout data is unavailable and winning projects did not geographically concentrate.

3.3 Actively acquire information or post-process inputs

A third approach is for the central system to actively invest in information acquisition to counter participation biases. When resident crowdsourcing informs resource allocation, for example, agencies can invest more active inspection resources or install sensors (e.g., flood sensors [Franchi et al. 2025]) in neighborhoods with lower reporting rates. Alternatively, given the public’s heterogeneous inputs, the system can make decisions that are nevertheless balanced. As an example of such post-processing, consider our work with the New York Public Library on the holds system, which allows patrons to request books from any system branch to be sent to their local neighborhood branch; we first found that heterogeneous usage of the holds system (even conditional on overall library usage) led to a large net outflow of books from lower-income neighborhoods to higher-income ones [Liu et al. 2024]. We then designed a routing prioritization scheme between branches to mitigate such disparities [Liu et al. 2025], so that all holds requests could be fulfilled without disproportionately depleting branches in lower holds-use neighborhoods.

However, these approaches are not always feasible. In deferred acceptance, where applicant preferences are directly used, it is unclear where active information acquisition can be incorporated or how matches can be post-processed. In democratic systems such as participatory budgeting, weighting votes may conflict with other design principles, such as ‘one-person-one-vote’ (as discussed by [Gelauff and Goel 2024a]). Such approaches may be feasible when constructing error bars or using vote outcomes to advise final decision-makers; however, the question of “representativeness” (and *of whom*) remains, especially when participation correlates with unobserved features [Chasalow and Levy 2021].

These solutions are analogous to those proposed in algorithmic fairness, to counter disparities in prediction accuracy that are caused by heterogeneous unobserved confounding or missing data. There, data may be actively acquired or post-processed while using demographics as features [Chen et al. 2018; Noriega-Campero et al. 2019; Caton and Haas 2020; Cai et al. 2020; Garg et al. 2021; Liu and Garg 2021; Movva et al. 2023; Zink et al. 2024; Balachandar et al. 2024; Dong et al. 2025; Chiang et al. 2025]. There as well, post-processing may be infeasible, due to legal constraints or a general preference for “group-unaware” approaches (e.g., the recent affirmative action ban in college admissions in the United States, which also affect algorithms in the admissions process [Lee et al. 2024]). More generally, I believe that the goal of countering heterogeneous participation may further connect market design to algorithmic fairness, cf. Finocchiaro et al. [2021].

All three approaches use central resources to counter heterogeneous participation and pursue “best of both worlds”: use elicited preferences, but mitigate allocation skews. Next, I discuss how researchers can contribute to the vision.

4. RESEARCH DIRECTIONS

Researchers have an important role to play in collaborating with practitioners on designing, deploying, and evaluating the above approaches. Researchers—including those who do not collaborate with practitioners—can also contribute in other ways. Below, I overview three approaches for a diverse range of skillsets: (a) empirically quantifying heterogeneous participation; (b) providing theoretical insight on

participation-allocation tradeoff and designing mechanisms to navigate it; (c) more directly considering human-computer-market interactions and interface design.

4.1 Empirically quantify heterogeneous participation

Academics – through open data, information requests, or practitioner collaboration – can help quantify participation heterogeneity. Methodologically, the challenge is that quantifying participation heterogeneity often requires disambiguating it from other, less concerning, explanations. In resident crowdsourcing, we must show that heterogeneous *conditions* cannot explain the discrepancy – that it is not the case that some neighborhoods report less because they encounter fewer incidents worth reporting [Liu et al. 2024; Agostini et al. 2024; Balachandar et al. 2025; Franchi et al. 2025]. In school matching, we must show that heterogeneous preferences – e.g., due to outside options or true heterogeneity in idiosyncratic preferences for certain schools or school characteristics – do not fully explain behavior, and that instead heterogeneous information plays an important role [Larroucau et al. 2024; Corradini and Idoux 2025; Corradini 2024]. This challenge often requires new statistical methods, analyzing natural experiments, or careful collection of “ground truth” data, such as through surveys and randomized controlled trials.

Such quantification helps provide an empirical underpinning with which interventions can be justified and well-engineered. For example, quantifying missing incident reports by neighborhood helps in the allocation of inspection and sensor resources in resident crowdsourcing, and quantifying heterogeneous awareness and behavior informs the design of personalized recommendations in school matching.

Finally, I note that empirically quantifying heterogeneous participation is related to two empirical lines of work: (1) preference estimation under strategic behavior [Agarwal and Somaini 2018; Calsamiglia et al. 2020], where the goal is to estimate preferences in non-strategyproof mechanisms, when (some) agents may be strategic; (2) empirical behavioral economics, that seeks to quantify how human behavior deviates from “optimal,” including in strategyproof mechanisms. Here, my focus is on quantifying *heterogeneous* behavior and its effects on downstream resource allocation, especially when there is no formal cost or strategic incentive.

4.2 Theoretically model allocation under heterogeneous participation and design mechanisms to explicitly navigate the participation-allocation tradeoff

Theoretical modeling of heterogeneous participation is a rich area for further study, to complement empirical measurement. Models can (a) elucidate welfare outcomes under heterogeneous participation; and (b) help design better mechanisms.

In the context of school matching, Kloosterman and Troyan [2020] analyze a setting in which some students are more informed than others about high quality options; under the model, such students may be *worse off* under deferred acceptance than without school choice; they then show that priorities may be designed in a way to avoid this outcome. Pathak and Sönmez [2008] analyze matching settings in which some students are “sophisticated” (strategic), while others are sincere despite strategic incentives; while “sincere students lose priority to sophisticated students” under the non-strategy-proof Boston mechanism, “any sophisticated student weakly prefers her assignment under the Pareto-dominant Nash equilibrium of the Boston mechanism to her assignment under the recently adopted student-optimal

stable mechanism.” It is essential to develop such models for other settings, as well as experiment with and deploy mechanisms with properties similar to the ones developed by Kloosterman and Troyan [2020]. More generally, some mechanisms may be more effective at supporting diverse participation.

One setting where such conceptual insights helped was the design of Feeding America’s market mechanism to allocate food to food banks. As detailed by Prendergast [2017], an essential consideration was to protect smaller food banks from heterogeneous participation, as they have “fewer resources and manpower ... relative to their larger counterparts, where there are often dozens of workers or volunteers.” The chosen mechanism avoided a continuous auction (which would benefit those with dedicated staff members) and allowed fractional bidding and storing of credits. It further effectively enabled a default option, giving food banks “the option to delegate bidding to an employee of Feeding America, where a food bank could simply outline in broad terms its needs to that person” [Prendergast 2017].

A related question suitable for modeling insight is: under what contexts is the participation-allocation tradeoff big, and when should we potentially abandon a mechanism? This question has recently been explored in the context of individual-level prediction to target resources: Shirali et al. [2024] argue that “prediction-based allocations outperform baseline methods using aggregate unit-level statistics only when between-unit inequality is low and the intervention budget is high,” i.e., that the cost of individualized predictions may not be worth it; Perdomo et al. [2023] empirically illustrate such ideas in the context of targeting interventions for students at risk of dropping out of school. Wang et al. [2024] argue against the legitimacy of decisionmaking that uses predictions of the future about individuals, due to reoccurring challenges regarding accuracy, disparate performance, and related concerns. Analogously, it may emerge in a model that eliciting preferences is only worth it when heterogeneity from preferences is larger than that from participation.

4.3 More directly consider human-computer-market interaction

Finally, the EconCS community should increase collaborations with human-computer interaction (HCI) researchers, to build interfaces that more effectively allow equal participation. Schank and McGuinness [2021] and Pahlka [2023] both pinpoint bad interface design as worsening government service. I posit that (1) good interfaces may be more effective than good theoretical properties in improving participation and systems, and (2) qualitative studies are important to understand participation. Here, I briefly overview my work and collaborations with HCI researchers.

In Bartle et al. [2025], we build and deploy an SMS-based system to help place patients being discharged from hospitals into nursing care homes. In our context in Hawai‘i, care homes are often run by retired nurses *out of their own homes*, with only one or two patients; whenever a patient needs to be placed, a full-time team of hospital social workers calls the approximately thousand nursing homes to see if they have capacity and can care for the given patient’s needs. This preference information is not centrally available because integration into a healthcare management system like Epic does not work for this rural, single-operator population. As we show, simply collecting capacity information – through SMS – and showing the data to hospital social workers trying to place patients into homes improves the process; my conjecture is that this data improvement – enabled by effective interface design

for care home operators – is far greater than improved call recommendations, such as through a matching optimization, would yield. However, one challenge is that many homes do not share their preferences over patient characteristics with the system. In followup work, we ran a randomized controlled trial and interviewed care home operators to understand this participation gap [Bartle et al. 2025]. The experiment revealed that nudges can (somewhat) increase the number of homes who share their preferences, and the interviews uncovered complex cultural phenomena as well as economic considerations that shape dynamic preferences. The mixed methods approach and collaboration across fields was essential in understanding participation and how interventions may increase it.

Substantial work has also shown that interface design can affect behavior in other systems studied by the community. In participatory budgeting, substantial work compares the behavioral and learning implications of the elicitation mechanism (i.e., whether they voters asked to rank or approve projects, or to modify a proposed budget or give a full budget) [Gelauff et al. 2018; Garg et al. 2019; Garg et al. 2019; Goel et al. 2019; Gelauff and Goel 2024b]—with the hypothesis that some mechanisms may be easier for voters to understand. In ratings systems, simple modifications such as the question that is asked can substantially affect behavior, by aligning different people on what “five stars” actually means [Garg and Johari 2019; 2021]. Similarly, I conjecture that interface design could reduce participation gaps in other systems. While simplifying interfaces is likely to be generally useful, open questions remain on how to best present information, including recommendations. Future work should experimentally evaluate interfaces and qualitatively interview participants regarding how they perceive a given interface and system design.

5. CONCLUSION

Economic and computational researchers have important roles to play in designing and analyzing societal systems [Roth 2002; Abebe et al. 2020]. Our community should be proud of our impact in influencing the deployment of so many real-world systems. Undoubtedly, many of these systems improve upon those that they replaced. However, just as we theoretically design mechanisms to be strategyproof, so that people *can safely* share their true preferences, we should focus on whether people *do* participate on equal footing, or can do so in the presence of heterogeneous time costs. We should further engineer our systems – theoretically, algorithmically, and through interface design – so that they do not inadvertently allocate scarce resources according to participation ability. In this article, I overviewed research, including much of my own, in pursuit of this goal. I believe that “best of both worlds” systems, that incorporate preferences without allowing heterogeneous participation to skew distributional outcomes, are possible and necessary.

Acknowledgments

I thank Rediet Abebe, Bailey Flanigan, Lodewijk Gelauff, Paul Gözl, Allison Koencke, Karen Levy, Irene Lo, and members of Cornell’s AI, Policy, and Practice Initiative for invaluable discussion and feedback, as well as my advisors, students, and collaborators. NG is supported by NSF CAREER IIS-2339427, and Cornell Tech Urban Tech Hub, Meta, and Amazon research awards.

REFERENCES

- ABDULKADIROĞLU, A., PATHAK, P. A., AND ROTH, A. E. 2005. The New York City High School Match. *American Economic Review* 95, 2 (Apr.), 364–367.
- ABEBE, R., BAROCAS, S., KLEINBERG, J., LEVY, K., RAGHAVAN, M., AND ROBINSON, D. G. 2020. Roles for computing in social change. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 252–260.
- AGARWAL, N. AND SOMAINI, P. 2018. Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica* 86, 2, 391–444.
- AGOSTINI, G., PIERSON, E., AND GARG, N. 2024. A bayesian spatial model to correct under-reporting in urban crowdsourcing. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. 21888–21896.
- AJAYI, K. AND SIDIBE, M. 2020. School choice under imperfect information. *Economic Research Initiatives at Duke (ERID) Working Paper* 294.
- ALLMAN, M., ASHLAGI, I., LO, I., LOVE, J., MENTZER, K., RUIZ-SETZ, L., AND O’CONNELL, H. 2022. Designing school choice for diversity in the san francisco unified school district. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 290–291.
- ARTEAGA, F., KAPOR, A. J., NEILSON, C. A., AND ZIMMERMAN, S. D. 2022. Smart matching platforms and heterogeneous beliefs in centralized school choice*. *The Quarterly Journal of Economics* 137, 3 (03), 1791–1848.
- ASHLAGI, I. AND ROTH, A. E. 2021. Kidney exchange: An operations perspective. *Management Science* 67, 9, 5455–5478.
- ASSOS, A., BAHARAV, C., FLANIGAN, B., AND PROCACCIA, A. 2025. Alternates, assemble! selecting optimal alternates for citizens’ assemblies. *Available at SSRN 5283438*.
- BAHARAV, C. AND FLANIGAN, B. 2024. Fair, manipulation-robust, and transparent sortition. In *Proceedings of the 25th ACM Conference on Economics and Computation*. 756–775.
- BALACHANDAR, S., GARG, N., AND PIERSON, E. 2024. Domain constraints improve risk prediction when outcome data is missing. In *The Twelfth International Conference on Learning Representations*.
- BALACHANDAR, S., SADHUKA, S., BERGER, B., PIERSON, E., AND GARG, N. 2025. Urban incident prediction with graph neural networks: Integrating government ratings and crowdsourced reports. *arXiv preprint arXiv:2506.08740*.
- BARTLE, V., DELL, N., AND GARG, N. 2025. Shopping around: An experiment in preferences and incentives for placing long-term patients.
- BARTLE, V., SHEARER, A., WROE, A., DELL, N., AND GARG, N. 2025. Faster information for effective long-term discharge: A field study in adult foster care. *Proceedings of the ACM on Human-Computer Interaction* 9, 2, 1–29.
- BENADÈ, G., GÖLZ, P., AND PROCACCIA, A. D. 2019. No stratification without representation. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 281–314.
- CAI, W., GAEBLER, J., GARG, N., AND GOEL, S. 2020. Fair allocation through selective information acquisition. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 22–28.
- CALSAMIGLIA, C., FU, C., AND GUELL, M. 2020. Structural estimation of a model of school choices: The boston mechanism versus its alternatives. *Journal of Political Economy* 128, 2 (Feb.), 642–680.
- CARAGIANNIS, I., MICHA, E., AND PETERS, J. 2024. Can a few decide for many? the metric distortion of sortition. In *Proceedings of the 41st International Conference on Machine Learning*. ICML’24. JMLR.org.
- CATON, S. AND HAAS, C. 2020. Fairness in Machine Learning: A Survey. *ACM Computing Surveys*.
- CHASALOW, K. AND LEVY, K. 2021. Representativeness in statistics, politics, and machine learning. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 77–89.
- CHEN, I., JOHANSSON, F. D., AND SONTAG, D. 2018. Why is my classifier discriminatory? *Advances in neural information processing systems* 31.

- CHIANG, E., SHANMUGAN, D., BEECY, A. N., SAYER, G., ESTRIN, D., GARG, N., AND PIERSON, E. 2025. Learning Disease Progression Models That Capture Health Disparities. In *Conference on Health, Inference, and Learning (CHIL '25)*.
- COHODES, S. R., CORCORAN, S. P., JENNINGS, J. L., AND SATTIN-BAJAJ, C. 2022. When Do Informational Interventions Work? Experimental Evidence from New York City High School Choice. *Educational Evaluation and Policy Analysis*, 01623737231203293.
- CORCORAN, S. P., JENNINGS, J. L., COHODES, S. R., AND SATTIN-BAJAJ, C. 2018. Leveling the playing field for high school choice: Results from a field experiment of informational interventions. Working Paper 24471, National Bureau of Economic Research. March.
- CORRADINI, V. 2024. Information and Access in School Choice Systems: Evidence from New York City.
- CORRADINI, V. AND IDOUC, C. M. 2025. Overcoming racial gaps in school preferences: The effect of peer diversity on school choice. Tech. rep.
- DONG, E., SCHEIN, A., WANG, Y., AND GARG, N. 2025. Addressing discretization-induced bias in demographic prediction. *PNAS nexus*, pgaf027.
- EBADIAN, S., KEHNE, G., MICHA, E., PROCACCIA, A. D., AND SHAH, N. 2022. Is sortition both representative and fair? *Advances in Neural Information Processing Systems* 35, 3431–3443.
- EBADIAN, S. AND MICHA, E. 2025. Boosting sortition via proportional representation. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems. AAMAS '25*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 667–675.
- FINOCCHIARO, J., MAIO, R., MONACHOU, F., PATRO, G. K., RAGHAVAN, M., STOICA, A.-A., AND TSIRTIS, S. 2021. Bridging machine learning and mechanism design towards algorithmic fairness. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 489–503.
- FISHKIN, J., GARG, N., GELAUFF, L., GOEL, A., MUNAGALA, K., SAKSHUWONG, S., SIU, A., AND YANDAMURI, S. 2019. Deliberative democracy with the online deliberation platform. In *Demo at the 7th AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2019 Demo)*.
- FISHKIN, J. S. 1991. *Democracy and deliberation: New directions for democratic reform*. Yale University Press.
- FLANIGAN, B., GÖLZ, P., GUPTA, A., HENNIG, B., AND PROCACCIA, A. D. 2021. Fair algorithms for selecting citizens' assemblies. *Nature* 596, 7873, 548–552.
- FLANIGAN, B., GÖLZ, P., GUPTA, A., AND PROCACCIA, A. D. 2020. Neutralizing self-selection bias in sampling for sortition. *Advances in Neural Information Processing Systems* 33, 6528–6539.
- FLANIGAN, B., GÖLZ, P., AND PROCACCIA, A. 2023. *Mini-Public Selection: Ask What Randomness Can Do for You*. Ash Institute for Democratic Governance and Innovation.
- FLANIGAN, B., KEHNE, G., AND PROCACCIA, A. D. 2021. Fair sortition made transparent. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds. Vol. 34. Curran Associates, Inc., 25720–25731.
- FLANIGAN, B., LIANG, J., PROCACCIA, A. D., AND WANG, S. 2024. Manipulation-robust selection of citizens' assemblies. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence*. AAAI Press.
- FRANCHI, M., GARG, N., JU, W., AND PIERSON, E. 2025. Bayesian modeling of zero-shot classifications for urban flood detection. *arXiv preprint arXiv:2503.14754*.
- GARG, N., GELAUFF, L. L., SAKSHUWONG, S., AND GOEL, A. 2019. Who is in your top three? optimizing learning in elections with many candidates. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*. Vol. 7. 22–31.
- GARG, N. AND JOHARI, R. 2019. Designing optimal binary rating systems. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 1930–1939.
- GARG, N. AND JOHARI, R. 2021. Designing informative rating systems: Evidence from an online labor market. *Manufacturing & Service Operations Management* 23, 3, 589–605.
- GARG, N., KAMBLE, V., GOEL, A., MARN, D., AND MUNAGALA, K. 2019. Iterative local voting for collective decision-making in continuous spaces. *Journal of Artificial Intelligence Research* 64, 315–355.

- GARG, N., LI, H., AND MONACHOU, F. 2021. Standardized tests and affirmative action: The role of bias and variance. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 261–261.
- GELAUFF, L. AND GOEL, A. 2024a. Opinion change or differential turnout: Changing opinions on the austin police department in a budget feedback process. *Digit. Gov.: Res. Pract.* 5, 3 (Sept.).
- GELAUFF, L. AND GOEL, A. 2024b. Rank, pack, or approve: Voting methods in participatory budgeting. In *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 18. 448–461.
- GELAUFF, L., GOEL, A., MUNAGALA, K., AND YANDAMURI, S. 2020. Advertising for demographically fair outcomes. *arXiv preprint arXiv:2006.03983*.
- GELAUFF, L., SAKSHUWONG, S., GARG, N., AND GOEL, A. 2018. Comparing voting methods for budget decisions on the assu ballot. Tech. rep., Technical report.
- GOEL, A., KRISHNASWAMY, A. K., SAKSHUWONG, S., AND AITAMURTO, T. 2019. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation (TEAC)* 7, 2, 1–27.
- GOLLIHER, D. 2025. The city council’s participatory budgeting by the numbers. Maximum New York.
- HARTIG, H., DANILLER, A., KEETER, S., AND VAN GREEN, T. 2023. Republican gains in 2022 midterms driven mostly by turnout advantage.
- HAYDUK, R., HACKETT, K., AND FOLLA, D. T. 2017. Immigrant engagement in participatory budgeting in new york city. *New Political Science* 39, 1, 76–94.
- HERD, P. AND MOYNIHAN, D. P. 2019. *Administrative burden: Policymaking by other means*. Russell Sage Foundation.
- JOHNSON, C., CARLSON, H. J., AND REYNOLDS, S. 2023. Testing the participation hypothesis: Evidence from participatory budgeting. *Political Behavior* 45, 1, 3–32.
- JONES, W. AND TEYTELBOYM, A. 2017. The international refugee match: A system that respects refugees’ preferences and the priorities of states. *Refugee Survey Quarterly* 36, 2, 84–109.
- KLOOSTERMAN, A. AND TROYAN, P. 2020. School choice with asymmetric information: Priority design and the curse of acceptance. *Theoretical Economics* 15, 3, 1095–1133.
- KOENECKE, A., GIANNELLA, E., WILLER, R., AND GOEL, S. 2023. Popular support for balancing equity and efficiency in resource allocation: A case study in online advertising to increase welfare program awareness. In *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 17. 494–506.
- LARROUCAU, T., RIOS, I., FABRE, A., AND NEILSON, C. 2024. College application mistakes and the design of information policies at scale.
- LEE, J., HARVEY, E., ZHOU, J., GARG, N., JOACHIMS, T., AND KIZILCEC, R. F. 2024. Ending affirmative action harms diversity without improving academic merit. In *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–17.
- LEE, S. 2024. Deliberative polling on constitutional amendments in mongolia. Tech. rep., Global Assembly.
- LIU, Z., BHANDARAM, U., AND GARG, N. 2024. Quantifying spatial under-reporting disparities in resident crowdsourcing. *Nature Computational Science* 4, 1, 57–65.
- LIU, Z. AND GARG, N. 2021. Test-optional policies: Overcoming strategic behavior and informational gaps. In *Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–13.
- LIU, Z. AND GARG, N. 2024. Redesigning service level agreements: Equity and efficiency in city government operations. In *Proceedings of the 25th ACM Conference on Economics and Computation*. 309–309.
- LIU, Z., RANKIN, S., AND GARG, N. 2024. Identifying and addressing disparities in public libraries with bayesian latent variable modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. 22258–22265.
- LIU, Z., ZHU, W., RANKIN, S., AND GARG, N. 2025. Optimizing library usage and browser experience: Application to the new york public library. *arXiv preprint arXiv:2503.23118*.

- MOVVA, R., SHANMUGAM, D., HOU, K., PATHAK, P., GUTTAG, J., GARG, N., AND PIERSON, E. 2023. Coarse race data conceals disparities in clinical risk score performance. In *Machine Learning for Healthcare Conference*. PMLR, 443–472.
- MULLAINATHAN, S. AND SHAFIR, E. 2013. *Scarcity: Why having too little means so much*. Macmillan.
- NEW YORK CITY. 2024. Heat vulnerability index rankings. Available at: <https://data.cityofnewyork.us/Health/Heat-Vulnerability-Index-Rankings/4mhf-duep>; Accessed: May 16, 2025.
- NEW YORK CITY. 2025a. 311 service requests from 2010 to present. Available at: <https://data.cityofnewyork.us/Social-Services/311-Service-Requests-from-2010-to-Present/erm2-nwe9>; Accessed: May 16, 2025.
- NEW YORK CITY. 2025b. Tree planting. Available at: <https://portal.311.nyc.gov/article/?kanumber=KA-01895>, Accessed: May 16, 2025.
- NORIEGA-CAMPERO, A., BAKKER, M. A., GARCIA-BULLE, B., AND PENTLAND, A. 2019. Active fairness in algorithmic decision making. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 77–83.
- PAHLKA, J. 2023. *Recoding America: why government is failing in the digital age and how we can do better*. Metropolitan Books.
- PATHAK, P. A. AND SÖNMEZ, T. 2008. Leveling the playing field: Sincere and sophisticated players in the boston mechanism. *American Economic Review* 98, 4, 1636–1652.
- PENG, K., RYU, E., KLEINBERG, J., TARDOS, E., AND GARG, N. 2025. Deviations from reach-match-safety strategies explain undermatching disparities in new york city high schools.
- PERDOMO, J. C., BRITTON, T., HARDT, M., AND ABEBE, R. 2023. Difficult lessons on social prediction from wisconsin public schools. *arXiv preprint arXiv:2304.06205*.
- PRENDERGAST, C. 2017. How food banks use markets to feed the poor. *Journal of Economic Perspectives* 31, 4, 145–162.
- ROTH, A. E. 2002. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* 70, 4, 1341–1378.
- SCHANK, H. AND MCGUINNESS, T. D. 2021. *Power to the public: The promise of public interest technology*. Princeton University Press.
- SHEN, Z., GELAUFF, L., GOEL, A., KOROLOVA, A., AND MUNAGALA, K. 2021. Robust allocations with diversity constraints. *Advances in Neural Information Processing Systems* 34.
- SHI, P. 2015. Guiding school-choice reform through novel applications of operations research. *Interfaces* 45, 2, 117–132.
- SHIRALI, A., ABEBE, R., AND HARDT, M. 2024. Allocation requires prediction only if inequality is low. *arXiv preprint arXiv:2406.13882*.
- SHYBALKINA, I. AND BIFULCO, R. 2019. Does participatory budgeting change the share of public funding to low income neighborhoods? *Public Budgeting & Finance* 39, 1, 45–66.
- STEWART, L. M., MILLER, S. A., HILDRETH, R., AND WRIGHT-PHILLIPS, M. V. 2014. Participatory budgeting in the united states: a preliminary analysis of chicago’s 49th ward experiment. *New Political Science* 36, 2, 193–218.
- TOMKINS, S., GROSSMAN, J., PAGE, L., AND GOEL, S. 2023. Showing high-achieving college applicants past admissions outcomes increases undermatching. *Proceedings of the National Academy of Sciences* 120, 45, e2306017120.
- WANG, A., KAPOOR, S., BAROCAS, S., AND NARAYANAN, A. 2024. Against predictive optimization: On the legitimacy of decision-making algorithms that optimize predictive accuracy. *ACM Journal on Responsible Computing* 1, 1, 1–45.
- ZEPIC, R., DAPP, M., AND KRCMAR, H. 2017. Participatory budgeting without participants: Identifying barriers on accessibility and usage of german participatory budgeting. In *2017 Conference for E-Democracy and Open Government (CeDEM)*. IEEE, 26–35.
- ZINK, A., OBERMEYER, Z., AND PIERSON, E. 2024. Race adjustments in clinical algorithms can help correct for racial disparities in data quality. *Proceedings of the National Academy of Sciences* 121, 34, e2402267121.