

Predictive Algorithms and Perceptions of Fairness: Parent Attitudes Toward Algorithmic Resource Allocation in K-12 Education

Rebecca A. Johnson,^a Simone Zhang^b

a) Georgetown University; b) New York University

Abstract: As institutions increasingly use predictive algorithms to allocate scarce resources, scholars have warned that these algorithms may legitimize inequality. Although research has examined how elite discourses position algorithms as fair, we know less about how the public perceives them compared to traditional allocation methods. We implement a vignette-based survey experiment to measure perceptions of algorithmic allocation relative to common alternatives: administrative rules, lotteries, petitions from potential beneficiaries, and professional judgment. Focusing on the case of schools allocating scarce tutoring resources, our nationally representative survey of U.S. parents finds that parents view algorithms as fairer than traditional alternatives, especially lotteries. However, significant divides emerge along socioeconomic and political lines—lower socioeconomic status (SES) and conservative parents favor the personal knowledge held by counselors and parents, whereas higher SES and liberal parents prefer the impersonal logic of algorithms. We also find that, after reading about algorithmic bias, parental opposition to algorithms is strongest among those who are most directly disadvantaged. Overall, our findings map cleavages in attitudes that may influence the adoption and political sustainability of algorithmic allocation methods.

Keywords: predictive algorithms; algorithmic decision-making; public perceptions; K-12 education; educational inequality; resource allocation

Reproducibility Package: The data underlying this article are available as part of our replication materials available at this link: <https://doi.org/10.7910/DVN/EUJ1YZ>. The data for the main analyses is the TESS .tab format file at this link: <https://dataverse.harvard.edu/file.xhtml?fileId=10796997&version=1.0>

Citation: Johnson, A. Rebecca, Simone Zhang. 2025. "Predictive Algorithms and Perceptions of Fairness: Parent Attitudes Toward Algorithmic Resource Allocation in K-12 Education" *Sociological Science* 12: 322-356.


Received: November 12, 2024

Accepted: January 24, 2025

Published: May 16, 2025

Editor(s): Ari Adut, Filiz Garip

DOI: 10.15195/v12.a15

Copyright: © 2025 The Author(s). This open-access article has been published under a Creative Commons Attribution License, which allows unrestricted use, distribution and reproduction, in any form, as long as the original author and source have been credited. 

THE methods that organizations and institutions use to allocate scarce resources shape material inequalities, confer recognition, and shift power (Auyero 2012; Lamont, Beljean, and Clair 2014). Scholars have documented a long-term trend toward more standardized, quantified, and technocratic allocation approaches (Porter 1996; Berman 2022), exemplified in recent years through the increasing use of predictive algorithms (Burrell and Fourcade 2021; Joyce et al. 2021; Levy, Chasalow, and Riley 2021). These algorithms — computational models that predict future outcomes of individuals based on historical data (Wang et al. 2024) — guide allocation decisions across domains of social life. They direct the attention of law enforcement and child protective services (Brayne and Christin 2021; Eiermann 2024); inform evaluations of merit and value in hiring, admissions, and tenant screening (Engler 2021; Rosen, Garboden, and Cossyleon 2021; Ajunwa 2023); and assess need for resources such as shelter beds, organ transplants, and health care services (Eubanks 2018; Obermeyer et al. 2019; Robinson 2022).

Scholarly explanations for the growing use of predictive algorithms suggest that they appeal to policymakers and organizational leaders because they embody two features commonly thought to enhance legitimacy: precise targeting of resources to optimize efficiency or reinforce dominant conceptions of deservingness (Steensland 2006; Watkins-Hayes and Kovalsky 2016; Berman 2022), and reliance on impersonal criteria rather than fallible human judgment (Porter 1996; Espeland and Vannebo 2007; Burrell and Fourcade 2021). Accordingly, scholars warn that algorithms may justify bias, discrimination, or inequality (Starr 2014; Fourcade and Healy 2017; Hirschman and Bosk 2020), as they are positioned as ideal “calculable rules” applied “without regard for persons” (Weber 2009 [1921]).

Although existing research examines how algorithms are presented and legitimated by those with authority to adopt and justify them, we understand less about how the public perceives their use, especially compared to common alternatives. Historically, organizations have repeatedly arrived at similar approaches to allocating scarce resources: professional judgment, petitions from prospective beneficiaries, standardized rules, and randomization. Through the case of school districts allocating scarce tutoring resources, we investigate public perceptions of algorithms against these recurring alternatives that vary in their impersonality and degree of targeting.

These public perceptions matter because they can shape other important downstream outcomes, such as parental trust and engagement in the context of schools (Roda and Sattin-Bajaj 2023). Additionally, legitimacy is audience-specific (Schoon 2022), and public opinion can influence an allocation method’s political sustainability. Furthermore, attitudes toward allocation procedures are connected to broader belief systems of interest to social scientists, such as trust in science and public institutions (Gauchat 2012; Glass 2019).

We fielded a vignette-based survey experiment with a nationally representative sample of U.S. parents. Our vignette is based on early warning systems, which use predictive algorithms to identify students at risk of adverse educational outcomes such as grade retention or not completing high school. These systems are increasingly used to direct prevention resources in the United States and around the world (U.S. Department of Education 2016; Bowers 2021; Feathers 2023; Perdomo et al. 2023; Trinidad 2024).

We compare parents’ perceptions of using an algorithm to allocate scarce tutoring resources against four status quo allocation methods commonly used in U.S. schools: counselor judgment that relies on professional expertise, parent requests that draw on familial knowledge, bureaucratic rules that create standardized categories of need, and lotteries that randomize access. We assess the perceived fairness of each approach and analyze open-ended survey responses to understand parents’ underlying reasoning. Additionally, we present information on a way that the algorithm is biased against disadvantaged students to examine whether and for whom information about algorithmic bias influences support for algorithmic resource allocation.

Overall, we find that parents rate algorithms as fairer than traditional allocation methods. However, we observe important cleavages in attitudes that diverge from what we might predict based on work studying dynamics within K-12 education.

For example, there is no consistent trend in how parents from different racial/ethnic groups view the fairness of school counselors compared to algorithms, despite research on how race/ethnicity influences counselors' interactions with parents and students (Lewis-McCoy 2014; Cartwright 2022). Instead, the sharpest cleavages emerge in evaluations of the fairness of parent requests based on parental class and political orientation.

Our analysis further reveals that parents justify their assessments of fairness by focusing on how effectively they believe a method targets resources based on need and what type of knowledge about need the method draws upon. While parents consistently favor greater targeting, they are mixed on the normative valence of the impersonal knowledge of an algorithm versus the more personal knowledge held by parents or counselors. For instance, some parents value impersonal algorithmic knowledge as a means to bridge educational inequalities. Others, meanwhile, see it as overlooking valuable information that parental expertise could offer. Finally, we find parents vary in the degree to which they support algorithmic allocation after being presented with information about algorithmic bias. Most parents who initially rated algorithms as fairer later opposed their use after reading the update about bias, but opposition was strongest among low-income parents and parents with lower levels of education, who would be most directly impacted by the form of bias we describe.

This study makes several contributions. First, by examining algorithms alongside status quo alternatives that consistently emerge across institutional contexts, we ground assessments of algorithmic fairness in real-world trade-offs rather than abstract evaluations of algorithms alone (Binns et al. 2018; Smith, 2018; Dodge et al. 2019; Waldman and Martin 2022; Martin and Waldman 2023) or against unfettered human judgment (Lee 2018; Langer, König, and Papathanasiou 2019; Miller and Keiser 2021; Bankins et al. 2022; Kennedy, Waggoner, and Ward 2022). Second, our study clarifies the features of allocation approaches that might enhance or diminish their perceived fairness among various social groups. Third, our nationally representative sample improves upon past studies, which have nearly universally relied on online convenience samples. Finally, our results suggest that those better insulated from an algorithm's adverse impacts tend to be less moved by information about bias, pointing to potential fractures in attitudes that could facilitate the disproportionate application of algorithms to the least powerful members of society (Madden et al. 2017; Eubanks 2018; Barabas et al. 2020; Rona-Tas 2020).

Overall, this research demonstrates the importance of analytically separating how algorithms are presented from how they are perceived and probing cleavages in perceived legitimacy. Although we analyze these issues with a vignette-based survey experiment, the findings have broader relevance for those studying real-world allocation contexts, providing hypotheses to help investigate where and what form acceptance of and resistance to algorithms might take.

Situating Perceptions of Algorithms within Status Quo Approaches

Algorithms operate within an existing set of ways that organizations allocate scarce resources aimed at benefiting individuals. Although specific allocation methods vary, similar solutions appear repeatedly across a wide array of bureaucratic

contexts. One common method is the judgment of professionals or street-level bureaucrats who individually evaluate those seeking help (Lipsky 1980). This approach is the dominant point of comparison to algorithms in existing research on public perceptions (Lee 2018; Langer et al. 2019; Miller and Keiser 2021; Bankins et al. 2022; Kennedy et al. 2022).

However, the overriding focus on comparisons between human judgment and algorithms neglects several other widely used allocation approaches. Another human-driven method is beneficiary requests, where organizational attention and resources are directed in response to stakeholder claims. Examples include 311-based systems that allocate local government time and services (Levine and Gershenson 2014; Kontokosta and Hong 2021; Hamel and Holliday 2024) and complaint systems that initiate investigations into discrimination and harassment (Faber and Kalbfeld 2019).

Other allocation methods are more formalized, reflecting concerns about the potential corrupting influence of human discretion. Manually created rule and point systems prioritize individuals based on predefined categories. These span early COVID-19 vaccine allocation systems that prioritized by factors such as age, healthcare worker status, and health comorbidities (Jain, Schwarz, and Lorgelly 2021); kidney allocation systems that weigh considerations like time on dialysis, possession of certain antibodies, and whether one is awaiting multiple organs (Hart et al. 2017); and local child care and housing voucher allocation policies that consider a variety of demographic, economic, and household characteristics (Bouek 2023; Zhang and Johnson 2023). Unlike predictive algorithms, these manual point systems rely on human deliberation over which categories to weigh, rather than training a model on historical data to optimally predict a chosen outcome (Johnson and Zhang 2022).

A final approach is lotteries, which have been used to allocate a wide range of goods, from slots in K-12 and professional schools (ten Cate 2021; Kim 2024) to research grants (Liu et al. 2020) to housing assistance (Bueno, Nunes, and Zucco 2024). Lotteries may be unweighted or weighted. For example, some jurisdictions used a weighted lottery to distribute Emergency Rental Assistance during the COVID-19 pandemic that gave higher odds to residents of census tracts with a high social vulnerability index (Collinson et al. 2024) and some school choice systems give lottery priority to students in foster care or receiving welfare benefits (Kim 2024).

Targeting, Impersonality, and Public Perceptions

What organizes how people assess the fairness of these allocation methods? Although many aspects of allocation approaches can influence perceptions of fairness, we focus on two that vary most starkly: their degree of targeting and the type of knowledge they draw upon. These two features can shape both perceptions of procedural fairness — whether the decision-making process is fair — and distributive fairness — whether the outcomes that result are fair (Tyler 1996). For instance, people might prefer professional human judgment because they value the holistic knowledge that human experts can draw upon or because they think that human judgment will produce outcomes they see as fairer. Although analytically distinct,

process and outcome concerns often intertwine in people's assessments (MacCoun 2005), and we do not aim to separate them in this study. Instead, we focus on the attributes of allocation approaches that are most salient to people when they evaluate fairness.

Scholars have documented a sustained trend of formal organizations and institutions seeking legitimacy through the use of increasingly impersonal and targeted procedures to determine who gets what (Porter 1996; Burrell and Fourcade 2021; Levy et al. 2021). Underpinning this embrace of rubrics, scoring systems, and, increasingly, algorithms is the notion that more impersonal ways of knowing curtail the influence of personal favor, thereby enhancing perceived fairness (Espeland and Vannebo 2007; Fourcade and Healy 2017; Hirschman and Bosk 2020). Precise targeting of resources is thought to secure public approval by enabling actors to align outcomes with popular conceptions of deservingness (Schneider and Ingram 1993; Steensland 2006; Katz 2013). As economic reasoning has gained traction in policy debates, scholars suggest that efficient resource targeting has become an orienting ideal for organizational action, crowding out more universalist approaches appealing to equity or rights (Berman 2022).

Predictive algorithms exemplify this shift to more targeted and impersonal resource allocation. Accordingly, scholarly and popular accounts often portray algorithms as enjoying greater legitimacy and perceived fairness than traditional alternatives (Hirschman and Bosk 2020; Burrell and Fourcade 2021; Joyce et al. 2021). Yet several streams of scholarship challenge the idea that impersonality and targeting are universally appealing to the public and predict cleavages in views of algorithms across social groups.

People may be skeptical of impersonal approaches due to a mismatch between how people traditionally evaluate moral worth and the case-based logic of impersonal systems like bureaucratic rules and algorithms (Heimer 2001; Kiviat 2023). These systems slot people into standardized categories or scores, allocating resources based on predictive criteria without regard for common notions of moral blameworthiness or causality (Starr 2014; Fourcade and Healy 2017; Hirschman and Bosk 2020). This approach contrasts with the narrative knowledge commonly required in everyday moral reasoning, where evaluators assess an individual's character and deservingness within the context of their circumstances (Kiviat 2023). Consequently, the impersonality of rules and algorithms, which excludes personalized narrative information, may conflict with public expectations for moral discernment (Grgic-Hlaca et al. 2018; Kiviat 2021). For example, automated hiring is often perceived as unfair because it ignores potentially relevant qualitative, contextual information about job applicants (Newman, Fast, and Hamon 2020). More broadly, survey respondents tend to question the fairness of applying algorithms to decisions deemed personal rather than mechanistic in nature (Lee 2018), as in the case explored here of evaluating students' need for tutoring.

Greater targeting, too, may inspire opposition, especially when it stratifies access to goods considered universal rights. The public could also express discomfort with institutional actors "playing God" (Zhang and Johnson 2023) or engaging in social engineering. Indeed, normative scholarship on fair allocation often positions less precise allocation mechanisms such as lotteries as a favorable option for limiting

potential for discrimination, favoritism, and perpetuating inequalities (Goodwin Barbara 1992; Fang and Casadevall 2016; Vong 2020; Wang et al. 2024). However, a lottery-based approach linking outcomes to chance could clash with commitments to meritocracy as a dominant ideal in American society (Sauder 2020). Research on public housing lotteries, for instance, shows lotteries are unpopular for directing resources less precisely than alternatives (Bueno et al. 2024).

Beyond these overarching influences, additional literatures suggest cleavages in how different social groups perceive algorithmic decision-making. Research on attitudes toward science highlights the possibility of political polarization. Because algorithms are often framed as a scientific approach to decision-making, conservatives — who tend to be more skeptical of science's role in public life (Gauchat 2015) — might view algorithms with greater suspicion. This skepticism could stem from broader concerns about the administrative-regulatory state, distrust of centralized authority, rejection of expertise, and tensions between scientific and traditional moral authority (Gauchat 2023). This could lead conservatives to prefer less scientifically driven and more decentralized selection approaches, such as parental requests, over algorithms.

Scholarship on unequal experiences with K-12 schools offers another lens for identifying potential cleavages in perceptions. Although this research has not directly measured attitudes toward algorithms and status quo methods, it documents inequalities that could inform attitudes. Schools tend to be more responsive to advantaged, white parents who are more likely to assert their influence (Lareau 2000, 2003; Calarco 2018). In contrast, Black and Hispanic parents, regardless of socioeconomic status (SES), often face barriers and marginalization when engaging with schools (Lewis-McCoy 2014; Cartwright 2022). These experiences could foster divergent views of the fairness of parental input: advantaged parents may favor it, believing it ensures their influence, while lower-SES and racially minoritized parents might view impersonal systems like algorithms as a more equitable alternative.

Perceptions of allocation by counselor judgment might be similarly divided. Although counselors can ideally incorporate personalized, qualitative information, role conflict and limited resources hinder their ability to assess students' needs holistically and equitably (Sattin-Bajaj et al. 2018; Blake 2020). This leads counselors to offer generic advice and resources (Gast 2021), fostering distrust (Holland 2015). More advantaged students, better equipped to assert their needs (Calarco 2011), may benefit disproportionately from counselor-based systems and thus view them more favorably. Accordingly, we might expect greater support for counselor discretion among higher SES and white parents than lower SES and racially minoritized parents.

The remaining impersonal status quo methods — bureaucratic rules and lotteries — can also produce inequalities, though often in less visible ways, which may result in more muted divides in opinion. Bureaucratic rules can collapse complex forms of student need into crude categories such as free or reduced price lunch eligibility (Domina, Penner, and Penner 2017; Singer 2023) or exclude high-need individuals who lack the resources to navigate administrative requirements for proving eligibility for categories like homelessness or disability (Miller and Bourgeois 2013;

Herd and Moynihan 2019; Mirza et al. 2022). Lotteries can similarly introduce inequalities, as eligibility for inclusion is often determined by burdensome screening processes that can exclude less advantaged potential beneficiaries (Frankenberg, Siegel-Hawley, and Wang 2011; Skinner 2014; Berends 2015). However, because these inequalities are more abstract and hidden, they may not provoke strong attitudinal cleavages.

Algorithms and the Legitimization of Inequality

A central contention in sociological scholarship on algorithms is that the tools may legitimize discrimination and bias. Although all allocation methods can produce biased outcomes, scholars argue that algorithms pose unique concerns owing to their associations with objectivity and claims of predictive validity (Benjamin 2019; Hirschman and Bosk 2020; Burrell and Fourcade 2021; Joyce et al. 2021). These associations can create an “appearance of objectivity that insulates controversial decisions” (Hirschman and Bosk 2020: 350) and enable algorithms to be “promoted and perceived as more objective or progressive than the discriminatory systems of a previous era” (Benjamin 2019: 5–6).

This theorized relationship between algorithmic authority and the legitimization of inequality raises important questions: Does exposure to information about algorithmic errors that disproportionately impact disadvantaged groups influence support for algorithms? Who is more likely to object to their use given such information? Although the notion that algorithms provide special cover for bias is a common refrain in the sociological literature on algorithms, it has largely gone unmeasured and untested. To address this gap, we evaluate whether parents express support for algorithm-based resource allocation even when confronted with evidence of algorithmic bias.

Some allied research offers indirect support for the idea that algorithms can make bias appear more legitimate. For example, Tilcsik (2021) demonstrates how theories of statistical discrimination—where biased decisions result from decision-makers relying on exaggerated realities of group differences when they lack the information they want—can lead people to rationalize discriminatory decision-making. The algorithmic context may amplify this dynamic, as the individualized nature of algorithmic prediction could be seen as an advance over group-based inferences, making such rationalizations particularly compelling.

However, a growing interdisciplinary body of scholarship suggests subgroup differences in the degree to which algorithmic inequalities are perceived as acceptable or legitimate. Studies on algorithmic decision-making in other domains reveal that liberals prioritize equitable treatment across social groups more than conservatives do (Jakesch et al. 2022) and are less accepting of using sensitive inputs such as gender and race (Grgić-Hlača et al. 2022). More highly educated respondents are also more likely to be concerned about disparate treatment (Kieslich, Keller, and Starke 2022). Research also indicates that self-interest and group identities shape responses to information about algorithmic bias (Pierson 2018; Bankins et al. 2022). Those disadvantaged by the highlighted biases may grow more opposed, whereas those believing they stand to benefit or remain unaffected may maintain support (c.f. Grgić-Hlača et al. 2022). Members of ethnoracial groups with greater experience of

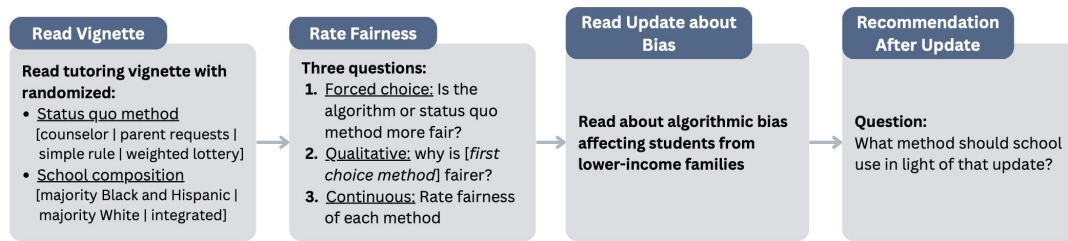


Figure 1: Survey Flow.

institutional discrimination, such as Black respondents in the U.S. context (Jakesch et al. 2022), may also be particularly attuned to bias against marginalized groups.

In this study, we examine whether support for algorithms changes in light of information about bias against disadvantaged groups. Although the sociological literature suggests that major drops in support for algorithms are unlikely, the interdisciplinary literature on attitudes predicts that support will wane unevenly: more among liberal, more highly educated, and/or racially minoritized respondents and less among more conservative, less highly educated, and/or white respondents. Additionally, the latter set of studies suggests that those closest to the population disadvantaged by an algorithm—lower income individuals in the vignette presented in this study—are more likely to withdraw their support than higher income individuals who stand to be less affected by algorithmic biases.

Experimental Design

Status Quo Methods and School Context

We investigate how a key stakeholder group in K-12 schools — parents — views the fairness of school districts using algorithms as a method to identify “high-need” students. In our vignette-based design (Figure 1 shows the survey flow), respondents read a vignette about COVID-19’s impact on learning loss. The vignette then describes a district that has decided to give some students tutors to help but lacks the resources to provide a tutor to all students in need. Thus, schools need a method to determine which students receive a tutor.

In the vignette, we indicate that schools in the district currently use a particular status quo method to allocate tutors, but the district is considering switching to an algorithm that predicts whether a student will need to repeat ninth grade. The vignette randomizes two components in a 4 × 3 design. First, and of primary interest, is the status quo method that algorithms would replace. Respondents are randomized to one of four status quo conditions: counselor judgment, a weighted lottery, parent requests, or a simple bureaucratic rule. This design tests our prediction that the perceived fairness of algorithms is shaped by the perceived fairness of the other allocation method being replaced.

To enable meaningful comparison, all methods that we present — both algorithmic and non-algorithmic — consider the same two general factors: family financial need and student academic performance. By keeping these factors as consistent as possible across methods, our design focuses respondents’ attention on procedural

differences in how each method processes this information: through operationalization in simple bureaucratic rules, as predictors in an outcome-focused algorithm, or through the human judgment of parents or counselors. However, while we strive for consistency in these factors, the fundamental nature of each method requires different levels of precision in how factors are defined. For instance, the simple rule specifies a family income cut-off, whereas the counselor judgment condition more broadly instructs counselors to consider family financial need—a difference that reflects the inherent contrast between rule- and human judgment-based approaches.

The vignette's second randomized component varies the school district's ethnoracial composition, addressing the possibility that respondents might have varying prototypical school districts in mind. We randomly present the district as either majority Black and Hispanic, majority white, or integrated. In our main analyses, we average across these ethnoracial composition conditions. Figure A1 and Table A1 in Appendix A contain the exact vignette wording.

Information about Algorithmic Bias

After completing an initial evaluation of the algorithm's relative fairness based on the preceding initial vignette, respondents read an update about bias in the algorithm. This helps us investigate whether and for whom information about bias informs evaluations of algorithms. The update is grounded in real-world challenges in algorithms trained on student data, which may contain incomplete records for students who frequently change schools—a problem that disproportionately affects lower income families (Welsh 2017):

The school district decided to use a predictive model. However, a year into using it, the district noticed an alarming pattern. The model worked fine for students who had been in the district since elementary school. However, for students whose families moved around a lot, the model incorrectly rated them as low need. That was because the model had no test score data or grades from the students' old districts. The students who moved around a lot often came from lower-income families. And the model never recommended that these students receive a tutor.

Dependent Variables

Our main dependent variable of interest is the respondent's binary answer to the question, "Which method for deciding which students get tutors is fairer?" Respondents compared the predictive model to their randomly assigned status quo method, with the two options presented in random order. Based on their answers to this binary question, we prompted an open-ended response where we asked: "Explain why you think [method they rated as fairer] is fairer than [other method]?" Then, we asked respondents to indicate their degree of certainty that their selected method was fairer on a continuous scale.

Finally, after reading the update about algorithmic bias, respondents were then asked, "With that update in mind, which method should the school district use to select which students get tutors?" Worth noting is that this final question asks

respondents to make an overall evaluation of which allocation method the district should use in light of the bias in the algorithm. Although we anticipate that these overall evaluations heavily weight the perceived fairness of each method given the focus of the previous questions and the update about algorithmic bias, respondents may also form preferences based on other factors such as perceived efficiency or expected personal benefit.

Heterogeneous Effects/Moderators

We explore heterogeneity in respondents' perceptions across several demographic and ideological attributes, all measured in the panel before randomization to the vignette. The survey provider added a screening question measuring respondents' parenting status.

We focus on two main attributes. First is the respondent's educational attainment as a measure of SES. Second is their political ideology (a 7-point scale ranging from extremely conservative to extremely liberal, with the center being moderate). In the online supplement, we show the robustness of the results to household income as the measure of SES and partisan identification (i.e., Republican vs. Democrat vs. Independent) as the measure of political orientation. We also present attitudes by race/ethnicity in the online supplement.

Data

Our survey experiment was administered by NORC and fielded on the AmeriSpeak panel through the Time-Sharing Experiments in the Social Sciences (TESS) program. AmeriSpeak is a probability-based panel where panelists are recruited through mail, telephone, and in-person field outreach. NORC makes special efforts to increase sample representativeness through supplementing U.S. Postal Service address lists with addresses that enhance sample coverage for rural and low-income households, and in-person field outreach that enhances sample representativeness of young adults, lower SES households, non-Internet households, and other hard-to-reach groups. From this panel, to help with the study's statistical power to detect differences between parents, NORC over-sampled current parents of K-12 students, providing survey weights that adjust for this over-sampling.

The survey was fielded in November 2021, when discourse around K-12 schools began to shift from debates about COVID-related closures to debates about recovery efforts. The final sample was comprised of $N = 5,606$ respondents who fall into three mutually exclusive categories: never or not yet parent ($N = 1,234$; 22 percent of the sample), current K-12 parent ($N = 2,665$; 48 percent of the sample), and parent but not K-12 (children younger or older) ($N = 1,713$; 30 percent of the sample). For our main results, we focus on the $N = 4,378$ current or ever parents to probe the views of the stakeholders who comprise the most relevant audience for algorithms in K-12 schools. The study was pre-registered at the following link: <https://osf.io/fjb56/>.

Figure 2 shows the demographic attributes of the parent respondents. All descriptive and inferential analyses and figures are re-weighted using either the

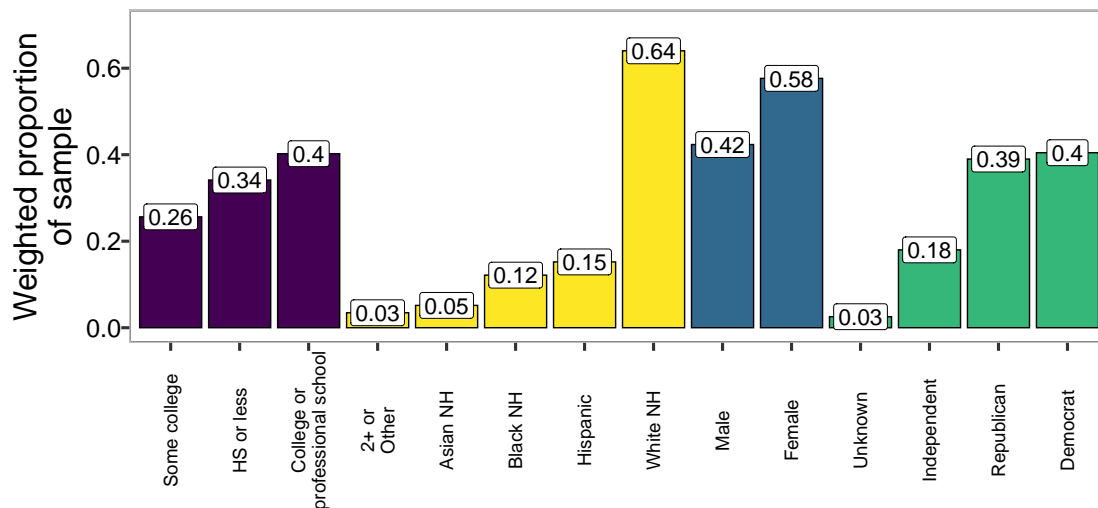


Figure 2: Sample composition: current or ever K-12 parents.

general population survey weight or the parent oversample weight. Important, for the forms of heterogeneity we examine, the sample demographics broadly match the U.S. population distribution. Notably, we have a nationally representative sample of politically conservative respondents, which stands in contrast to the skewed political composition of convenience samples such as Amazon Mechanical Turk commonly used in past studies of attitudes toward algorithms.

Analytic approach

Quantitative-dependent variables

We first examine whether respondents' fairness ratings differ significantly between allocation methods, conducting two-tailed tests of differences-in-proportions separately for each randomly assigned status quo method. Then, we estimate the relationship between respondent attributes and the method that respondents rated as fairer, again separately for each status quo method. We use logistic regression where the dependent variable is whether a respondent rated the status quo method as fairer and the predictor is the respondent's value for a focal attribute (e.g., educational attainment). Finally, among respondents who initially rated the algorithm as fairer, we examine which groups were mostly likely to recommend against using an algorithm after reading the update about algorithmic bias. All analyses are weighted to account for the national sample and parent oversample.

Large language model-assisted coding of open-ended qualitative responses

We combined manual human coding with machine learning to analyze respondents' open-ended explanations of why they saw their chosen allocation method as fairer.

The two authors each manually coded a separate random sample of open-ended responses in the parent analytic sample ($N = 422$ total; $N = 211$ per coder). The random sample was stratified along three dimensions: the status quo condition to which the respondent was randomized, whether the respondent rated the algorithm as more fair, and whether the respondent supported using the algorithm after reading the informational update about algorithmic bias (with the third stratifying factor kept blind to us as coders). We first read a small subset of free responses to identify recurring themes in the reasons offered. Based on this initial review of responses, we developed a coding scheme that coded the (1) salience and (2) normative valence (positive or negative) of two attributes of allocation methods: their degree of targeting and whether they drew on impersonal, rather than personal, knowledge. We manually coded all open-ended answers in the stratified random sample following this coding scheme.

We then extended the coding scheme to the remaining survey responses using a large language model (LLM). LLMs are AI systems that process and generate human language after being trained on vast amounts of text. Recent research has shown that when provided effective instructions (called a “prompt”), LLMs can code texts with a high degree of accuracy even with limited specialized task-specific training (Chae and Davidson 2023; Törnberg 2024).

We randomly split our human-coded responses 80-20: we used 80 percent of the coded responses to iteratively test and refine a detailed prompt that conveyed our coding scheme and reserved 20 percent to evaluate the model’s performance on new observations after we refined the prompt. We tested prompts on several LLMs, including GPT-4o, GPT-4o mini, Claude 3 Opus, and Claude 3.5 Sonnet. Overall, we found Claude 3.5 Sonnet coded the survey responses most accurately, with the best performing prompt achieving accuracies of 86–91 percent and weighted F1 scores of 0.87–0.93 on the validation set across the coding tasks (see Appendix B for more complete predictive performance information). The full prompt that we used is presented in the online supplement. We applied this prompt to code all remaining uncoded responses in our data set ($N = 3,890$; excludes empty, fully numeric, and nonsensical responses).

Results

Main Patterns and Cleavages

A majority of parents think the predictive algorithm is a fairer way to ascertain student need than each status quo method (Figure 3 and Table S1 in the online supplement; see Table S2 in the online supplement showing demographic balance across the conditions). The lottery was perceived as the least fair, with 81 percent of parents thinking the algorithm was fairer and with respondents randomized to that status quo method also rating the algorithm highest on the continuous measure of fairness. As we discuss later in our analysis of qualitative data, the majority of parents thought algorithms were fairer because they were concerned that a lottery would waste resources through its lack of targeting. We see no differences in these fairness ratings across the school racial/ethnic contexts respondents were assigned to (see Table S3 in online supplement).

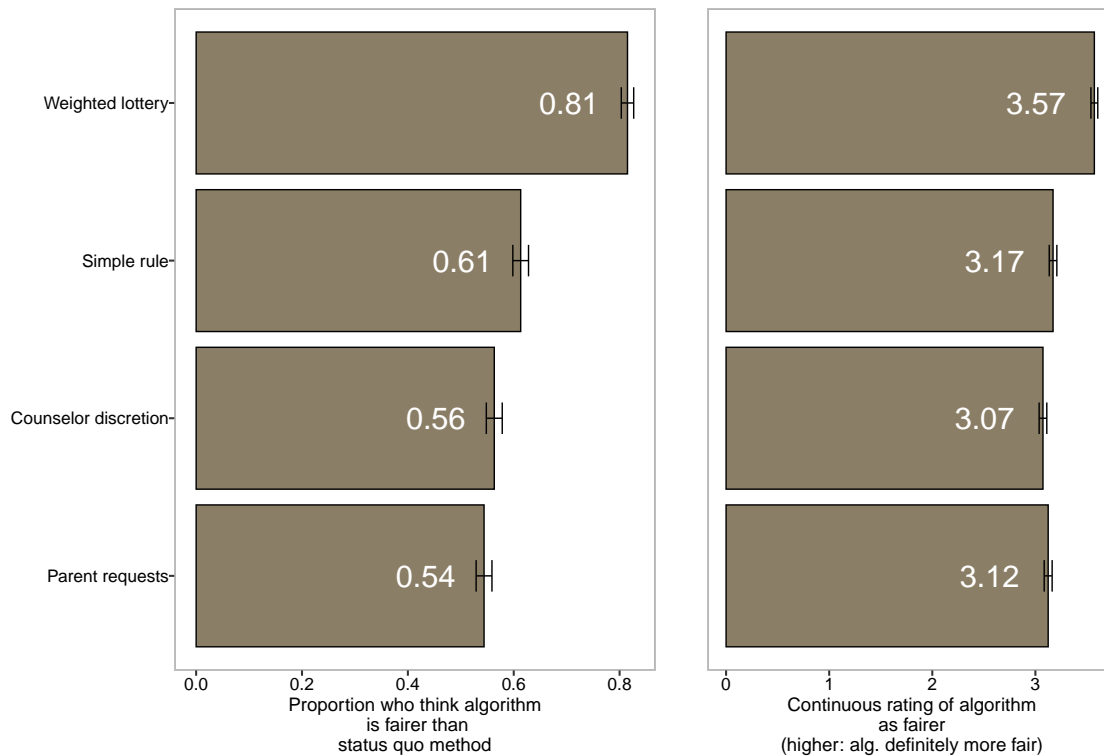


Figure 3: Proportion of parents who think the algorithm is fairer than each status quo method and mean fairness ratings (weighted): The continuous rating has a value of 1 = status quo method definitely more fair to 5 = algorithm definitely more fair, with higher values showing more confidence in the algorithm's fairness.

Yet the overall perception of algorithms as the fairest way to allocate tutors conceals significant heterogeneity among parents for some methods. Although we pre-registered that we expected a parent's race/ethnicity to shape their views of fairness, our results show few consistent patterns along that dimension (Table S4 and Figure S2 in the online supplement). Instead, we see that two background attributes—parent educational attainment and parent political ideology—are more significant sources of cleavages in fairness perceptions (Figure 4; Tables S5 and S7 in the online supplement).

In contrast to our initial expectations based on work on opportunity hoarding by higher SES parents in K-12 education (Calarco 2018; Lareau 2000, 2003), higher SES parents are substantially more likely than lower SES parents to rate algorithms as fairer than parent requests and counselor discretion. While 70 percent of parents with a college degree or higher rate algorithms as fairer than parent requests, only 40 percent of parents with a high school (HS) degree or less do. Similarly, more highly educated parents are more likely to think algorithms are fairer than counselors compared to less highly educated parents. Table S6 and Figure S3 in the online supplement show that these results are robust for measuring SES using income. While we hypothesized that lower SES parents might view parent requests and

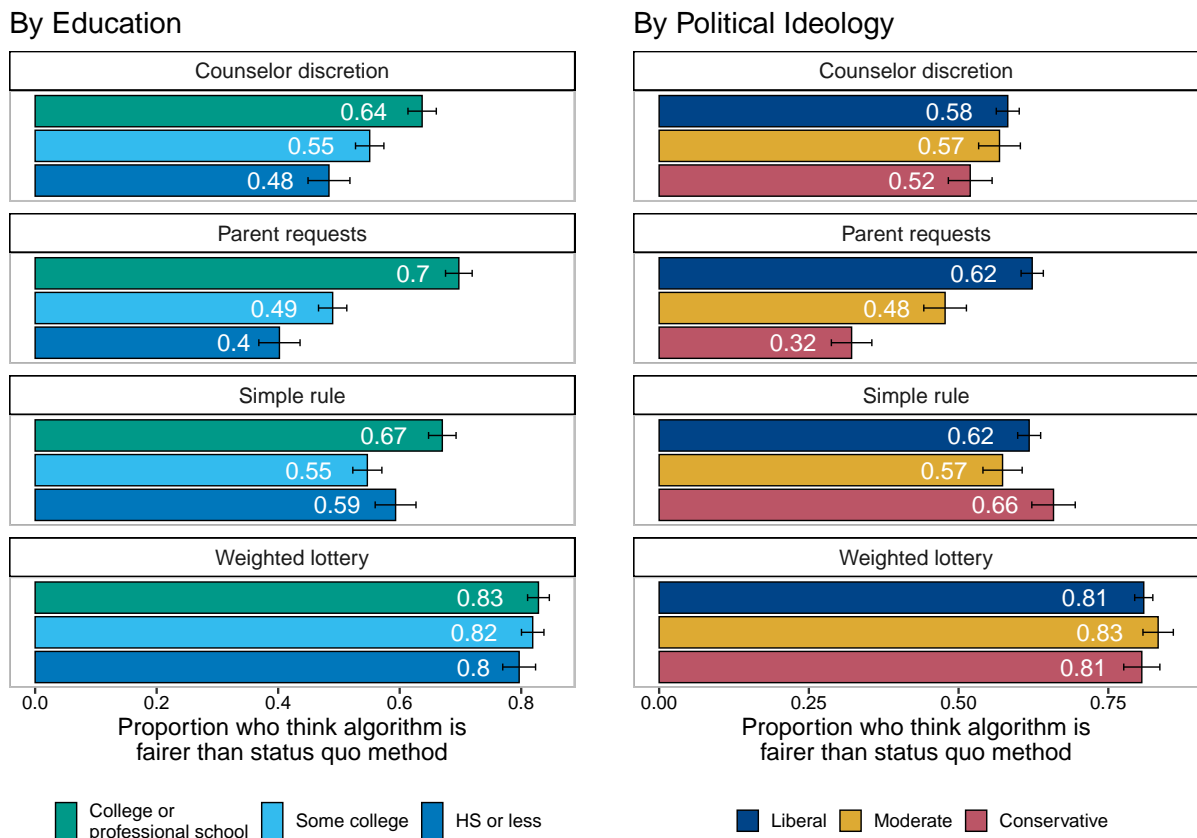


Figure 4: Proportion of parents who think the algorithm is fairer by educational attainment and political ideology (weighted).

counselor discretion as unfair because those parents are more attuned to how such approaches might favor higher SES parents, we see the opposite pattern. Lower SES parents instead view counselor discretion and parent requests as fairer despite research showing that the methods disadvantage these same parents.

Consistent with literature on political polarization in attitudes toward science and skepticism of elite authority, politically conservative parents are far less likely than liberal parents to rate algorithms as fairer than parent requests, with this result robust to measures of partisan identification (Table S8 and Figure S4 in the online supplement). While only 32 percent of conservative parents view algorithms as fairer than parent requests, 62 percent of liberal parents do. Unlike with SES, however, we do not observe polarization around counselor discretion, the other more personal method. Combined, these findings suggest that conservative parents particularly value decentralized personal expertise.

We find less divergence between parents in views of lotteries and rules, aligning with our expectations. Lotteries, in particular, were consistently rated as less fair than algorithms by a large majority of respondents across groups. Table S9 in the online supplement shows the results among parents are similar to those within the

general population sample; while parents are more likely than non-parents to view parent requests as fair, all groups rated the algorithm as fairer than lotteries and simple rules.

Mechanisms: Which Attributes of Selection Methods Are Salient and Valued or Devalued?

The previous results show cleavages in the perceived fairness of algorithms that are more or less sharp depending on what status quo method algorithms are contrasted with. Nearly all respondents view algorithms as fairer than a lottery. Meanwhile, cleavages emerge among parents when they compare algorithms to other selection methods such as parent requests and counselor discretion.

These main quantitative results leave unanswered why parents rate allocation methods as they do. From systematic coding of parents' open-ended explanations, we identify two key attributes parents regularly focus on to assess fairness: (1) how finely a method targets tutors based on student need and (2) whether a method draws more on impersonal or personal forms of knowledge about students. In this section, we examine when these attributes — targeting and impersonality — are most salient to parents judging the relative fairness of methods, as well as whether these attributes are thought to enhance or diminish a method's perceived fairness.

Focusing first on targeting, we find that the primary driver of respondents' dislike for lotteries is a belief that the method insufficiently distinguishes between students based on need. As one respondent (white non-Hispanic, moderate, high school (HS) or less) explained, the algorithm was fair "because picking a lottery may miss a student that really needs it." A broad cross-section of respondents shared similar views, perceiving allocating tutors based on luck as arbitrary and inefficient. For instance, one respondent (multiracial, liberal, and some college) noted that "The predictive model would select the students most likely to benefit. The lottery would reward student's luck, if selected. Theoretically, the lottery could assign a tutor to a student not needing one." Another respondent (white non-Hispanic, conservative, college or professional school) similarly observed, "The lottery randomly selects students, some of which won't even graduate. The model looks at those that will [sic] benefit and make better use of the tutor."

Parent concerns about targeting also contributed to their views of algorithms as fairer than a simple rule. Parents reasoned that the algorithm was fairer because they believed it took more information into account. For instance, one parent (white non-Hispanic, conservative, college or professional school) argued: "Test scores alone do not paint a complete picture of (sic) a student. The predictive model seems more in-depth. It analyzes the data over a longer period of time. Students are more than just one test score, and their family financial situation. The predictive model looks at multiple data points in order to consider which students would benefit most from this individualized support." Others objected to the way that family income played a large near-deterministic role in the simple rule, with one respondent (Black non-Hispanic, liberal, college or professional school) in this category arguing: "A student's ability to excel in the classroom should be determined by their exams,

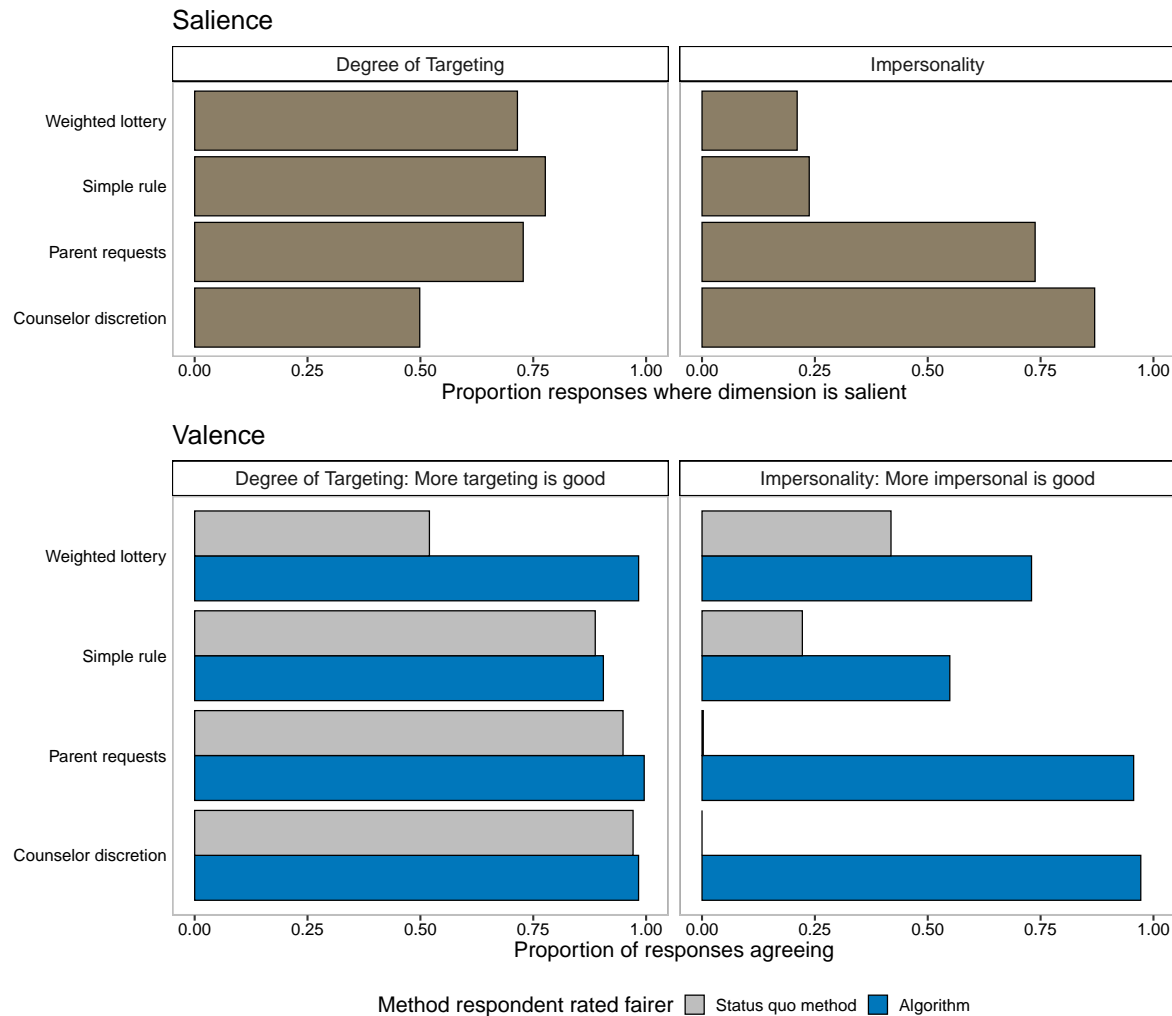


Figure 5: Saliency and valence of targeting and impersonality in open-ended responses comparing algorithms and status quo methods. Proportion agreeing more targeting/impersonal is good is calculated among those for whom the dimension was salient and who gave a response with a clear salience (either positive or negative).

test scores, attendance and attention span. Their family income should not be a determining factor.”

Zooming out to the full qualitative sample, Figure 5 shows nearly uniformly high support among parents for the idea that greater targeting is good, with the attribute made most salient when the algorithm is contrasted with a lottery.

In contrast, when we turn to impersonality and what form of knowledge should guide assessments of need, we observe sharp cleavages between parents who saw parent requests or counselor discretion as fairer—who earlier results showed tend to be more conservative and have lower educational attainment and income—and those who saw algorithms as fairer. Figure 5 shows that while nearly every respondent who rated algorithms as fairer thought impersonal knowledge enhanced

Table 1: Contrasting views about the fairness of algorithms versus parent requests.

Quote	Respondent Demographics	Rated as Fairer
"Sometimes the students who need the most do not have parents asking for help. Data needs to be put in place to make sure [sic] students who truly need the help get it!"	White non-Hispanic, Liberal, College or professional school	Algorithm
"The algorithm is a neutral party, it is looking purely at data and statistics and doesn't care who the student is, what color they are, etc., a parent deciding is too personal, a parent might want their kid to get a tutor when the kid doesn't really need one or visa versa."	White non-Hispanic, Liberal, High school or less	Algorithm
"If a parent sees that their child struggles, they should be able to request additional help vs. data and statistics determining who gets additional help. It is fairer when tutors are available by request of a parent or a guardian."	White non-Hispanic, Conservative, College or professional school	Parent requests
"If the district only has a limited amount of money to spend on tutors than [sic] it should go to those families who see the need and reach out to the schools. The act of reaching out shows parental involvement at a level that will most likely ensure the tutoring won't be wasted on students and families not committed to the students success."	White non-Hispanic, Conservative, Some college	Parent requests

fairness, nearly every respondent who rated parent requests or counselor discretion as fairer thought impersonal knowledge reduced fairness.

To illustrate in the case of parent requests, Table 1 shows the responses of four parents who all highlighted the impersonality of the algorithm and/or the personal nature of parent knowledge in their open-ended explanations. The first two parents, who saw the algorithm as more fair, attach a negative normative valence to personal parent knowledge ("a parent deciding is too personal") and worry that personal knowledge can be a source of inequality. They see the impersonal nature of algorithmic knowledge as a useful tool for combating inequality. In contrast, the latter two parents attach a positive normative valence to personal parental knowledge. They assert that parental knowledge is fairer as a basis for claims-making or suggest that parent involvement itself predicts tutoring having a more beneficial impact on the student. These normative valences undergird parents' perceptions of the relative fairness of the algorithm and parent requests.

In the case of counselor discretion, we similarly see that the personal nature of counselor knowledge is salient to all parents, but parents attach different normative valences to that knowledge. Table 2 shows how some parents worry that a counselor's personal perspective and feelings will "cloud their judgment" and engender bias. Others instead see counselors knowing students and families with their "innate sense" as a source of more fine-grained, contextualized, and accurate information about student need.

Table 2: Contrasting views about the fairness of algorithms versus counselor discretion.

Quote	Respondent Demographics	Rated following as Fairer
"It takes away the personal perspective. A person's feelings about a kid won't cloud their judgement on if they get additional help or not."	Hispanic, Liberal, Some college	Algorithm
"The counselor can have personal feelings toward a student that may prevent students who really need tutors from receiving them. A predictive model allows us to look at factual data and form a logical conclusion."	White non-Hispanic, Liberal, Some college	Algorithm
"People have an innate sense and, while they could certainly use a computer program to analyze data, their experience and simply asking the child how they are doing will probably trump data-base knowledge in most every school setting."	Hispanic, Conservative, College or professional school	Counselor discretion
"The predictive model doesn't SEE what's going on. A person does. They know the people, they know their situations, their behaviors, and sometimes, what's actually going on in their homes. They can make the decisions based on what they observe..."	White non-Hispanic, Conservative, HS or less	Counselor discretion

Quantitative analyses show that, beyond these four examples, placing normative value in the personal knowledge of counselors or parents serves as a key mechanism that explains why lower SES, politically conservative parents view these methods as fairer than algorithms. Table S10, focused on these two status quo methods, shows that these groups, along with Black and Hispanic parents, are significantly more likely to see more impersonal forms of knowledge in a negative light.

Which parents recommend against the algorithm after an update about algorithmic bias?

The qualitative results show that two factors shape parents' evaluations of algorithmic fairness: the importance of targeting resources based on student need (nearly universally endorsed among parents) and whether knowledge of that need should come from data or the personal expertise of parents or counselors. The update parents then read about algorithmic bias, which notes how the algorithm fails to flag low-income students who move around a lot as needing help, challenges one of the core stated advantages of the algorithm: its accurate assessment of student need. Which parents, despite reading about this bias, recommend that the district use the algorithm?

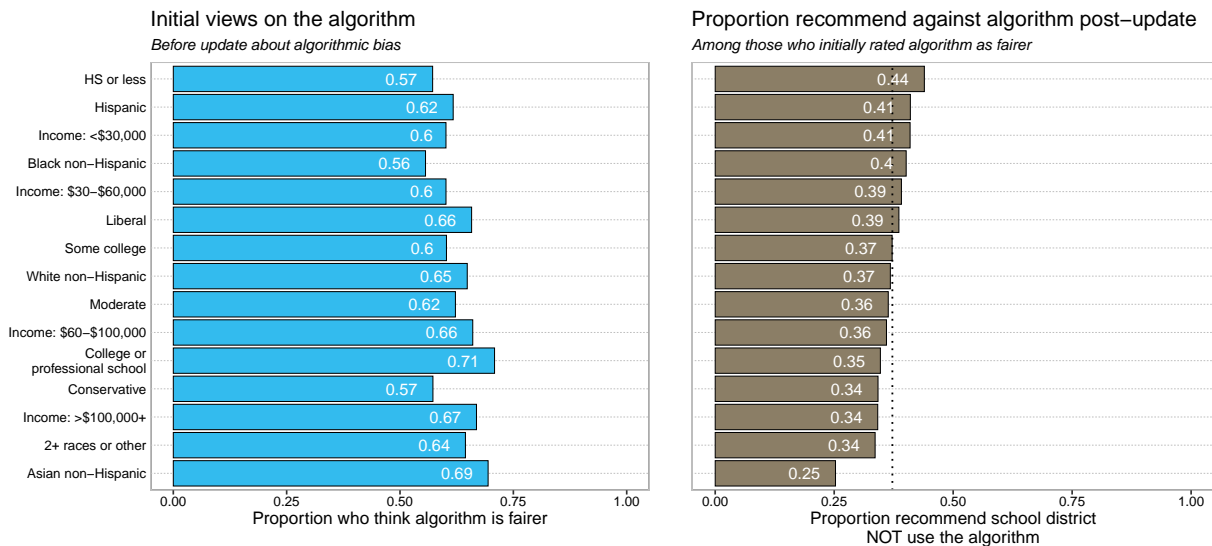


Figure 6: Which groups are most likely to recommend against the algorithm's use after reading the informational update about algorithmic bias? The left panel displays the proportion of respondents who initially deemed the algorithm fairer before the update. The right panel shows the proportion of respondents who initially rated the algorithm as fairer who then recommend against its use following the update about bias. The dotted line indicates the overall proportion of respondents who oppose the algorithm's use.

Overall, 63 percent of parents initially rated the algorithm as the fairer method compared to the status quo. After reading the update about bias, 43 percent of parents recommended that the district use the algorithm. Among those who initially rated the algorithm as fairer, 37 percent recommended against its use after reading the update.

Figure 6 and Table 3 disaggregate these views across groups. After reading the update, parents who initially saw the algorithm as fairer varied in their likelihood of supporting the algorithms' use. Those with incomes under \$30,000—the group closest to the algorithm's failure for low-income students—were among those most likely to oppose implementation. We also find that parents with a high school level of education or less were more likely to oppose the algorithm than parents with a college degree or more, contradicting our initial expectation that parents with more education might be more concerned about bias. This result is instead consistent with social groups who are more proximate to those whom an algorithm would disadvantage being more likely to contest its use.

Furthermore, we find that among those who initially rated the algorithm as fairer, parents identifying as Asian non-Hispanic, multiracial, or another category were more likely to recommend its use after the update than other parents. In particular, Asian parents initially rated algorithms as fairer than did other ethnoracial groups and show among the highest support for implementation after the update, perhaps informed by broader political debates about anti-Asian discrimination in U.S. education that were highly salient when this survey was fielded (Lee 2021). Meanwhile, Hispanic and Black non-Hispanic parents' support for algorithms after the update was comparable to that of white non-Hispanic parents. Finally, in

Table 3: Among respondents who rated an algorithm as fairer, who recommends against an algorithm's use after learning about algorithmic bias? Positive = more likely to turn against algorithms.

	Dependent Variable: Recommends against Algorithm			
	Race/Ethnicity	Educ.	Income	Ideology
Multiracial or other (ref: non-Hispanic white)	−0.111 (0.050) $p = 0.027^*$			
Asian (ref: non-Hispanic white)	−0.089 (0.040) $p = 0.028^*$			
Black non-Hispanic (ref: non-Hispanic white)	0.036 (0.030) $p = 0.236$			
Hispanic (ref: non-Hispanic white)	0.047 (0.026) $p = 0.074$			
High school or less (ref: College+)		0.099 (0.021) $p = 0.00001^{***}$		
Some college (ref: College+)		0.039 (0.023) $p = 0.093$		
Income: <\$30,000 (ref: Income >\$100,000)			0.077 (0.026) $p = 0.004^{**}$	
Income: \$30-\$60,000 (ref: Income >\$100,000)			0.025 (0.026) $p = 0.326$	
Income: \$60-\$100,000 (ref: Income >\$100,000)			−0.011 (0.025) $p = 0.668$	
Moderate (ref: liberal)				0.004 (0.025) $p = 0.881$
Conservative (ref: liberal)				−0.030 (0.026) $p = 0.237$
Constant	0.376 (0.011) $p = 0.000^{***}$	0.338 (0.014) $p = 0.000^{***}$	0.358 (0.017) $p = 0.000^{***}$	0.387 (0.012) $p = 0.000^{***}$
Observations	2,791	2,791	2,791	2,689

Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

contrast to our initial predictions based on research showing that liberals are more sensitive to between-group inequalities, we find no significant differences in the rate at which parents recommend against the algorithm by political ideology after the update.

Discussion

Predictive algorithms exemplify the ongoing shift toward more impersonal, quantified, and technocratic allocation procedures (Porter 1996; Burrell and Fourcade 2021; Levy et al. 2021; Berman 2022). Although some people associate these characteristics with fairer allocation, others find them objectionable. We show that contrasting algorithms with traditional allocation methods helps clarify the factors that drive public perceptions of the fairness of using predictive models to distribute society's benefits and burdens.

Using the case of algorithms to assess student need in K-12 schools, we found that most parents judged algorithms as fairer than conventional methods. However, this overall pattern concealed significant divides in perceptions. Although we did not observe major differences by race/ethnicity, lower SES and politically conservative parents rated counselor discretion (SES only) and parent requests (both groups) as fairer. Open-ended responses revealed mechanisms behind these evaluations. Nearly all parents valued more precise targeting, but parents disagreed on the fairness of assessing need via an algorithm's impersonal approach versus a parent or counselor's personal approach. Lower SES conservative parents favored personal knowledge, whereas higher SES liberal respondents preferred the impersonal knowledge of algorithms, which they saw as more reliable for identifying high-need students who lack strong personal advocates. Thus, judgments of algorithms in schools reflect broader cleavages in trust in technocratic expertise. Other research shows how reformers advocating for the adoption of early warning systems emphasize the efficiency gains of using data for targeting and having that data draw on impersonal knowledge about students (Trinidad 2024). Our findings show how the fairness perceptions of higher SES liberal parents mirror the reformers' logic but that the logic is far from universal among all parents.

Our findings also highlight a disconnect between how scholars or elites assess the fairness of allocation approaches and how the public perceives them. Although some research has explored subtle ways lotteries can create inequalities (Frankenberg et al. 2011; Skinner 2014; Berends 2015), lotteries are commonly elevated in scholarly discourses as a fair, inequality-reducing method because they tie outcomes to luck rather than individual circumstance (Goodwin 1992; Fang and Casadevall 2016; Vong 2020; Wang et al. 2024). However, parents in our study do not share this view, consistently rating lotteries as substantially less fair than other approaches across all demographic groups. Moreover, despite research warning that linking assistance to parental advocacy likely exacerbates inequality (Lareau 2000, 2003; Lewis-McCoy 2014; Calarco 2020; Cartwright 2022), this approach enjoyed high fairness ratings. This was especially true among lower SES parents, the very group the literature predicts is most likely to be disadvantaged by such a system. These results demonstrate the need for empirical studies of perceptions, as they may

reveal discrepancies between theoretical portrayals of allocation methods and their reception by stakeholders on the ground.

Our finding that greater targeting is nearly universally valued aligns with accounts of the deep institutionalization of efficiency-based approaches in policy-making (Berman 2022), indicating the take-up of such logics among the general public who worry about wasting resources. However, our finding that impersonality polarizes the public challenges the widespread assumption that quantitative, impersonal approaches are broadly seen as more legitimate due to their associations with objectivity. The divide we observe — higher SES liberal respondents viewed impersonality as an asset for reducing bias, whereas lower SES conservative respondents saw it as a liability that misses crucial information about the whole person — mirrors broader cleavages in American public attitudes toward science. Gauchat (2011) posits that one source of lower trust in science is institutional alienation, or feeling “discontent with abstract bureaucratic systems” (p. 755). Those with lower trust view scientific knowledge as only one credible way of knowing alongside commonsense and religious knowledge. Conservative parents’ open-ended responses expressed concern that algorithms overlook crucial information, such as a parent’s engagement level or receptivity to help, which they believe is essential for the effectiveness of assistance. These responses reveal how attitudes toward major institutions such as science shape or guide perceptions of related technologies and methods.

Scholars warn that algorithmic decision-making may be disproportionately applied to the least powerful segments of society (Madden et al. 2017; Eubanks 2018; Barabas et al. 2020; Rona-Tas 2020). This study points to a mechanism that could sustain such a dynamic. We find uneven responsiveness to information about bias against vulnerable groups: lower income parents and parents with lower levels of education, groups more closely connected to those harmed by the algorithm’s bias, opposed algorithms more than did higher income and more highly educated respondents. This suggests that information about algorithmic bias may be less persuasive for more powerful members of society, who tend to be less exposed to its harms, possibly perpetuating the application of algorithms to marginalized groups.

From a policy perspective, this article is among the first to examine public attitudes toward a widespread but often hidden technology: early warning prediction models that identify students at risk of adverse academic outcomes (U.S. Department of Education 2016; Feathers 2023; Perdomo et al. 2023). The political polarization we find over algorithms versus parent requests suggests potential spillover from polarization on issues such as race, gender, and sexuality in school curricula, which have galvanized parents along political lines and intensified debates about “parent rights” in the United States (Baldwin Clark 2023; Filimon and Ivănescu 2023; Kelly 2023). Indeed, this study was fielded at a time when parent rights discourse was especially salient (Figure S1 in the online supplement shows the overlap with the fall 2021 midterm elections). One limitation of the present survey-based study is that it leaves the question of whether these expressed attitudes translate into real-world policy inputs. Future research should investigate whether the attitudes uncovered here align with positions taken in public forums like school board meetings.

Future research should also compare perceptions of fairness against the actual distributive or procedural fairness of different allocation methods. Unlike other areas of attitudinal research where perceptions can be juxtaposed against widely accepted social science indicators or stylized facts, such as established rates of intergenerational mobility or political polarization (Levendusky and Malhotra 2016; Cheng and Wen 2019), there is a more limited empirical research base on the fairness implications of algorithms. Not only are there numerous competing, sometimes mutually incompatible notions of fairness in the literature (Friedler, Scheidegger, and Venkatasubramanian 2021; Corbett-Davies et al. 2023), the impacts of using algorithms rather than other allocation methods can vary considerably across decision-making domains and local contexts (Albright 2023; Eiermann 2024).

Moving forward, future research could pair studies of on-the-ground fairness impacts in specific contexts with perceptual surveys of stakeholders and the public. Doing so in the context of schools, for instance, could reveal if the present results indicate that parents are overly optimistic or pessimistic about algorithms relative to status quo methods. Exploring when and where perceptions diverge from the outcomes of algorithms can help both refine operationalizations of algorithmic fairness (e.g., some parents viewing it as essential to incorporate a human component) and reveal how the framing of allocation methods can hide resultant inequalities.

Methodologically, this work contributes to the burgeoning interdisciplinary literature on lay perceptions of algorithms by asking respondents to evaluate algorithms against a range of real-world alternatives (Starke et al. 2022). This approach moves beyond the human versus machine dichotomy prevalent in past research. Our findings demonstrate that support for algorithms and the factors driving fairness evaluations vary depending on the counterfactual presented. Although many critiques and analyses of algorithms typically contrast them with unfettered human discretion or assess their properties in isolation, our study highlights the value of acknowledging that stakeholders with firsthand experiences of institutions adopting these technologies may evaluate algorithms in relation to other status quo alternatives they have encountered.

Although one strength of our design was its comparison to real-world allocation methods, future research could further refine how to probe the impacts of learning about algorithmic bias. Our study used a clearly marked “update” and then immediately asked respondents about their views. Although this likely shifted all groups toward more negative views of using an algorithm, not biasing our interpretation of between-group differences, future research could investigate whether more subtle ways of signaling bias also sway views. In addition, the update about bias only highlighted biases in the algorithm, not biases in each of the status quo methods. Future research could explore how presenting information about biases in both status quo and algorithmic methods affects views.

This study focused on probing the perceived fairness of allocation methods and parents’ recommendations for which method to use. However, the study design may have influenced the results. Participants first answered questions about algorithmic fairness and received information about algorithmic bias before being asked which method districts should use. This sequencing likely primed respondents to prioritize fairness considerations in their recommendations. Future

research could explore how highlighting other attributes, such as an allocation method's "efficiency" or ease of implementation, influences preferences.

Future research should also explore how perceptions of fairness, and the importance placed on targeting and impersonality, differ across the various kinds of algorithms used in education, such as those used to predict academic success or self-harm, and in different institutional and national contexts where other status quo approaches might prevail (Bergman, Kopko, and Rodriguez 2021; Collins et al. 2021; Engler, 2021). A promising direction could be to investigate how the social meaning of the allocated resource — be that a kidney, housing assistance, or a scholarship — shapes the perceived moral appropriateness of technocratic targeting versus universalist allocation. In the school context, scholars could examine how attitudes vary when the allocated good is a highly sought-after competitive resource for getting ahead, like spots in selective magnet schools, as opposed to the compensatory resource for "catching up" that we study. Would liberal, high-SES parents maintain their support for algorithms if they were used to assess merit rather than need in such contexts?

Amidst calls for greater stakeholder engagement in the design and implementation of decision-making algorithms, our study sheds light on the potential contours of stakeholder responses. In the case of algorithms to predict need for help in K-12 schools, we find that overall acceptance of algorithms masks significant divisions that align with broader social and political cleavages around trust in science, expertise, and public institutions. Our findings also suggest that the concerns about inequality and bias that animate scholarly and elite media critiques of algorithmic decision-making resonate unevenly with the public. As algorithms increasingly make critical decisions across many domains of social life, understanding the reasons people embrace, question, or reject algorithms will shed light on factors that could fuel or limit their adoption and influence.

Appendix

Appendix A: Survey Wording

A [school district where 90% of students are white/school district where 90% of students are Black or Latino/school district where 45% of students are white and 45% of students are Black or Latino] is facing COVID-19 learning losses. Although some students are doing fine, others are struggling.

To help struggling students, the district is pairing some ninth-grade students with a tutor who meets with the student multiple times per week during the school day to help the student catch up academically.

The problem: Unfortunately, one-on-one tutors are very expensive and the district only has enough money to provide tutors to 15 percent of the many students who have fallen behind.

How do schools decide which students get tutoring? [randomized to status quo method outlined below]

However, an analytics team within the district has proposed a new method: switching to an algorithm / predictive model. The predictive model would analyze the student records of every student in the district from the past 10 years—such as test scores, grades, attendance, and family financial need—to learn what factors predict whether a student is likely to fail ninth grade. The model would then use what it learned to identify current students who are most likely to need to repeat ninth grade. School counselors would then provide tutors to students the model recommends.

Summary: We want your opinion about how the school district's leadership should decide which students get tutors:

- **How the school district initially gave tutors:** school counselors have [used other method]
- **How the school district could give tutors:** school counselors would use an algorithm or predictive model

Figure A1: Vignette wording

Table A1: Vignette and dependent variable wording.

Category	Order	Wording
<i>Status quo methods in vignette (randomized to one)</i>		
Educator discretion	1	"Initially, guidance counselors have been using their judgment and personal knowledge of students to decide which students to provide with tutors. The district encouraged guidance counselors to weigh students' academic records and family financial need when selecting students."
Lottery	1	"Initially, the school district has been using a lottery. All students were eligible, but the district gave students who demonstrated academic and financial need better odds of being randomly chosen."
Parent requests	1	"Initially, the district has been using parents' requests. The district has encouraged parents to only request tutoring if they believe that their child needs it based on their academic record and the family does not have the financial means to pay for a tutor."
Simple rule	1	"Initially, guidance counselors have been using a test score and family income cutoff set by the school district to decide which students get tutors."
<i>Comprehension check</i>		
Stating in their own words what the predictive model is	2	"Can you explain briefly in your own words what it means for the school counselor to use a predictive model to choose which students get tutors? As a reminder, a predictive model would analyze the student records of every student in the district from the past 10 years—such as test scores, grades, and family financial need—to learn what factors predict whether a student is likely to fail ninth grade to predict which students have the highest need."
<i>Quantitative ratings</i>		
Forced choice between predictive model and other decision-making method	3	"Which method for deciding which students get tutors is more fair?"
Binary choice of efficiency	5	"You've selected (inserts method they chose as more fair) as more fair. The district is pressed for time. Which method do you think would save them the most time in selecting which students get tutors?"
Continuous rating of predictive model	6	"When comparing [inserts other method] to the predictive model, how would you rate how certain you are about which is fairer?" Answer choices: 1 = [Insert other method is definitely more fair], 3 = I'm not sure which is more fair, and 5 = Predictive model is definitely more fair
<i>Qualitative response</i>		
Explanation for why chosen method is more fair	4	"Explain why you think the [inserts method they chose as more fair] is more fair than the [inserts method they said was less fair]"
Question after status update	7	With that update in mind, which method should the school district use to select which students get tutors?
<i>Additional demographics</i>		
Parenting status	8	"Please select which of the following best describes your parenting status: (1) I am a current parent of school-age children (0-18), (2) I am a former parent of school-age children (children 19+), (3) I am not a current or former parent of school-age children"
ZIP code	9	"Please provide your ZIP code"

Appendix B Large Language Model Performance on Coding Open-Ended Survey Responses

Table B1: Performance on identifying the salience of targeting and impersonality in survey responses in the training and validation sets: The training set observations were used to develop the prompt for Claude 3.5 Sonnet (model version from June 20, 2024), whereas the test set observations were reserved for evaluation. Salience was a binary classification task (either salient or not salient).

	Training Set				Validation Set			
	Accuracy	Precision	Recall	F1 Score	Accuracy	Precision	Recall	F1 Score
Targeting	91%	0.90	0.97	0.94	90%	0.94	0.91	0.93
Impersonality	91%	0.96	0.88	0.92	86%	0.94	0.85	0.89

Table B2: Performance on identifying the valence of targeting and impersonality in survey responses in the training and validation sets: The training set observations were used to develop the prompt for Claude 3.5 Sonnet (model version from June 20, 2024), whereas the test set observations were reserved for evaluation. Valence was a multi-class classification task. For each dimension, we pooled observations where the dimension was not salient and where its valence was unclear as we do not distinguish between them in the main analysis presented in the article. Macro F1 score is the simple average of the F1 scores for each class, weighing all classes equally. Weighted F1 score is the average of the F1 scores for each class, weighted by the number of true instances for each class to account for imbalance in the frequency of the classes.

	Training Set			Validation Set		
	Accuracy	Macro F1	Weighted F1	Accuracy	Macro F1	Weighted F1
Targeting	89%	0.82	0.89	88%	0.78	0.87
Impersonality	91%	0.91	0.91	91%	0.91	0.91

References

- Ajunwa, Ifeoma. 2023. *The Quantified Worker: Law and Technology in the Modern Workplace*. Cambridge University Press. <https://doi.org/10.1017/9781316888681>
- Albright, Alex. 2023. "The Hidden Effects of Algorithmic Recommendations." <https://doi.org/10.21034/iwp.104>
- Auyero, Javier. 2012. *Patients of the State: The Politics of Waiting in Argentina*. Duke University Press. <https://doi.org/10.1515/9780822395287>
- Baldwin Clark, LaToya. 2023. "The Critical Racialization of Parents' Rights."
- Bankins, Sarah, Paul Formosa, Yannick Griep, and Deborah Richards. 2022. "AI Decision Making with Dignity? Contrasting Workers' Justice Perceptions of Human and AI Decision Making in a Human Resource Management Context." *Information Systems Frontiers* 24:857–75. <https://doi.org/10.1007/s10796-021-10223-8>
- Barabas, Chelsea, Colin Doyle, J. B. Rubinovitz, and Karthik Dinakar. 2020. "Studying up: reorienting the study of algorithmic fairness around issues of power." Pp. 167–76

- in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20*, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3351095.3372859>
- Benjamin, Ruha. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. John Wiley & Sons.
- Berends, Mark. 2015. "Sociology and School Choice : What We Know After Two Decades of Charter Schools." *Annual Review of Sociology* 41:159–80. <https://doi.org/10.1146/annurev-soc-073014-112340>
- Bergman, Peter, Elizabeth Kopko, and Julio E. Rodriguez. 2021. "A Seven - College Experiment Using Algorithms to Track Students : Impacts and Implications for Equity and Fairness." <https://doi.org/10.3386/w28948>
- Berman, Elizabeth Popp. 2022. *Thinking like an Economist: How Efficiency Replaced Equality in U.S. Public Policy*. Princeton University Press. <https://doi.org/10.1515/9780691226606>
- Binns, Reuben, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 2018. "It's Reducing a Human Being to a Percentage': Perceptions of Justice in Algorithmic Decisions." Pp. 1–14 in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18*. New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.31235/osf.io/9wqxr>
- Blake, Mary Kate. 2020. "Other Duties as Assigned : The Ambiguous Role of the High School Counselor." *Sociology of Education* 93:315–30. <https://doi.org/10.1177/0038040720932563>
- Bouek, Jennifer W. 2023. "The Wait List as Redistributive Policy : Access and Burdens in the Subsidized Childcare System." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 9:76–97. <https://doi.org/10.7758/RSF.2023.9.5.04>
- Bowers, Alex J. 2021. "Early Warning Systems and Indicators of Dropping out of Upper Secondary School: The Emerging Role of Digital Technologies." In *OECD Digital Education Outlook 2021 Pushing the Frontiers with Artificial Intelligence , Blockchain and Robots: Pushing the Frontiers with Artificial Intelligence*. Blockchain and Robots. <https://doi.org/10.1787/c8e57e15-en>
- Brayne, Sarah and Angèle Christin. 2021. "Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts." *Social Problems* 68:608–24. <https://doi.org/10.1093/socpro/spaa004>
- Bueno, Natália S., Felipe Nunes, and Cesar Zucco. 2024. "Benefits by Luck: A Study of Lotteries as a Selection Method for Government Programs." *Comparative Political Studies* p. 00104140241290203. <https://doi.org/10.1177/00104140241290203>
- Burrell, Jenna and Marion Fourcade. 2021. "The Society of Algorithms." *Annual Review of Sociology* 47:213–37. <https://doi.org/10.1146/annurev-soc-090820-020800>
- Calarco, Jessica McCrory. 2011. "'I Need Help !' Social Class and Children 's Help -Seeking in Elementary School." *American Sociological Review* 76:862–82. <https://doi.org/10.1177/0003122411427177>
- Calarco, Jessica McCrory. 2018. *Negotiating Opportunities: How the Middle Class Secures Advantages in School*. Oxford University Press.
- Calarco, Jessica McCrory. 2020. "Avoiding Us versus Them: How Schools 'Dependence on Privileged "Helicopter " Parents Influences Enforcement of Rules." *American Sociological Review* 85:223–46. <https://doi.org/10.1177/0003122420905793>
- Cartwright, Ashleigh. 2022. "A Theory of Racialized Cultural Capital." *Sociological Inquiry* 92:317–40. <https://doi.org/10.1111/soin.12479>

- Chae, Youngjin (YJ) and Thomas Davidson. 2023. "Large Language Models for Text Classification: From Zero - Shot Learning to Instruction-Tuning." <https://doi.org/10.31235/osf.io/sthwk>
- Cheng, Siwei and Fangqi Wen. 2019. "Americans Overestimate the Intergenerational Persistence in Income Ranks." *Proceedings of the National Academy of Sciences* 116:13909–914. <https://doi.org/10.1073/pnas.1814688116>
- Collins, Sara , Park, Jasmine , Reddy, Anisha , Sharifi, Yasamin , and Vance, Amelia. 2021. "The Privacy and Equity Implications of Using Self - Harm Monitoring Technologies."
- Collinson, Robert, Anthony A. DeFusco, John Eric Humphries, Benjamin J. Keys, David C. Phillips, Vincent Reina, Patrick S. Turner, and Winnie van Dijk. 2024. "The Effects of Emergency Rental Assistance During the Pandemic: Evidence from Four Cities." <https://doi.org/10.3386/w32463>
- Corbett-Davies, Sam, Johann D. Gaebler, Hamed Nilforoshan, Ravi Shroff, and Sharad Goel. 2023. "The measure and mismeasure of fairness." *J. Mach. Learn. Res.* 24:312:14730–846.
- Dodge, Jonathan, Q. Vera Liao, Yunfeng Zhang, Rachel K. E. Bellamy, and Casey Dugan. 2019. "Explaining models: an empirical study of how explanations impact fairness judgment." Pp. 275–85 in *Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI '19*. New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3301275.3302310>
- Domina, Thurston, Andrew Penner, and Emily Penner. 2017. "Categorical Inequality: Schools As Sorting Machines." *Annual Review of Sociology* 43:311–30. <https://doi.org/10.1146/annurev-soc-060116-053354>
- Eiermann, Martin. 2024. "Algorithmic Risk Scoring and Welfare State Contact Among US Children." *Sociological Science* 11:707–42. <https://doi.org/10.15195/v11.a26>
- Engler, Alex. 2021. "Enrollment algorithms are contributing to the crises of higher education." <https://policycommons.net/artifacts/4142700/enrollment-algorithms-are-contributing-to-the-crises-of-higher-education/4951896/>
- Espeland, Wendy Nelson and Berit Irene Vannebo. 2007. "Accountability, Quantification, and Law." *Annual Review of Law and Social Science* 3:21–43. <https://doi.org/10.1146/annurev.lawsocsci.2.081805.105908>
- Eubanks, Virginia. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Publishing Group.
- Faber, Jacob William and Jessica Rose Kalbfeld. 2019. "Complaining While Black: Racial Disparities in the Adjudication of Complaints against the Police." *City & Community* 18:1028–1067. <https://doi.org/10.1111/cico.12388>
- Fang, Ferric C. and Arturo Casadevall. 2016. "Research Funding: the Case for a Modified Lottery." *mBio* 7: <https://doi.org/10.1128/mBio.00694-16>
- Feathers, Todd. 2023. "False Alarm: How Wisconsin Uses Race and Income to Label Students 'High Risk'." *The Markup*.
- Filimon, Luiza-Maria and Mihaela Ivănescu. 2023. "Bans, sanctions, and dog-whistles: a review of anti-critical race theory initiatives adopted in the United States since 2020." *Policy Studies* 0:1–22. <https://doi.org/10.1080/01442872.2023.2214088>
- Fourcade, Marion and Kieran Healy. 2017. "Seeing like a market." *Socio-Economic Review* 15:9–29. <https://doi.org/10.31235/osf.io/jzy7k>
- Frankenberg, Erica, Genevieve Siegel-Hawley, and Jia Wang. 2011. "Choice without Equity: Charter School Segregation." *Education Policy Analysis Archives* 19:1–1. <https://doi.org/10.14507/epaa.v19n1.2011>

- Friedler, Sorelle A., Carlos Scheidegger, and Suresh Venkatasubramanian. 2021. "The (Im)possibility of fairness: different value systems require different mechanisms for fair decision making." *Commun. ACM* 64:136–43. <https://doi.org/10.1145/3433949>
- Gast, Melanie Jones. 2021. "'You're Supposed to Help Me': The Perils of Mass Counseling Norms for Working-Class Black Students." *Urban Education* 56:1429–455. <https://doi.org/10.1177/0042085916652178>
- Gauchat, Gordon. 2011. "The cultural authority of science: Public trust and acceptance of organized science." *Public Understanding of Science* 20:751–70. <https://doi.org/10.1177/0963662510365246>
- Gauchat, Gordon. 2012. "Politicization of Science in the Public Sphere: A Study of Public Trust in the United States, 1974 to 2010." *American Sociological Review* 77:167–87. <https://doi.org/10.1177/0003122412438225>
- Gauchat, Gordon. 2015. "The Political Context of Science in the United States: Public Acceptance of Evidence - Based Policy and Science Funding." *Social Forces* 94:723–46. <https://doi.org/10.1093/sf/sov040>
- Gauchat, Gordon W. 2023. "The Legitimacy of Science." *Annual Review of Sociology* 49:263–279. <https://doi.org/10.1146/annurev-soc-030320-035037>
- Glass, Jennifer. 2019. "Why Aren't We Paying Attention? Religion and Politics in Everyday Life." *Sociology of Religion* 80:9–27. <https://doi.org/10.1093/socrel/sry041>
- Goodwin, Barbara. 1992. *Justice by Lottery*. Chicago, IL: University of Chicago Press.
- Grgić-Hlača, Nina, Gabriel Lima, Adrian Weller, and Elissa M. Redmiles. 2022. "Dimensions of Diversity in Human Perceptions of Algorithmic Fairness." In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '22, pp. 1–12, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3551624.3555306>
- Grgic-Hlaca, Nina, Elissa M. Redmiles, Krishna P. Gummadi, and Adrian Weller. 2018. "Human Perceptions of Fairness in Algorithmic Decision Making: A Case Study of Criminal Risk Prediction." Pp. 903–12 in *Proceedings of the 2018 World Wide Web Conference*, WWW '18. Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/3178876.3186138>
- Hamel, Brian T. and Derek E. Holliday. 2024. "Unequal Responsiveness in City Service Delivery: Evidence from 42 Million 311 Calls." *Quarterly Journal of Political Science* 19:243–74. <https://doi.org/10.1561/100.00022089>
- Hart, A., S. K. Gustafson, M. A. Skeans, P. Stock, D. Stewart, B. L. Kasiske, and A. K. Israni. 2017. "OPTN/SRTR 2015 Annual Data Report: Early effects of the new kidney allocation system." *American Journal of Transplantation: Official Journal of the American Society of Transplantation and the American Society of Transplant Surgeons* 17(Suppl 1):543–64. <https://doi.org/10.1111/ajt.14132>
- Heimer, Carol A. 2001. "Cases and Biographies: An Essay on Routinization and the Nature of Comparison." *Annual Review of Sociology* 27:47–76. <https://doi.org/10.1146/annurev.soc.27.1.47>
- Herd, Pamela and Donald P. Moynihan. 2019. *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation. <https://doi.org/10.7758/9781610448789>
- Hirschman, Daniel and Emily Adlin Bosk. 2020. "Standardizing Biases: Selection Devices and the Quantification of Race." *Sociology of Race and Ethnicity* 6:348–64. <https://doi.org/10.1177/2332649219844797>

- Holland, Megan M. 2015. "Trusting Each Other: Student -Counselor Relationships in Diverse High Schools." *Sociology of Education* 88:244–62. <https://doi.org/10.1177/0038040715591347>
- Jain, Vageesh, Lara Schwarz, and Paula Lorgelly. 2021. "A Rapid Review of COVID - 19 Vaccine Prioritization in the U.S.: Alignment between Federal Guidance and State Practice." *International Journal of Environmental Research and Public Health* 18:3483. <https://doi.org/10.3390/ijerph18073483>
- Jakesch, Maurice, Zana Bućinca, Saleema Amershi, and Alexandra Olteanu. 2022. "How Different Groups Prioritize Ethical Values for Responsible AI." In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, pp. 310–23, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533097>
- Johnson, Rebecca Ann and Simone Zhang. 2022. "What is the Bureaucratic Counterfactual? Categorical versus Algorithmic Prioritization in U.S. Social Policy." In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, pp. 1671–682, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533223>
- Joyce, Kelly, Laurel Smith-Doerr, Sharla Alegria, Susan Bell, Taylor Cruz, Steve G. Hoffman, Safiya Umoja Noble, and Benjamin Shestakofsky. 2021. "Toward a Sociology of Artificial Intelligence: A Call for Research on Inequalities and Structural Change." *Socius* 7:2378023121999581. <https://doi.org/10.1177/2378023121999581>
- Katz, Michael B. 2013. *The Undeserving Poor: America's Enduring Confrontation with Poverty: Fully Updated and Revised*. Oxford University Press.
- Kelly, Laura Beth. 2023. "What Do So - Called Critical Race Theory Bans Say?" *Educational Researcher* 52:248–50. <https://doi.org/10.3102/0013189X231159382>
- Kennedy, Ryan P., Philip D. Waggoner, and Matthew M. Ward. 2022. "Trust in Public Policy Algorithms." *The Journal of Politics* 84:1132–148. <https://doi.org/10.1086/716283>
- Kieslich, Kimon, Birte Keller, and Christopher Starke. 2022. "Artificial intelligence ethics by design. Evaluating public perception on the importance of ethical design principles of artificial intelligence." *Big Data & Society* 9:20539517221092956. <https://doi.org/10.1177/20539517221092956>
- Kim, Sarah. 2024. "Top D. C. Public Schools Increase At - Risk Student Enrollment After Policy Change." *DCist*.
- Kiviat, Barbara. 2021. "Which Data Fairly Differentiate? American Views on the Use of Personal Data in Two Market Settings." *Sociological Science* 8:26–47. <https://doi.org/10.15195/v8.a2>
- Kiviat, Barbara. 2023. "The Moral Affordances of Construing People as Cases: How Algorithms and the Data They Depend on Obscure Narrative and Noncomparative Justice." *Sociological Theory* 41:175–200. <https://doi.org/10.1177/07352751231186797>
- Kontokosta, Constantine E. and Boyeong Hong. 2021. "Bias in smart city governance: How socio-spatial disparities in 311 complaint behavior impact the fairness of data-driven decisions." *Sustainable Cities and Society* 64:102503. <https://doi.org/10.1016/j.scs.2020.102503>
- Lamont, Michèle, Stefan Beljean, and Matthew Clair. 2014. "What is missing? Cultural processes and causal pathways to inequality." *Socio-Economic Review* 12:573–608. <https://doi.org/10.1093/ser/mwu011>

- Langer, Markus, Cornelius J. König, and Maria Papathanasiou. 2019. "Highly automated job interviews: Acceptance under the influence of stakes." *International Journal of Selection and Assessment* 27:217–34. <https://doi.org/10.1111/ijsa.12246>
- Lareau, Annette. 2000. *Home Advantage: Social Class and Parental Intervention in Elementary Education*. Rowman & Littlefield.
- Lareau, Annette. 2003. *Unequal Childhoods: Class, Race, and Family Life*. University of California Press.
- Lee, Jennifer. 2021. "Asian Americans , Affirmative Action & the Rise in Anti - Asian Hate." *Daedalus* 150:180–98. https://doi.org/10.1162/daed_a_01854
- Lee, Min Kyung. 2018. "Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management." *Big Data & Society* 5:2053951718756684. <https://doi.org/10.1177/2053951718756684>
- Levendusky, Matthew S. and Neil Malhotra. 2016. "(Mis)perceptions of Partisan Polarization in the American Public." *Public Opinion Quarterly* 80:378–91. <https://doi.org/10.1093/poq/nfv045>
- Levine, Jeremy R. and Carl Gershenson. 2014. "From Political to Material Inequality: Race , Immigration , and Requests for Public Goods." *Sociological Forum* 29:607–27. <https://doi.org/10.1111/socf.12106>
- Levy, Karen, Kyla E. Chasalow, and Sarah Riley. 2021. "Algorithms and Decision - Making in the Public Sector." *Annual Review of Law and Social Science* 17:309–34. <https://doi.org/10.1146/annurev-lawsocsci-041221-023808>
- Lewis-McCoy, R. L'Heureux. 2014. *Inequality in the Promised Land: Race, Resources, and Suburban Schooling*. Stanford University Press.
- Lipsky, Michael. 1980. *Street Level Bureaucracy: Dilemmas of the Individual in Public Services*. Russell Sage Foundation. <https://doi.org/10.2307/1288305>
- Liu, Mengyao, Vernon Choy, Philip Clarke, Adrian Barnett, Tony Blakely, and Lucy Pomeroy. 2020. "The acceptability of using a lottery to allocate research funding: a survey of applicants." *Research Integrity and Peer Review* 5:3. <https://doi.org/10.1186/s41073-019-0089-z>
- MacCoun, Robert J. 2005. "VOICE , CONTROL , AND BELONGING: The Double-Edged Sword of Procedural Fairness." *Annual Review of Law and Social Science* 1:171–201. <https://doi.org/10.1146/annurev.lawsocsci.1.041604.115958>
- Madden, Mary, Michele Gilman, Karen Levy, and Alice Marwick. 2017. "Privacy, Poverty, and Big Data: A Matrix of Vulnerabilities for Poor Americans." *Washington University Law Review* 95:53.
- Martin, Kirsten and Ari Waldman. 2023. "Are Algorithmic Decisions Legitimate? The Effect of Process and Outcomes on Perceptions of Legitimacy of AI Decisions." *Journal of Business Ethics* 183:653–70. <https://doi.org/10.1007/s10551-021-05032-7>
- Miller, Peter M. and Alexis K. Bourgeois. 2013. "Considering the Geographic Dispersion of Homeless and Highly Mobile Students and Families." *Educational Researcher* 42:242–49. <https://doi.org/10.3102/0013189X12474064>
- Miller, Susan M and Lael R Keiser. 2021. "Representative Bureaucracy and Attitudes Toward Automated Decision Making." *Journal of Public Administration Research and Theory* 31:150–65. <https://doi.org/10.1093/jopart/muaa019>
- Mirza, Mansha, Emily Keating, Amy Krischer, and Jessica Pinto. 2022. "Care coordination Experiences of Low-Income Parents of Children and Youth with Special Health Care

- Needs: An Exploratory Study." *Journal of Health Care for the Poor and Underserved* 33:1925–948. <https://doi.org/10.1353/hpu.2022.0146>
- Newman, David T., Nathanael J. Fast, and Derek J. Harmon. 2020. "When eliminating bias isn't fair: Algorithmic reductionism and procedural justice in human resource decisions." *Organizational Behavior and Human Decision Processes* 160:149–67. <https://doi.org/10.1016/j.obhdp.2020.03.008>
- Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting racial bias in an algorithm used to manage the health of populations." *Science* 366:447–53. <https://doi.org/10.1126/science.aax2342>
- Perdomo, Juan C., Tolani Britton, Moritz Hardt, and Rediet Abebe. 2023. "Difficult Lessons on Social Prediction from Wisconsin Public Schools." arXiv:2304.06205 [cs].
- Pierson, Emma. 2018. "Demographics and Discussion Influence Views on Algorithmic Fairness." arXiv:1712.09124 [cs].
- Porter, Theodore M. 1996. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton University Press. <https://doi.org/10.1515/9780691210544>
- Robinson, David G. 2022. *Voices in the Code: A Story about People, Their Values, and the Algorithm They Made*. Russell Sage Foundation. <https://doi.org/10.7758/9781610449144>
- Roda, Allison and Carolyn Sattin-Bajaj. 2023. "Meritocracy and Advantaged Parents' Perceptions of the Fairness of School Choice Policies." *Educational Policy* p. 08959048231174878. <https://doi.org/10.1177/08959048231174878>
- Rona-Tas, Akos. 2020. "Predicting the Future: Art and Algorithms." *Socio-Economic Review* 18:893–911. <https://doi.org/10.1093/ser/mwaa040>
- Rosen, Eva, Philip M. E. Garboden, and Jennifer E. Cossyleon. 2021. "Racial Discrimination in Housing: How Landlords Use Algorithms and Home Visits to Screen Tenants." *American Sociological Review* 86:787–822. <https://doi.org/10.1177/00031224211029618>
- Sattin-Bajaj, Carolyn, Jennifer L. Jennings, Sean P. Corcoran, Elizabeth Christine Baker-Smith, and Chantal Hailey. 2018. "Surviving at the Street Level: How Counselors' Implementation of School Choice Policy Shapes Students' High School Destinations." *Sociology of Education* 91:46–71. <https://doi.org/10.1177/0038040717751443>
- Sauder, Michael. 2020. "A Sociology of Luck." *Sociological Theory* 38:193–216. <https://doi.org/10.1177/0735275120941178>
- Schneider, Anne and Helen Ingram. 1993. "Social Construction of Target Populations: Implications for Politics and Policy." *American Political Science Review* 87:334–47. <https://doi.org/10.2307/2939044>
- Schoon, Eric W. 2022. "Operationalizing Legitimacy." *American Sociological Review* 87:478–503. <https://doi.org/10.1177/00031224221081379>
- Singer, Jeremy. 2023. "School Choice, Socioeconomic Status, and Stratified Enrollment Among Low-Income and Racially Minoritized Students: A Mixed-Methods Study in Detroit." <https://doi.org/10.2139/ssrn.4390504>
- Skinner, Rebecca. 2014. "Charter school programs authorized by the Elementary and Secondary Education Act (ESEA Title VB): A primer." Congressional Research Service, Washington, DC. https://www.everycrsreport.com/files/20140422_R43493_ac25b0bc5009b35d5870abc3935f017b0afad72b.pdf
- Smith, Aaron. 2018. "Public Attitudes Toward Computer Algorithms." <https://policycommons.net/artifacts/617047/public-attitudes-toward-computer-algorithms/1597793/>

- Starke, Christopher, Janine Baleis, Birte Keller, and Frank Marcinkowski. 2022. "Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature." *Big Data & Society* 9:20539517221115189. <https://doi.org/10.1177/20539517221115189>
- Starr, Sonja B. 2014. "Evidence- Based Sentencing and the Scientific Rationalization of Discrimination." *Stanford Law Review* 66:803–72.
- Steensland, Brian. 2006. "Cultural Categories and the American Welfare State: The Case of Guaranteed Income Policy." *American Journal of Sociology* 111:1273–326. <https://doi.org/10.1086/499508>
- ten Cate, Olle. 2021. "Rationales for a Lottery Among the Qualified to Select Medical Trainees: Decades of Dutch Experience." *Journal of Graduate Medical Education* 13:612–15. <https://doi.org/10.4300/JGME-D-21-00789.1>
- Tilcsik, András. 2021. "Statistical Discrimination and the Rationalization of Stereotypes." *American Sociological Review* 86:93–122. <https://doi.org/10.1177/0003122420969399>
- Trinidad, Jose Eos. 2024. "Positionality in the Politics of Education: Mapping the Strategies and Constraints of Education Reformers." *Educational Policy* p. 08959048231220038. <https://doi.org/10.1177/08959048231220038>
- Törnberg, Petter. 2024. "Large Language Models Outperform Expert Coders and Supervised Classifiers at Annotating Political Social Media Messages." *Social Science Computer Review* p. 08944393241286471. <https://doi.org/10.1177/08944393241286471>
- Tyler, Tom R. 1996. "The relationship of the outcome and procedural fairness: How does knowing the outcome influence judgments about the procedure?" *Social Justice Research* 9:311–25. <https://doi.org/10.1007/BF02196988>
- U.S. Department of Education. 2016. "Issue Brief: Early Warning Systems." Office of Planning, Evaluation and Policy Development, US Department of Education. <https://eric.ed.gov/?id=ED571990>
- Vong, Gerard. 2020. "Weighing Up Weighted Lotteries: Scarcity , Overlap Cases , and Fair Inequalities of Chance." *Ethics* 130:320–48. <https://doi.org/10.1086/707212>
- Waldman, Ari and Kirsten Martin. 2022. "Governing algorithmic decisions: The role of decision importance and governance on perceived legitimacy of algorithmic decisions." *Big Data & Society* 9:20539517221100449. <https://doi.org/10.1177/20539517221100449>
- Wang, Angelina, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2024. "Against Predictive Optimization: On the Legitimacy of Decision -making Algorithms That Optimize Predictive Accuracy." *ACM J. Responsib. Comput.* 1:9:1–9:45. <https://doi.org/10.1145/3636509>
- Watkins-Hayes, Celeste and Elyse Kovalsky. 2016. "The Discourse of Deservingness: Morality and the Dilemmas of Poverty Relief in Debate and Practice." In *The Oxford Handbook of the Social Science of Poverty* , edited by David Brady and Linda M. Burton. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199914050.013.10>
- Weber, Max. 2009(1921). *From Max Weber: Essays in Sociology*. Routledge.
- Welsh, Richard O. 2017. "School Hopscotch: A Comprehensive Review of K -12 Student Mobility in the United States." *Review of Educational Research* 87:475–511. <https://doi.org/10.3102/0034654316672068>
- Zhang, Simone and Rebecca A. Johnson. 2023. "Hierarchies in the Decentralized Welfare State: Prioritization in the Housing Choice Voucher Program." *American Sociological Review* 88:114–53. <https://doi.org/10.1177/00031224221147899>

Acknowledgments: Thanks to the following students for excellent research assistance—Collin Crane, Liz Moison, Morgan Welch, and Rosy Zhong—and to Leah Jones, Katherine Christie, and Tyler Simko for related collaborations/discussions. We are also grateful for feedback from the following audiences: Sociology of Education Association annual meeting; Georgetown McCourt School of Public Policy seminar series; Georgetown Sociology colloquium; the Notre Dame Center for Research on Educational Opportunity; APPAM and Sean Reardon as a discussant; Lydia Liu's AI, Society, and Education Seminar at Princeton University; and James Druckman and anonymous reviewers via the TESS process. This research received funding from the Dartmouth Neukom Institute for Computational Science, the NSF TESS Young Investigators Special Competition (NSF Grant 0818839; Jeremy Freese and James Druckman, Principal Investigators), and the Spencer/NAEd Postdoctoral Fellowship.

Rebecca A. Johnson: Equal first authorship. McCourt School of Public Policy (affiliate: Department of Sociology), Georgetown University. E-mail: rj545@georgetown.edu.

Simone Zhang: Equal first authorship. Department of Sociology, New York University. E-mail: simone.zhang@nyu.edu.