

MITIGATING INCLUDED- AND OMITTED-VARIABLE BIAS IN ESTIMATES OF DISPARATE IMPACT

BY JONGBIN JUNG, SAM CORBETT-DAVIES, JOHANN D. GAEBLER^{1,a} ,
RAVI SHROFF^{2,b}, AND SHARAD GOEL^{3,c} 

¹*Department of Statistics, Harvard University, jgaebler@fas.harvard.edu*

²*Applied Statistics, Social Science, and Humanities, New York University, ravi.shroff@nyu.edu*

³*Harvard Kennedy School, sgoel@hks.harvard.edu*

In studies of discrimination, one popular strategy is to estimate disparities after adjusting for observed covariates, typically with a regression model. This approach, however, suffers from two key statistical challenges. First, omitted-variable bias can skew results if the model does not adjust for all relevant factors; second, and conversely, included-variable bias—a less well known phenomenon—can skew results if the set of covariates includes irrelevant factors. Here we introduce a new, three-step statistical method, which we call risk-adjusted regression, to address both concerns in settings where decision makers have clearly measurable objectives. In the first step, we use all available covariates to estimate the value, or inversely, the *risk*, of taking a certain action, such as approving a loan application or hiring a job candidate. Second, we measure disparities in decisions after adjusting for these risk estimates alone, mitigating the problem of included-variable bias. Finally, in the third step, we assess the sensitivity of results to potential mismeasurement of risk, addressing concerns about omitted-variable bias. To do so, we develop a novel, non-parametric sensitivity analysis that yields tight bounds on the true disparity in terms of the average gap between true and estimated risk—a single interpretable parameter that facilitates credible estimates. We demonstrate this approach on a detailed dataset of 2.2 million police stops of pedestrians in New York City, and show that traditional statistical tests of discrimination can substantially underestimate the magnitude of disparities.

1. Introduction. Studies of discrimination generally start by assessing whether certain groups, particularly those defined by race and gender, receive favorable decisions more often than others. For example, one might examine whether loan applications from white candidates are granted more often than those from racial minorities, or whether male employees are promoted more often than women. Although observed disparities may be the result of bias, it is also possible that they stem from statistical differences between groups. In particular, if some groups contain disproportionately many qualified members, then one would also expect those groups to receive disproportionately many favorable decisions, even in the absence of discrimination.

To tease apart these two possibilities—group differences versus discrimination—the most popular statistical approach is ordinary linear or logistic regression. In the banking context, for instance, one could examine race-contingent lending rates after adjusting for relevant factors, such as income and credit history. Disparities that persist after accounting for such factors are often interpreted as evidence of discrimination. This basic statistical strategy has been used in numerous studies to test for bias across domains, including education ([Espenshade, Chung and Walling, 2004](#); [Grossman et al., 2023](#)), employment ([Polachek, 2008](#)), criminal

Keywords and phrases: discrimination, omitted-variable bias, included-variable bias, risk-adjusted regression, sensitivity analysis.

justice (Gaebler et al., 2022; Gelman, Fagan and Kiss, 2007; MacDonald and Raphael, 2019), and medicine (Balsa, McGuire and Meredith, 2005).

Despite the ubiquity of such regression-based tests for discrimination, the approach suffers from two serious statistical limitations. First, the well-known problem of omitted-variable bias arises when decisions are based in part on relevant factors that correlate with group membership, but which are omitted from the regression (Angrist and Pischke, 2008). For example, if lending officers consider an applicant’s payment history, and if payment history correlates with race but is not recorded in the data (and thus cannot be included in the regression), the results of the regression can suggest discrimination where there is none, or *vice versa*. Unfortunately, omitted-variable bias is the rule rather than the exception. It is generally prohibitive to measure every variable relevant to a decision, and it is likely that most unmeasured variables are at least weakly correlated with demographic attributes, skewing results.

A second problem with regression-based tests is what Ayres (2005, 2010) calls *included-variable bias*, an issue as important as omitted-variable bias in studies of discrimination but one that receives far less attention. To take an extreme example, it is problematic to include control variables in a regression that are obvious proxies for legally protected attributes—such as vocal register as a proxy for gender—when examining the extent to which observed disparities stem from group differences in qualification. Including such proxies will typically lead one to underestimate the true magnitude of discrimination in decisions. But what counts as a “proxy” is not always clear. For example, given existing patterns of residential segregation, one might argue that ZIP codes are a proxy for race, and thus should be excluded when testing for racial bias. But one could also argue that ZIP code provides legitimate information relevant to a decision, and so excluding it would lead to omitted-variable bias. Ayres (2010) proposes a middle ground, suggesting that potential proxies should be included, but their coefficients capped to a “justifiable” level; in practice, however, it is difficult to determine and defend specific constraints on regression coefficients.¹

Here we present a statistically principled and logistically straightforward method for measuring discrimination that addresses both omitted- and included-variable bias. Our method, which we call risk-adjusted regression, proceeds in three steps. In the first step, we use all available information, including potential proxies of protected traits, to estimate the value—or, equivalently, the risk—of taking a particular action. For example, in the lending context, we might estimate an applicant’s risk of default if granted a loan, conditional on all available covariates. In the second step, we assess disparities by regressing decisions (e.g., loan offers) on individual-level risk estimates and protected traits alone, allowing us to measure the extent to which similarly qualified individuals are treated differently. This strategy can be seen as formalizing the coefficient-capping procedure of Ayres (2010)—with covariates used only to the extent that they are statistically justified by risk—and thus circumvents the problem of

¹Such problems have prompted a search for alternatives to regression-based approaches. Most prominently, Becker (1993) proposed the outcome test, which is based not on the rate at which decisions are made, but on the success rate of those decisions. In the context of banking, Becker argued that even if, hypothetically, racial minorities were less creditworthy than white applicants, minorities who were granted loans should still be found to repay their loans at the same rate as white applicants who were granted loans. If loans to minorities had a higher repayment rate than loans to white borrowers, it suggests that lenders applied a double standard (intentionally or not), granting loans only to exceptionally qualified minorities—potentially violating disparate impact laws. The outcome-based approach has been applied almost as broadly as simple regression to study discrimination in, for example, policing (Ayres, 2002; Goel, Rao and Shroff, 2016a; Goel et al., 2017; Knowles, Persico and Todd, 2001), lending (Berkovec et al., 2018), and scientific publication (Smart and Waldfogel, 1996). Outcome tests, however, have their own significant statistical shortcomings, most notably the problem of infra-marginality (Ayres, 2002; Pierson, Corbett-Davies and Goel, 2018; Simoiu, Corbett-Davies and Goel, 2017), which can lead the test to incorrectly suggest an absence of discrimination (Pierson et al., 2020).

included-variable bias. Finally, we assess the sensitivity of results to potential mismeasurement of risk. In particular, we derive tight analytic bounds on risk-adjusted disparities as a function of the extent to which risk estimates differ from true risk.

To demonstrate this approach, we examine 2.2 million stops of pedestrians conducted by the New York City Police Department between 2008 and 2011. After adjusting for a stopped individual’s statistical risk of carrying a weapon—based in part on detailed behavioral indicators recorded by officers—we find that stopped Black and Hispanic pedestrians are searched for weapons substantially more often than stopped white individuals. We find that these risk-adjusted disparities are considerably larger than disparities suggested by a standard regression that adjusts for all available covariates, underscoring the importance of accounting for included-variable bias. Finally, we show that our results are robust to potentially large errors in risk estimates.

2. Theories of Discrimination. There are two main legal doctrines of discrimination in the United States: disparate treatment and disparate impact. The first, disparate treatment, derives force from the Equal Protection Clause of the U.S. Constitution’s Fourteenth Amendment, and it prohibits government agents from acting with “discriminatory purpose” ([Washington v. Davis, 1976](#)). Although equal protection law bars policies undertaken with animus, it allows for the limited use of protected attributes to further a compelling government interest. For example, until recently, certain affirmative action programs for college admissions were legally permissible to further the government’s interest in promoting diversity. In 2023, the U.S. Supreme Court overturned the legality of such affirmative action programs, ruling that it was unlawful to explicitly consider race in college admissions decisions ([SFFA v. Harvard, 2023](#)).

The most widespread statistical test of such intentional discrimination is ordinary linear or logistic regression, in which one estimates the likelihood of favorable (or unfavorable) decisions across groups defined by race, gender, or other legally protected traits. In this approach, the investigator adjusts for all potentially relevant risk factors, excluding only clear proxies for the protected attributes. Barring omitted-variable bias, non-zero coefficients on the protected traits suggest those factors influenced the decision maker’s actions; in the absence of a compelling justification, such evidence is suggestive of a discriminatory purpose. It is difficult—and perhaps impossible—to rigorously define the *influence*, or causal effect, of largely immutable traits like race on decisions ([Greiner and Rubin, 2011](#); [VanderWeele and Robinson, 2014](#)), but a regression of this type is nevertheless considered a reasonable first step to identify discriminatory motive, both by criminologists and by legal scholars ([Fagan, 2010](#); [Gaebler et al., 2022](#)). For an equal protection claim to succeed in court, however, one typically needs additional documentary evidence (e.g., acknowledgement of an illegitimate motive) to bolster the statistical evidence.

In contrast to disparate treatment, the disparate impact doctrine is concerned with the effects of a policy, not a decision maker’s intentions, and it is the primary form of discrimination we study in this paper. Under the disparate impact standard, a practice may be deemed discriminatory if it has an unjustified adverse effect on protected groups, even in the absence of explicit categorization or animus. The doctrine stems from statutory rules, rather than constitutional law, and applies only in certain contexts, such as employment (via Title VII of the 1964 Civil Rights Act) and housing (via the Fair Housing Act of 1968). Apart from federal statutes, some states have passed more expansive disparate impact laws, including Illinois and California.

The disparate impact doctrine was formalized in the landmark U.S. Supreme Court case *Griggs v. Duke Power Co.* (1971). In 1955, the Duke Power Company instituted a policy that mandated employees have a high school diploma to be considered for promotion, which

had the effect of drastically limiting the eligibility of Black employees. The Court found that this requirement had little relation to job performance, and thus deemed it to have an unjustified disparate impact. Importantly, the employer’s motivation for instituting the policy was irrelevant to the Court’s decision; even if enacted without discriminatory purpose, the policy was deemed discriminatory in its effects and hence illegal.

As discussed above, the standard statistical test for disparate treatment is a “kitchen sink” regression, where one examines the residual explanatory power of protected group status after including all other available covariates as controls. That approach, however, is ill-suited to assess whether practices are rationally justified, which is the relevant standard in disparate impact claims. [Ayres \(2005\)](#) makes the point persuasively in the context of the original *Griggs* decision:

“One could imagine running a regression to test whether an employer was less likely to hire African American applicants than white applicants. It would be possible to control in this regression for whether the applicant had received a high-school diploma. Under the facts of *Griggs*, such a control would likely have reduced the racial disparity in the hiring rates. But including in the regression a variable controlling for applicants’ education would be inappropriate. The central point of *Griggs* was to determine whether the employer’s diploma requirement had a disparate racial impact. The possibility that including a diploma variable would reduce the estimated race effect in the regression would in no way be inconsistent with a theory that the employer’s diploma requirement disparately excluded African Americans from employment.”

In short, by including educational status in the regression, one would mask the policy’s unjustified disparate impact.

To assess claims of disparate impact—in *Griggs* and beyond—one would ideally compare decision rates for similarly *qualified* groups of applicants (e.g., similarly qualified white and Black candidates). Unfortunately, if one does not, or cannot, adjust for sufficiently many covariates, omitted-variable bias may skew results; conversely, if one does adjust for a rich set of covariates, included-variable bias may corrupt conclusions.²

3. A Statistical Approach to Assessing Disparate Impact. We now formally describe our approach to measuring disparate impact—a procedure we call risk-adjusted regression. We imagine that every individual in some population of n people is subject to a binary decision A , with X denoting the information on which that decision was based. We further assume that each individual is a member of one of m protected classes $C \in \{1, \dots, m\}$ (e.g., race, gender, or combinations thereof). While not statistically necessary, it is convenient to take $C = 1$ to represent membership in some reference group of interest against whose treatment members of other groups will be compared. Finally, we assume each individual has a binary property of interest W , which decision makers typically seek to infer via their available information X . For each individual, the decision maker observes X , which may or may not include the individual’s protected class C . Conversely, the analyst observes the decision A and the protected class C , but typically only observes an approximate version of X , which we call \tilde{X} . Importantly, we also assume the analyst observes W for the subset of individuals for whom $A = 1$.

As a concrete example, consider estimating risk-adjusted racial disparities in police searches of pedestrians for weapons—the application we discuss in more detail below. In

²A related body of work conceptualizes discrimination as race-specific differences in decision error rates ([Arnold, Dobbie and Hull, 2020, 2021](#); [Bohren, Hull and Imas, 2022](#); [Coston et al., 2020](#); [Hardt, Price and Srebro, 2016](#)). For example, one might consider differences in loan rejection rates between white and Black applicants who in reality would repay their loans (i.e., a counterfactual false negative rate). Recent work, however, has observed that by conditioning on *ex post* (potential) outcomes—as opposed to *ex ante* risk, as we do—such approaches suffer from infra-marginality and associated statistical issues, rendering them problematic measures of common legal and policy understandings of discrimination ([Ayres, 2005](#); [Corbett-Davies et al., 2023](#); [Mayson, 2018](#); [Nilforoshan et al., 2022](#)).

this case, A_i , X_i , C_i , and W_i , respectively, indicate whether the i -th stopped pedestrian was searched, the information available to the officer when deciding whether to conduct a search, the stopped individual’s race, and whether the individual was in possession of a weapon. Further, \tilde{X}_i denotes the information available to analysts from administrative records of the i -th stop. Our statistical assumption on W above corresponds to the analyst knowing whether searched individuals have a weapon, which, in practice, is generally the case. Although we do not consider them further here, this general framing applies to a wide variety of settings, such as hiring (where W represents a prospective employee’s productivity and A whether or not they were hired) and lending (where W represents the amount a prospective borrower will repay and A represents the bank’s lending decision).

Given this setup, we define *ex ante risk* to be:

$$(1) \quad R = \Pr(W = 1 \mid X).$$

In our policing example, R_i is the probability that the i -th stopped individual, with covariates X_i , is carrying a weapon.³ Our goal is to quantify whether decisions systematically differ for individuals at the same level of risk. Many estimands summarizing such potential differences are possible, but, for simplicity, we consider the following non-parametric quantity:

$$(2) \quad \mathbb{E}[\Pr(A = 1 \mid C = j, R) - \Pr(A = 1 \mid C = 1, R)].$$

Eq. (2) defines risk-adjusted disparities to be the difference in the probability of taking an action for group $C = j$ relative to the reference group $C = 1$, after accounting for potential differences in risk across groups.

To facilitate computation—and, in particular, the sensitivity analysis presented below—we approximate the estimand in (2) by a linear probability model. Specifically, we estimate $\Pr(A = 1 \mid C, R)$ as

$$(3) \quad \sum_{j=1}^m \hat{\beta}_j \cdot \mathbf{1}(C = j) + \hat{\beta}_R \cdot R,$$

where $\mathbf{1}(C = j)$ indicates membership in group j and $\hat{\beta}_j$ is its corresponding fitted coefficient. The difference between fitted coefficients, $\hat{\beta}_j - \hat{\beta}_1$, yields an estimate of (2)—see the Appendix for further details on the statistical quality of this estimator.

Under this model, if $\hat{\beta}_j$ is greater than $\hat{\beta}_1$, this indicates that members of group j are more likely to receive action $A = 1$ than members of the base group with similar estimated risk. In our policing example, this means that members of group j are searched more often than members of the reference group who were equally likely to be carrying a weapon, and we would say that such elevated search rates are *unjustified* by risk. We note, however, that $\hat{\beta}_j > \hat{\beta}_1$ does not imply *intentional* discrimination—as in *Griggs*, unjustified disparate impact is possible even under a facially neutral policy undertaken without animus.

4. Sensitivity analysis. Our definition of disparate impact depends on the true risks, R_i . But, in practice, analysts observe only partial information on X and W , and so at best can construct only imperfect estimates of risks, \hat{R}_i . For instance, in our running police example,

³To allow for the possibility of statistical discrimination (Arrow, 1973; Chaudhuri and Sethi, 2008; Phelps, 1972) in the decision maker’s course of action, risk may be conditioned on group membership as well as other observables—formally, by including C among the covariates X . However, whether the inclusion of race or other protected attributes in risk models is legally appropriate is uncertain. Importantly, our theoretical and empirical results remain essentially unchanged whether or not one assumes X encodes C ; see Figures 13 through 19 in the Appendix.

officers might base search decisions in part on subtle behavioral cues that are not documented in the data; and analysts typically would not know whether individuals who were *not* searched were in fact carrying a weapon. In Section 5, we discuss various approaches to estimating risk in light of these challenges. Whatever approach one adopts, the accuracy of one’s conclusions depends critically on the accuracy of one’s estimates of risk. To address this issue, we develop a novel sensitivity analysis, drawing inspiration from methods for sensitivity analysis popular in the causal inference literature (Jung et al., 2020; Rosenbaum and Rubin, 1983)—though we emphasize that our framework does not itself involve estimating causal effects.

To start, we assume that there is a known constant $\epsilon \geq 0$ such that

$$(4) \quad \frac{1}{n} \sum_{i=1}^n |R_i - \hat{R}_i| \leq \epsilon;$$

that is, that the true risks and the estimated risks differ on average by at most ϵ . Given this assumption, the goal of our sensitivity analysis is to understand how different our estimate of disparate impact, $\hat{\beta}_j - \hat{\beta}_1$, could be if we had access to the true risk R_i instead of the estimated risk \hat{R}_i . In practice, an analyst would examine the robustness of conclusions to different choices of ϵ .

We could attempt to bound our estimate of disparate impact by searching over all possible choices of R_i satisfying the constraint in Eq. (4). But such an approach is overly conservative, as the observed data themselves rule out possible values of R_i . In our policing example, for instance, we expect the average true risk of searched individuals to approximately equal the proportion of searched individuals carrying a weapon. Since, by assumption, the observed data tell us about this latter quantity, they constrain the possible risk distributions we must consider.

To formalize our approach, a key quantity to consider is the average (true) risk on each stratum,

$$(5) \quad \frac{1}{|\mathcal{S}_{j,a}|} \cdot \sum_{i \in \mathcal{S}_{j,a}} R_i,$$

where the stratum

$$\mathcal{S}_{j,a} = \{i : C_i = j, A_i = a\}$$

denotes all those individuals in group j for whom the decision maker took action a . By the law of large numbers, the average risk on stratum $\mathcal{S}_{j,a}$ is approximated by $\mathbb{E}[R | C = j, A = a]$. Further,

$$\begin{aligned} \mathbb{E}[R | C = j, A = a] &= \mathbb{E}[\mathbb{E}[W | X] | C = j, A = a] \\ &= \mathbb{E}[W | C = j, A = a], \end{aligned}$$

where the first equality follows by definition, and the second by the law of iterated expectations as well as the fact that A and C are functions of X .⁴ When $A = 1$, by our assumption in Section 3, the above quantity is identified by data available to the analyst. In particular, in our policing example, a consistent estimator of $\mathbb{E}[W | C = j, A = 1]$ is the proportion of searched individuals possessing a weapon, among those individuals belonging to group j . We denote by ρ_j the analyst’s estimate of $\mathbb{E}[W | C = j, A = 1]$ based on the observed data.⁵

⁴In some contexts, rather than thinking of A as a deterministic function of X and C , it is useful to assume that decisions depend also on some exogenous source of variation (e.g., an officer might flip an appropriately weighted coin to decide whether to conduct a search). Importantly, this randomness cannot reveal anything about W , so $W \perp\!\!\!\perp A | X$, and the equality still holds.

⁵One might alternatively estimate $\mathbb{E}[W | C = j, A = 1]$ by the average *estimated* risks on the stratum, namely $\frac{1}{|\mathcal{S}_{j,1}|} \cdot \sum_{i \in \mathcal{S}_{j,1}} \hat{R}_i$. The alternative, which we adopt in our empirical analysis in Section 5, has the interpretive advantage that our sensitivity analysis exactly recovers the correct answer when $\epsilon = 0$.

Finally, putting the pieces together, for each particular j^* we seek the largest and smallest values of $\hat{\beta}_{j^*} - \hat{\beta}_1$, as estimated with R_i , subject to two key constraints: (1) the average absolute deviation between R_i and \hat{R}_i is less than ϵ ; and (2) R_i is consistent with the observed data on each stratum $\mathcal{S}_{j,1}$. This is succinctly expressed as an optimization problem (henceforth **“the base problem”**):

$$(6) \quad \begin{aligned} & \underset{\mathbf{R} \in \mathbb{R}^n}{\text{Optimize}} && \hat{\beta}_{j^*} - \hat{\beta}_1 \\ & \text{s.t.} && \frac{1}{n} \sum_{i=1}^n |R_i - \hat{R}_i| \leq \epsilon \\ & && \frac{1}{|\mathcal{S}_{j,1}|} \cdot \sum_{i \in \mathcal{S}_{j,1}} R_i = \rho_j, \quad (j = 1, \dots, m) \\ & && R_i \leq u_i, \quad (i = 1, \dots, n) \\ & && R_i \geq \ell_i, \quad (i = 1, \dots, n), \end{aligned}$$

where, by “Optimize,” we simply mean both maximize and minimize, and \hat{R}_i and ρ_j are fixed by the data as discussed above.⁶ The additional upper and lower bounds on R_i arise from the fact that the risks must be probabilities, i.e., $0 \leq R_i \leq 1$, though our approach accommodates tighter, individual-specific bounds $\ell_i \leq R_i \leq u_i$ set by the analyst.

There are two main obstacles to solving the base problem above. First, the dimension of the search space is $n - m$, where n is the number of observations and m is the number of groups. Second, the regression coefficients $\hat{\beta}_j$ are highly non-convex functions of the true risk vector $\mathbf{R} = (R_1, \dots, R_n)$. Addressing these twin challenges involves a detailed analysis of the underlying geometry of the problem. Here we present an outline of our approach that illustrates the key ideas, with the full exposition in the Appendix.

It is useful to think of the base problem as an optimization over subproblems (henceforth **“the parameterized problem”**) where the average (true) risk is assumed known and equal to some τ_j for *unobserved* strata $\mathcal{S}_{j,0}$ as well. Using the closed form of OLS regression, we can show that solving the parameterized problem reduces to—although is not equivalent to—solving the following optimization problem (henceforth **“the simplified problem”**):

$$(7) \quad \begin{aligned} & \underset{\mathbf{R} \in \mathbb{R}^n}{\text{Optimize}} && \frac{1}{n} \sum_{i=1}^n R_i^2 \\ & \text{s.t.} && \frac{1}{n} \sum_{i=1}^n |R_i - \hat{R}_i| \leq \epsilon, \\ & && \frac{1}{|\mathcal{S}_{j,1}|} \cdot \sum_{i \in \mathcal{S}_{j,1}} R_i = \rho_j, \quad (j = 1, \dots, m) \\ & && \frac{1}{|\mathcal{S}_{j,0}|} \cdot \sum_{i \in \mathcal{S}_{j,0}} R_i = \tau_j, \quad (j = 1, \dots, m) \\ & && R_i \leq u_i, \quad (i = 1, \dots, n) \\ & && R_i \geq \ell_i. \quad (i = 1, \dots, n) \end{aligned}$$

An efficient method for solving the simplified problem then yields an efficient method for solving the base problem simply by searching over this much smaller m -dimensional parameter space of (τ_1, \dots, τ_m) . In practice, m typically equals two or three, and so this final optimization step can be solved either using a grid search, as we do in Section 5, or using other optimization techniques for low-dimensional non-convex problems.

⁶It may seem inappropriate to fix the average risks on the observed stratum $\mathcal{S}_{j,1}$ to be exactly ρ_j , since ρ_j is estimated with error. One could instead allow ρ_j to vary within some range (e.g., a 95% confidence interval), solving the parameterized problem for various ρ_j as well as τ_j , as detailed below. However, doing so doubles the dimension of the search space, which is computationally costly, and does not fully take advantage of available information about the distribution of the estimation error. A more principled and computationally feasible approach to dealing with estimation error is to bootstrap confidence intervals for the bounds following [Zhao, Small and Bhattacharya \(2019\)](#), as we do in Section 5.2 below, where ρ_j is re-estimated in each bootstrapped resample, and the base problem is solved with corresponding equality constraints.




Fig 1: *Minimization normal form.* Black dots represent estimated risks \hat{R}_i , the error bars represent the allowable range $[\ell_i, u_i]$, and red dots indicate the corresponding value of the normal form vector R_i . Here Δ denotes the gap between the upper and lower thresholds.

The objective of the simplified problem, $\frac{1}{n} \sum_{i=1}^n R_i^2$, is convex, as are the other constraints on \mathbf{R} , and so, in principle, minimization in the simplified problem can be achieved efficiently and accurately using interior-point methods (Boyd and Vandenberghe, 2004). However, by carefully examining the KKT conditions (Karush, 1939; Kuhn and Tucker, 1951), it is possible to see that the minimizing risk vector must adhere to a special form, which we term “minimization normal form.” Specifically, there exist thresholds $t_{j,a}^{lwr}$ and $t_{j,a}^{upr}$ for each stratum such that the minimizing R_i is obtained by “pulling up” low estimated risks to $t_{j,a}^{lwr}$, “pulling down” high estimated risks to $t_{j,a}^{upr}$, and leaving unchanged intermediate risks. An example of a risk vector in minimization normal form is shown in Figure 1.

This observation reduces the dimension of the search space from $n - 2m$ to $2m$; that is, by the above, one need only search over the $2m$ thresholds. One can, however, do better. A careful examination of the KKT conditions reveals that while the thresholds $t_{j,a}^{lwr}$ and $t_{j,a}^{upr}$ may vary by stratum, the gap between them $\Delta = t_{j,a}^{upr} - t_{j,a}^{lwr}$ must be the same across strata. Moreover, once Δ is fixed, the thresholds $t_{j,a}^{lwr}$ and $t_{j,a}^{upr}$ are determined by the strata-specific average risks ρ_j and τ_j . As a result, finding the minimizing risk vector for the simplified problem reduces to optimizing over a single parameter, Δ , which, in our setting, can be done very efficiently. In particular, we give an $O(n \cdot \log(m))$ algorithm that finds the exact solution.

Maximization requires more care, since the maximization problem is non-convex, and the general problem of maximizing a quadratic objective over a convex set is NP-hard (Sahni, 1974). Using a similar, but more delicate, version of the techniques used to minimize the simplified problem, we give an $O(m^3 \cdot (\epsilon/\delta)^2 + n)$ -time algorithm for approximating the maximum of the simplified algorithm to within δ . The key steps described above for optimizing the base problem are given in Algorithm 1.

As we noted previously, our approach to sensitivity analysis for risk-adjusted regression is related to sensitivity analysis methods developed in the causal inference literature, particularly to recent work on L_2 sensitivity bounds (Huang and Pimentel, 2022; Zhang and Zhao, 2022). Our approach, however, is distinct in important ways that make it uniquely advantageous in our setting. First, methods for sensitivity analysis in the causal inference literature that can be adapted to our setting often involve a large number of parameters—making them difficult to credibly calibrate—or require parametric assumptions on the form of the confounding itself (e.g. Cinelli and Hazlett, 2020; Jung et al., 2020; Rosenbaum, 2002; Rosenbaum and Rubin, 1983). In contrast, the single parameter in our sensitivity analysis—

Algorithm 1 (Sensitivity analysis)

Input: The estimated risk vector $\hat{\mathbf{R}}$, the lower and upper bounds ℓ and \mathbf{u} , the average absolute difference ϵ , and the approximation error δ .

Output: The minimum and maximum values of $\hat{\beta}_j - \hat{\beta}_1$ for all groups $j = 1, \dots, m$.

- 1: Set $\rho_j \leftarrow \frac{1}{|S_{j,1}|} \cdot \sum_{i \in S_{j,1}} \hat{R}_i$ for $j = 1, \dots, m$
- 2: **repeat**
- 3: Calculate \mathbf{R}_{\min}^* minimizing the simplified problem with ρ_1, \dots, ρ_m and τ_1, \dots, τ_m
- 4: Calculate \mathbf{R}_{\max}^* maximizing the simplified problem to within δ with ρ_1, \dots, ρ_m and τ_1, \dots, τ_m
- 5: Calculate $\hat{\beta}_j - \hat{\beta}_1$ using both \mathbf{R}_{\min}^* and \mathbf{R}_{\max}^*
- 6: **until** The grid of (τ_1, \dots, τ_m) is exhausted
- 7: **return** The largest and smallest differences $\beta_j - \beta_1$ observed for all $j = 2, \dots, m$

ϵ , the average absolute deviation between the true and estimated risks—is straightforward to reason about and explain, particularly to policymakers or other non-technical stakeholders.

Further, most past sensitivity methods that do involve a single parameter would, if suitably adapted to our setting, be parameterized in terms of an L_∞ -bound Γ on the log odds ratio between true and estimated risks (e.g., [Dorn and Guo, 2022](#); [Zhao, Small and Bhattacharya, 2019](#)). However, in realistic models of confounding, the change in the odds ratio typically cannot be bounded in this way for all units: some—potentially very small—set of units will have an estimated risk of approximately 0 but a true risk of approximately 1 or *vice versa*. As a result, one may need to set Γ to be very large to strictly satisfy the assumptions of the approach, yielding sensitivity bounds that are, in practice, quite conservative. More recent methods involving L_2 bounds avoid this problem ([Huang and Pimentel, 2022](#); [Zhang and Zhao, 2022](#)), but would not yield tight bounds on the estimand of interest in our setting.

5. An Application to Policing. We apply our approach above to investigate the “stop-and-frisk” practices of the New York City Police Department (NYPD).⁷ Police officers in the United States may stop and question pedestrians if they have “reasonable and articulable” suspicion of criminal activity; officers may additionally conduct a “frisk” (i.e., a brief pat-down of one’s outer garments) if they believe the stopped individual is carrying a weapon. Although a policy of stopping and frisking individuals is not inherently illegal, in *Floyd v. City of New York* (2013), a federal district court ruled that the NYPD carried out such stops with racial animus, violating the Equal Protection Clause of the Fourteenth Amendment.

The court in *Floyd* was interested in assessing claims of disparate treatment; here we re-analyze the data with a focus on disparate impact. We specifically consider frisk decisions, as they have a clear goal of ensuring officer safety by recovering weapons, and a well-measured outcome—whether a weapon was in fact found. We study 2.2 million pedestrian stops that occurred between 2008 and 2011. For each stop, we have detailed information on the date, time, and location of the stop; the demographics of the stopped individual (e.g., age, gender, and race); the suspected crime; the reasons prompting the stop (e.g., “furtive movements” or “suspicious bulge”); and additional circumstances surrounding the stop (e.g., evasive responses to questioning, witness reports, or evidence of criminal activity in the vicinity).⁸

To start, we note that 1.7% of frisks turn up a weapon. White pedestrians are frisked in 44% of police stops, whereas Black and Hispanic pedestrians are frisked in 57% and 58% of stops, respectively, a 13–14 percentage point gap. These raw disparities are computed without adjusting for any potentially explanatory variables, and so represent an extreme case of

⁷Reproduction materials for this analysis are available at <https://github.com/jgaeb/rar-paper>.

⁸This information is recorded in a standardized way on UF-250 forms that officers are required to complete after each stop. A copy of the form can be [found online](#).

omitted-variable bias. At the other extreme is the “kitchen sink regression”, which adjusts for all observed pre-frisk covariates in a standard linear probability model. In this case, stopped Black and Hispanic pedestrians are 3–4 percentage points more likely to be frisked, relative to white pedestrians with similar recorded characteristics. These kitchen-sink disparities are suggestive of disparate treatment (and similar evidence was indeed presented to the court in *Floyd* to support such an allegation), but they may understate the extent to which the policy imposes an unjustified disparate impact on racial minorities, due to included-variable bias.

5.1. Estimating risk-adjusted disparities. The key ingredient in applying risk-adjusted regression is estimating the risk $R = \Pr(W = 1 \mid X)$, as in Eq. (1), where W indicates whether a stopped individual has a weapon and X is the information available to the officer when making their frisk decision. There are, however, two challenges to estimating this quantity. First, the information available to officers typically differs from that recorded in administrative datasets available to analysts. Second, W remains unobserved for individuals who were not frisked. These challenges can be mitigated, but not eliminated. Our approach is thus to estimate these risks as best we can, and then gauge the robustness of our results to estimation error. For ease of exposition, we adopt a simple estimation strategy, regressing weapon possession (W) on the covariates \tilde{X} in the recorded data, fit on the subset of individuals who were, in reality, frisked.⁹ These estimates suffer from some degree of omitted-variable bias and selection effects, an issue we return to in our sensitivity analysis below.

We start by dividing the original stops into two subsets: (1) the approximately 1 million stops that occurred in 2008 and 2009, which we use to train a risk model; and (2) the remaining 1.2 million stops that occurred in 2010 and 2011, to which we apply the fitted risk model to estimate disparities in frisk decisions. To estimate risk of weapon possession, we use gradient boosted decision trees, a non-linear model popular in the machine learning community for its predictive performance, restricting to stops in the first subset of data in which a frisk was conducted. Predictive performance and model checks presented in Figure 12 in the Appendix indicate that the model yields predictions with reasonable performance and that predictions are well-calibrated across groups. We use this fitted model to produce an estimate of *ex ante* risk \hat{R}_i for every pedestrian stopped in the second subset of the data, including those who were not frisked. These risk estimates are shown in Figure 10 and 11

Figure 2 shows frisk rates as a function of estimated risk, disaggregated by race. At every level of risk, stopped Black and Hispanic pedestrians are frisked at a much higher rate than stopped white individuals, a gap that is suggestive of disparate impact in frisk decisions. To add quantitative detail to these qualitative results, we now compute risk-adjusted disparities, fitting the linear probability model in Eq. (3) on the second half of the NYPD data, computing estimates for the Black-white disparity $\hat{\beta}_{\text{Black}} - \hat{\beta}_{\text{White}}$ and Hispanic-white disparity $\hat{\beta}_{\text{Hispanic}} - \hat{\beta}_{\text{White}}$. Figure 3 shows the results, together with the raw disparities and those estimated from a kitchen-sink model. We find that stopped Black and Hispanic pedestrians were about 15 percentage points more likely to be frisked than white pedestrians who were equally likely to be carrying a weapon. Further, the risk-adjusted disparities are in fact *greater* than the raw disparities in frisk rates. To understand why, we note that stopped white pedestrians were, on average, more likely to be carrying a weapon yet less likely to be frisked than racial

⁹Our approach to estimating risk is not the only one, or even the best in any given instance. For example, one could, in theory, partially address the selection problem by requiring that some subset of individuals be randomly frisked, as is sometimes done at roadside checkpoints and airports—though this raises legal, ethical, and logistical difficulties. Alternatively, one could look for officers who frisk almost everyone they stop—similar to recent approaches for studying judicial decision-making (e.g., Dobbie, Goldin and Yang, 2018; Kleinberg et al., 2018). Both of these alternatives come with their own limitations, meaning that sensitivity analysis is still critical to credible inference.




Fig 2: *Frisk rates vs. risk, as estimated via logistic regression curves fit separately for each race group. Across risk levels, stopped Black and Hispanic pedestrians are frisked substantially more frequently than comparably risky white individuals, indicative of disparate impact.*

minorities; as a result, the risk-adjusted gap in frisk rates is even larger than the raw, unadjusted gap. Finally, we see that the kitchen-sink regression dramatically underestimates the extent of disparate impact faced by minorities. In this case, the kitchen-sink model adjusts for a variety of features—including whether the stopped individual made “furtive movements”—that are correlated with race but are poor predictors of weapon possession, skewing estimates of disparate impact.¹⁰

5.2. Sensitivity analysis. The disparities computed above aim to circumvent included-variable bias by adjusting for each individual’s estimated risk. But our estimates of risk may themselves be skewed if officers observe factors that are predictive of risk but are not recorded in our data. To account for this possibility, Figure 4 displays worst-case bounds on our estimate of disparate impact as a function of the mismatch between true and estimated risk—operationalized in terms of ϵ , as defined in Eq. (4); see the Appendix for further details on how the bounds are computed. To ease interpretation, the horizontal axis in the plot is expressed in terms of the *relative* average absolute difference between the true and estimated risks: ϵ divided by the overall weapon recovery rate among frisked individuals (1.7%). Our analysis shows that large risk-adjusted disparities remain even if we allow the true risk to differ considerably from our risk estimates. In particular, we would find large disparate impacts for both Black and Hispanic pedestrians even if the true risks differed from our risk estimates by 50% of the base rate.

It is impossible to know the precise extent to which our risk estimates differ from the true risk. To understand the plausible magnitude of the discrepancy, we conduct a simulation in which we remove a key set of risk-relevant covariates from the data, estimate risk based on

¹⁰We exclude hair color and eye color from the kitchen-sink model, since these are obvious proxies for race that are effectively unrelated to risk, and would thus be excluded in most traditional legal and statistical analyses of discrimination. As we would expect, including these variables as controls exacerbates the problem of included-variable bias, but our results show that such bias can occur even if obviously problematic variables are excluded.




Fig 3: Racial gaps in frisk rates adjusting for different sets of covariates, where the y-axis shows the percentage point difference relative to stopped white individuals. The left panel shows the raw disparities in frisk rates. As a measure of discrimination, raw disparities suffer from omitted-variable bias: there may, in theory, be legitimate reasons why Black and Hispanic pedestrians are more likely to be frisked. The middle panel shows the estimated race effects in a kitchen-sink regression, adjusting for all pre-frisk covariates. These estimates suffer from included-variable bias because they adjust for features that are correlated with race but unrelated to risk. The right panel shows the results of our risk-adjusted regression, adjusting exclusively for estimated risk of weapon possession. In all cases, estimated standard errors are less than 0.2 percentage points, and so are not visible in the plot.

the reduced information, and then measure differences between the original and new risk estimates. Specifically, we remove variables listed in two sections of the UF-250 stop forms that describe the “circumstances” prompting the encounter. These sections each consist of 10 binary variables—including, for example, “fits description”, “actions indicative of casing”, and “changing direction at sight of officer”—that are crucial for establishing the legal basis of the stop. We then compute the average absolute difference between our original risk estimates based on the full, uncensored data and the risk estimates based on the redacted data.¹¹ We find that this value is 0.7 percentage points—or 44% of the base rate—which we take as an estimate of ϵ in a scenario with severe unobserved confounding. As shown in Figure 4, this level of error (indicated by the dashed vertical lines) would yield an estimate of disparate impact that is, at a minimum, greater than 7.5 percentage points for both Black and Hispanic individuals. This sensitivity analysis suggests our results are robust to substantial unobserved confounding.

6. Discussion. We have sought to develop a simple, intuitive framework for addressing the most serious concerns of included- and omitted-variable bias in disparate impact studies. On a detailed dataset of police stops, we found that these concerns are more than hypothetical possibilities. In particular, regressions that adjust for all available covariates—in line

¹¹As above, we use gradient boosted decision trees to estimate the probability of frisk conditional on the remaining, uncensored covariates.




Fig 4: *Sensitivity of the risk-adjusted disparities in frisk decisions to mismeasurement of risk. The blue bands bound our estimates of disparate impact as a function of the average absolute difference between the true and estimated risks ϵ , relative to the base rate (1.7%). The dotted line at 44% ($\epsilon = 0.7$ p.p.) corresponds to a simulated situation with severe confounding. The small grey bands along the top and bottom of the blue bands represent 95% percentile bootstrapped confidence intervals ($N = 1000$; Zhao, Small and Bhattacharya, 2019). The step size for the grid search over average risks of the unobserved strata was 1 p.p. for all groups, and $\delta = 0.06$ p.p. (or $\gamma = 0.01$ p.p.; see the appendix) was the approximation parameter for the maximization routine.*

with common legal and statistical convention—can substantially skew estimates of disparate impact.

Throughout our analysis, we have estimated “disparate impact” by a regression coefficient on protected-group identity in a model that adjusts for estimated risk. This procedure is analogous to current practice in discrimination studies, where we simply replace the usual set of control variables with a single variable capturing risk. As seen by our formalization of disparate impact in Eq. (2), we are effectively measuring a particular weighted average of differences in decision rates across individuals of similar risk. While intuitively reasonable, this definition raises subtle questions of law and policy.

Consider, for example, Figure 2, where we plot race-specific frisk rates as a function of risk. Stopped Black and Hispanic pedestrians are frisked more often than stopped white pedestrians at every level of risk. As a result, one would find that racial minorities face disparate impact regardless of how one averages across risk levels; the precise number might change, but the qualitative conclusion would remain the same. However, comparing Black and Hispanic pedestrians, the direction of the disparity depends on the risk level one considers.¹² Low-risk Hispanic individuals are frisked more often than low-risk Black individuals, but high-risk Black individuals are frisked more often than their high-risk Hispanic counterparts. Consequently, a conclusion of disparate impact between Black and Hispanic pedestrians would depend heavily on the precise definition applied. The analysis is further

¹²Such a comparison between minority groups is unusual in disparate impact cases, but it illustrates the underlying theoretical issue.

complicated if the risk distributions differ substantially between groups. If, hypothetically, stopped Hispanic pedestrians were mostly low-risk and stopped Black pedestrians mostly high-risk, majorities of both groups could argue that they were treated more harshly than members of the other group who were equally likely to be carrying a weapon.¹³

The crossing of risk curves that we see in Figure 2 is a potentially widespread phenomenon, and, to our knowledge, disparate impact law has not yet resolved the underlying conceptual ambiguity it invokes. Many discussions of disparate impact tacitly assume that policies either consistently harm or help groups defined by protected traits. Such thinking can be seen in the original *Griggs* ruling, where the Supreme Court aimed to proscribe policies that acted as “built-in headwinds” for racial minorities. But, formally, disparate impact law concerns facially race-neutral policies, not intentional discrimination, and there is no theoretical or empirical guarantee that such policies will adversely impact all members of a particular group.

A related issue is the extent to which concern for unjustified disparities compels decision makers to act optimally. For example, Figure 2 suggests that officers are only marginally responsive to risk, with the lowest-risk individuals still frisked more than 40% of the time. If, instead, officers frisked only the people with high probability of carrying a weapon, they could frisk far fewer individuals—and, in particular, far fewer minority individuals—while recovering almost the same number of weapons (Goel, Rao and Shroff, 2016b). A more efficient frisk strategy could thus reduce the burdens of policing on racial minorities while still maintaining public safety. Such efficiency is indeed one of the aims of statistical risk assessment tools that are now used in the criminal justice system and beyond to guide high-stakes decisions (Chouldechova et al., 2018; Corbett-Davies et al., 2023; Goel et al., 2021; Monahan and Skeem, 2016; Shroff, 2017). If these tools are shown to reduce racial disparities, are policymakers obliged—legally or ethically—to adopt them? The role of efficiency in disparate impact claims has largely gone unanswered by the courts, and adds yet another subtlety to defining and measuring disparities.

By foregrounding the role of risk in understanding disparities, we have aimed to clarify some of the thorny conceptual issues at the heart of disparate impact analysis. While there are still important unresolved questions, we believe that our statistical approach provides practitioners with a tractable way to assess disparities in many domains while avoiding some important pitfalls of traditional methods. Looking forward, we hope this work spurs further theoretical and empirical research on discrimination at the intersection of statistics, economics, law, and public policy.

¹³Some scholars have similarly investigated interactions between race and other decision-making criteria. For example, Espenshade, Chung and Walling (2004) find that preferences for underrepresented minorities in college admissions is greatest for applicants with SAT scores in the 1200–1300 range, and the effect is attenuated at lower scores. That analysis, however, found no score ranges where minority applicants faced an absolute disadvantage relative to white applicants with equal scores. We do not know of any research that has found a change in the direction of the disparities like we see between Black and Hispanic pedestrians in Figure 2.

APPENDIX A: SENSITIVITY ANALYSIS

In what follows, we formally justify the sensitivity analysis laid out in Section 4, deriving algorithms to efficiently solve the optimization problem it gives rise to and rigorously verifying their correctness and run-times. This discussion is structured in three parts. In the first part, by introducing a new low-dimensional family of parameters, we show how to reduce the base optimization problem of Eq. (6) into the simplified problem of Eq. (7). In the second and third parts, we show how to efficiently find the maximum and minimum values of the simplified optimization problem, respectively.

A.1. Simplifying the objective. To simplify the optimization problem in Eq. (6), we can avail ourselves of the fact that OLS has a closed-form solution. We augment the notation in Section 4 as follows:

$$\mathcal{G}_j = \mathcal{S}_{j,0} \cup \mathcal{S}_{j,1} = \{i : C_i = j\},$$

i.e., \mathcal{G}_j denotes the set of indices belonging to the j -th group.

LEMMA A.1. *Let $n_j = |\mathcal{G}_j|$ denote the number of individuals in the j -th group; let σ_j denote the search rate in group j , i.e., $\frac{1}{n_j} \sum_{i \in \mathcal{G}_j} A_i$; and let r_j and t_j denote*

$$r_j = \sum_{i \in \mathcal{S}_{j,1}} R_i, \quad t_j = \sum_{i \in \mathcal{S}_{j,0}} R_i,$$

i.e., the total risk of the searched and unsearched individuals belonging to the j -th group, respectively.

If R_i is constant within each of the m groups, then the OLS estimate need not exist due to colinearity. Otherwise, the OLS estimate of the coefficients $\hat{\beta}$ in Eq. (3) satisfies

$$(8) \quad \hat{\beta} = \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_m \\ 0 \end{bmatrix} + \frac{\frac{1}{n} \left[\sum_{j=1}^m \sigma_j \cdot t_j - (1 - \sigma_j) \cdot r_j \right]}{\frac{1}{n} \left[\sum_{i=1}^n R_i^2 - \sum_{j=1}^m \frac{(r_j + t_j)^2}{n_j} \right]} \begin{bmatrix} \frac{r_1 + t_1}{n_1} \\ \vdots \\ \frac{r_m + t_m}{n_m} \\ -1 \end{bmatrix}.$$

PROOF. Recall that our design matrix and outcome variables take the form

$$\mathbf{X} = \begin{bmatrix} \mathbf{1}(C_1 = 1) & \dots & \mathbf{1}(C_1 = m) & R_1 \\ \vdots & \ddots & \vdots & \vdots \\ \mathbf{1}(C_n = c_1) & \dots & \mathbf{1}(C_n = m) & R_n \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} A_1 \\ \vdots \\ A_n \end{bmatrix}.$$

The OLS estimate of the coefficients $\hat{\beta}$ is then simply $(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y}$.

Note that by definition, we have that

$$\mathbf{X}^\top \mathbf{X} = \begin{bmatrix} n_1 & & & \sum_{i \in \mathcal{G}_1} R_i \\ & \ddots & & \vdots \\ & & n_m & \sum_{i \in \mathcal{G}_m} R_i \\ \sum_{i \in \mathcal{G}_1} R_i & \dots & \sum_{i \in \mathcal{G}_m} R_i & \sum_{i=1}^n R_i^2 \end{bmatrix} = \begin{bmatrix} n_1 & & & r_1 + t_1 \\ & \ddots & & \vdots \\ & & n_m & r_m + t_m \\ r_1 + t_1 & \dots & r_m + t_m & \sum_{i=1}^n R_i^2 \end{bmatrix}.$$

Similarly,

$$\mathbf{X}^\top \mathbf{Y} = \begin{bmatrix} \sum_{i \in \mathcal{G}_1} A_i \\ \vdots \\ \sum_{i \in \mathcal{G}_m} A_i \\ \sum_{i=1}^n R_i \cdot A_i \end{bmatrix} = \begin{bmatrix} n_1 \cdot \sigma_1 \\ \vdots \\ n_m \cdot \sigma_m \\ \sum_{j=1}^m r_j \end{bmatrix}.$$

Note that $\mathbf{X}^\top \mathbf{X}$ has the form

$$\begin{bmatrix} D & v^\top \\ v & \sum_{i=1}^n R_i^2 \end{bmatrix}$$

where D is diagonal. Then, we see that since D is invertible, it follows that $\mathbf{X}^\top \mathbf{X}$ is invertible if the Schur complement of the bottom right-hand entry is invertible, i.e., non-zero; that is, if

$$\sum_{i=1}^n R_i^2 - \frac{\sum_{j=1}^m (r_j + t_j)^2}{n_j} \neq 0.$$

Recalling the definitions of r_j and t_j , we see that the left-hand expression equals

$$\sum_{j=1}^m \sum_{i \in \mathcal{G}_j} \left(R_i - \frac{s_j + t_j}{n_j} \right)^2 = \sum_{j=1}^m n_j \cdot \text{VAR}((R_i)_{i \in \mathcal{G}_j}).$$

Now, the variance of a set is only zero if all the elements of the set are equal, i.e., if for all $i \in \mathcal{G}_j$, R_i equals the average risk in that group, $\frac{r_j + t_j}{n_j}$.

Assume this is not the case, i.e., that $\mathbf{X}^\top \mathbf{X}$ is invertible. Since $\mathbf{X}^\top \mathbf{X}$ is an arrowhead matrix, we can invert it using the Sherman-Morrison Formula. In particular, we obtain that

$$(9) \quad (\mathbf{X}^\top \mathbf{X})^{-1} = \begin{bmatrix} \frac{1}{n_1} & & & \\ & \ddots & & \\ & & \frac{1}{n_m} & \\ & & & 0 \end{bmatrix} + \frac{1}{\sum R_i^2 - \sum_{j=1}^m \frac{(r_j + t_j)^2}{n_j}} \begin{bmatrix} \frac{r_1 + t_1}{n_1} \\ \vdots \\ \frac{r_m + t_m}{n_m} \\ -1 \end{bmatrix} \begin{bmatrix} \frac{r_1 + t_1}{n_1} \\ \vdots \\ \frac{r_m + t_m}{n_m} \\ -1 \end{bmatrix}^\top.$$

Combining this with the expression derived for \mathbf{Y} above and dividing the numerator and denominator by n yields that

$$(10) \quad \hat{\beta} = \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_m \\ 0 \end{bmatrix} + \frac{\frac{1}{n} \left[\sum_{j=1}^m \sigma_j \cdot t_j - (1 - \sigma_j) \cdot r_j \right]}{\frac{1}{n} \left[\sum R_i^2 - \sum_{j=1}^m \frac{(r_j + t_j)^2}{n_j} \right]} \begin{bmatrix} \frac{r_1 + t_1}{n_1} \\ \vdots \\ \frac{r_m + t_m}{n_m} \\ -1 \end{bmatrix}.$$

□

We note that Lemma A.1 provides the following useful condition on when the base problem has a meaningful solution:

COROLLARY A.2. *The maximum and minimum values of the base problem are both finite if and only if it is not the case that $\ell_i \leq \rho_j \leq u_i$ for all $i \in \mathcal{G}_i$ and $j = 1, \dots, m$, and in addition*

$$\frac{1}{n} \sum_{j=1}^m \sum_{i \in \mathcal{G}_j} |\hat{R}_i - \rho_j| \leq \epsilon.$$

PROOF. The proof follows immediately from noting that since the averages of the risks on the observed strata $\mathcal{S}_{j,1}$ are fixed at ρ_j , the risks for the whole j -th group can only all be equal if the risks of the observed *and* the unobserved individuals are uniformly equal to ρ_j . □

From the analytical expression for the regression coefficients given in Lemma A.1, we can derive the following useful corollary.

COROLLARY A.3. *Let F be a compact subset of \mathbb{R}^n which is not constant within all m groups. Let $\hat{\beta}$ be as in Eq. (3), and r_j and t_j as in Lemma A.1.¹⁴ Then, we have that for all $\mathbf{R} \in F$,*

$$\hat{\beta}_j - \hat{\beta}_1 \leq \frac{\sigma_j}{n_j} - \frac{\sigma_1}{n_1} + \max_{\mathbf{R} \in F} \left(\frac{\frac{1}{n} \left[\sum_{j=1}^m \sigma_j \cdot t_j - (1 - \sigma_j) \cdot r_j \right]}{\frac{1}{n} \left[\sum_{i=1}^n R_i^2 - \sum_{j=1}^m \frac{(r_j + t_j)^2}{n_j} \right]} \left(\frac{r_j + t_j}{n_j} - \frac{r_1 + t_1}{n_1} \right) \right),$$

and

$$\hat{\beta}_j - \hat{\beta}_1 \geq \frac{\sigma_j}{n_j} - \frac{\sigma_1}{n_1} + \min_{\mathbf{R} \in F} \left(\frac{\frac{1}{n} \left[\sum_{j=1}^m \sigma_j \cdot t_j - (1 - \sigma_j) \cdot r_j \right]}{\frac{1}{n} \left[\sum_{i=1}^n R_i^2 - \sum_{j=1}^m \frac{(r_j + t_j)^2}{n_j} \right]} \left(\frac{r_j + t_j}{n_j} - \frac{r_1 + t_1}{n_1} \right) \right).$$

The proof follows directly from Eq. (8).

The following corollary completes the simplification of the optimization problem in Eq. (6) in the case where the average true risk in the unobserved strata is also assumed known (“the parameterized problem” of Section 4), i.e., of

$$(11) \quad \begin{aligned} & \text{Optimize}_{\mathbf{R} \in \mathbb{R}^n} && \hat{\beta}_{j^*} - \hat{\beta}_1 \\ & \text{s.t.} && \frac{1}{n} \sum_{i=1}^n |R_i - \hat{R}_i| \leq \epsilon, \\ & && \frac{1}{|\mathcal{S}_{j,1}|} \cdot \sum_{i \in \mathcal{S}_{j,1}} R_i = \rho_j, \quad (j = 1, \dots, m) \\ & && \frac{1}{|\mathcal{S}_{j,0}|} \cdot \sum_{i \in \mathcal{S}_{j,0}} R_i = \tau_j, \quad (j = 1, \dots, m) \\ & && R_i \leq u_i, \quad (i = 1, \dots, n) \\ & && R_i \geq \ell_i. \quad (i = 1, \dots, n) \end{aligned}$$

COROLLARY A.4. *The optima of the parameterized problem in Eq. (11), if they are both well-defined, are equivalently the solutions of the simplified problem in Eq. (7).*

PROOF. In the notation of Lemma A.1 and Corollary A.3, we note that the constraints of the parameterized problem fix $r_j = \rho_j$ and $t_j = \tau_j$ for all feasible \mathbf{R} .

Consequently, Corollary A.3, the objective function of the parameterized problem has the form

$$(12) \quad a + \frac{b}{x - c},$$

where

$$\begin{aligned} a &= \frac{\sigma_j}{n_j} - \frac{\sigma_1}{n_1}, \\ b &= \frac{1}{n} \left[\sum_{j=1}^m \sigma_j \cdot \tau_j - (1 - \sigma_j) \cdot \rho_j \right] \cdot \left(\frac{\rho_j + \tau_j}{n_j} - \frac{\rho_1 + \tau_1}{n_1} \right), \\ c &= \frac{1}{n} \sum_{j=1}^m \frac{(\rho_j + \tau_j)^2}{n_j}, \\ x &= \frac{1}{n} \sum_{i=1}^n R_i^2. \end{aligned}$$

¹⁴In particular, strictly speaking, $\hat{\beta}$, r_j , and t_j should be written $\hat{\beta}(\mathbf{R})$, $r_j(\mathbf{R})$, and $t_j(\mathbf{R})$, since they depend on \mathbf{R} .

By continuity and the fact that the feasible region is convex, and hence connected, and compact, the range of the objective in Eq. (7) over the feasible region is a closed interval I . As shown in the proof of Lemma A.1 and Corollary A.2, $x \geq c$, and $x > c$ if the optima are both well-defined. Consequently, we can assume the interval is strictly positive, i.e., $I \subseteq (0, \infty)$. The derivative of Eq. (12) exists everywhere on I and is, moreover, non-zero and continuous, and so $a + \frac{b}{x-c}$ is strictly monotone on I . Therefore, the maximum and minimum must occur at the endpoints, i.e., at the maximum and minimum values of x . \square

We note that since b is not guaranteed to be positive, the maximum of the simplified problem in Eq. (7) does not necessarily correspond to the the maximum of the base problem in Eq. (6) and *vice versa*.

A.2. Minimization. In the case of minimization, the simplified problem in Eq. (7) defines a convex optimization problem, which, as noted in Section 4 above, can consequently be solved efficiently using standard interior-point methods. However, these results can be improved in a problem-specific way to yield a simpler, and, in practice, more computationally tractable method of finding exact solutions. Moreover, they parallel the solution method for maximization, which is not a convex problem, and hence to which standard interior point methods cannot be applied.

A.2.1. *Optimizing over a single stratum.* We begin by solving a simpler version of the problem, where we restrict to a single stratum. In the case of minimization—as in the case of maximization, as we show below—solutions to the minimization problem must be in a certain normal form. While searching over all risk vectors is prohibitively difficult, searching over normal form risk vectors can be done straightforwardly in linear time.

DEFINITION A.5 (Minimization normal form). We say that a risk vector \mathbf{R} is in *minimization normal form* if there exist thresholds $t_{a,j}^{\text{LWR}}$ and $t_{a,j}^{\text{UPR}}$ such that for $i \in \mathcal{S}_{a,j}$,

$$R_i = \begin{cases} \min(u_i, t_{a,j}^{\text{LWR}}) & \hat{R}_i \leq t_{a,j}^{\text{LWR}}, \\ \hat{R}_i & t_{a,j}^{\text{LWR}} \leq \hat{R}_i \leq t_{a,j}^{\text{UPR}}, \\ \max(\ell_i, t_{a,j}^{\text{UPR}}) & t_{a,j}^{\text{UPR}} \leq \hat{R}_i, \end{cases}$$

and, in addition, $t_{a,j}^{\text{UPR}} - t_{a,j}^{\text{LWR}}$ is equal to some Δ for all $a = 0, 1$ and $j = 1, \dots, m$.

In the context of a single stratum, we refer to t^{UPR} and t^{LWR} for notational simplicity. We also assume that $\mathcal{S}_{a,j} = \{1, \dots, n\}$ for the same reason.

Minimization normal form results from “pushing” \hat{R}_i up to t^{LWR} or down to t^{LWR} as far as possible. An example of a risk vector in minimization normal form is shown in Figure 1. As the name suggests, risk vectors minimizing the objective must be in minimization normal form.

LEMMA A.6. *Consider the optimization problem (“the single stratum minimization problem”)*

$$(13) \quad \begin{aligned} & \underset{\mathbf{R} \in \mathbb{R}^n}{\text{Minimize}} && \frac{1}{n} \sum_{i=1}^n R_i^2 \\ & \text{s.t.} && \frac{1}{n} \sum_{i=1}^n |R_i - \hat{R}_i| \leq \epsilon, \\ & && \frac{1}{n} \sum_{i=1}^n R_i = \mu, \\ & && R_i \leq u_i, \quad (\forall i) \\ & && R_i \geq \ell_i. \quad (\forall i) \end{aligned}$$

Suppose that \mathbf{R}^* is a minimizer. Then, \mathbf{R}^* is unique and in minimization normal form. Moreover, \mathbf{R}^* exhausts the L_1 budget, in the sense that either $\frac{1}{n}\|\mathbf{R}^* - \hat{\mathbf{R}}\|_1 = \epsilon$ or $t^{\text{LWR}} = t^{\text{UPR}} = \mu$.

PROOF. The proof proceeds by examining the first-order KKT conditions to derive weak necessary conditions that solutions must satisfy. These conditions are then strengthened to minimization normal form by directly comparing the objective function at different points satisfying the weak conditions. Finally, we show that the minimum is the point satisfying the “strengthened” conditions which is most distant from $\hat{\mathbf{R}}$, i.e., that exhausts the budget.

The problem is simplified by rewriting the final constraint as a collection of linear (and hence everywhere differentiable) constraints. In particular, $\frac{1}{n}\|\mathbf{R} - \hat{\mathbf{R}}\|_1 \leq \epsilon$ is equivalent to the 2^n constraints of the form $\mathbf{S}^\top(\mathbf{R} - \hat{\mathbf{R}}) \leq \frac{\epsilon}{n}$ where for all $\mathbf{S} \in \{-1, 1\}^n$. We assume throughout, without loss of generality, that $\ell_i < u_i$.

Weak conditions. Let \mathbf{e}_i denotes the i -th standard basis vector and $\mathbf{1} = \sum_{i=1}^n \mathbf{e}_i$. Note that the gradient of the objective is $\frac{2}{n} \cdot \mathbf{R}$, the gradient of the first constraint is \mathbf{S} ; the second, $\frac{1}{n} \cdot \mathbf{1}$; the third \mathbf{e}_i ; and the fourth $-\mathbf{e}_i$. The first-order necessary KKT conditions therefore require that¹⁵

$$(14) \quad 2 \cdot \mathbf{R}^* - \lambda \cdot \mathbf{1} + \sum_{i=1}^n \mu_{0,i} \cdot \mathbf{e}_i - \sum_{i=1}^n \mu_{1,i} \cdot \mathbf{e}_i + \sum_{k=1}^{2^n} \nu_k \cdot \mathbf{S}_k = 0$$

for some arbitrary λ , and for non-negative $\mu_{0,i}$, $\mu_{1,i}$, and ν_k for $k = 1, \dots, 2^n$ satisfying complementary slackness, i.e., such that $\mu_{0,i} \cdot (x_i - u_i) = 0$, $\mu_{1,i} \cdot (\ell_i - x_i) = 0$, and $\nu_k \cdot (\mathbf{S}_k^\top(\mathbf{R} - \hat{\mathbf{R}}) - \epsilon) = 0$ for all $i = 1, \dots, n$ and $k = 1, \dots, 2^n$.

These consequences allow us to prove the following weak characterization of \mathbf{R} . Let $\Delta = \sum_{k=1}^{2^n} \nu_k$. Then, for all i ,

$$(15) \quad R_i^* \in \left\{ \hat{R}_i, u_i, \ell_i, \frac{1}{2} \cdot [\lambda - \Delta], \frac{1}{2} \cdot [\lambda + \Delta] \right\}.$$

For, from (14), we get that

$$(16) \quad 2 \cdot R_i^* - \lambda + \mu_{0,i} - \mu_{1,i} + \sum_{k=1}^n \nu_k \cdot S_{k,i} = 0.$$

If $R_i^* = \hat{R}_i$, there is nothing to prove. Therefore, assume that $R_i^* \neq \hat{R}_i$. We note that by complementary slackness, $\nu_k > 0$ only if $S_{k,i} \cdot (R_i^* - \hat{R}_i) \geq 0$ for all $i = 1, \dots, n$. For, supposing without loss of generality that $R_i^* > \hat{R}_i$, if $S_{k,i} \cdot (R_i^* - \hat{R}_i) < 0$, then, since $\mathbf{S} + 2\mathbf{e}_i \in \{-1, 1\}^n$, we would have that

$$(\mathbf{S} + 2\mathbf{e}_i)^\top(\mathbf{R}^* - \hat{\mathbf{R}}) = \epsilon + 2(R_i^* - \hat{R}_i) > \epsilon,$$

which violates the constraints. Hence $S_{k,i} = S_{k',i}$ for all k and k' such that $\nu_k, \nu_{k'} > 0$. Therefore, our expression simplifies to

$$2R_i^* - \lambda + \mu_{0,i} - \mu_{1,i} + \Delta \cdot s,$$

where $s = \pm 1$. By complementary slackness, if $\mu_{0,i} > 0$, then $R_i^* = u_i$; if $\mu_{1,i} > 0$, then $R_i^* = \ell_i$. Therefore, we need only consider the case where $\mu_{0,i} = \mu_{1,i} = 0$, i.e., where

$$2 \cdot R_i^* - \lambda_{a,j} + s \cdot \Delta = 0,$$

whence $R_i^* = \frac{1}{2}[\lambda_{a,j} + s \cdot \Delta]$. Therefore, for all i , (15) holds.

¹⁵For notational simplicity, we have multiplied through by n and absorbed constants into the corresponding Lagrange multipliers.

Strengthening conditions. Next, we strengthen (15) to minimization normal form. First, we show that if $R_{i_0}^* > \hat{R}_{i_0}$ and $R_{i_0}^* > R_{i_1}^*$, then $R_{i_1}^* = u_{i_1}$. For, suppose not. Then, there exists some $\delta > 0$ such that $R_{i_0}^* - \delta > \hat{R}_{i_0}$ and $R_{i_1}^* + \delta < u_{i_1}$. Define $\mathbf{R}' = \mathbf{R}^* + \delta \cdot (\mathbf{e}_{i_1} - \mathbf{e}_{i_0})$. Then, we have that \mathbf{R}' is feasible—in particular, $\|\mathbf{R}^* - \hat{\mathbf{R}}\|_1 \geq \|\mathbf{R}' - \hat{\mathbf{R}}\|_1$ —and

$$\begin{aligned} \sum_{i=1}^n (R_i^*)^2 - (R_i')^2 &= (R_{i_0}^*)^2 - ((R_{i_0}^*)^2 + \delta^2 - 2 \cdot R_{i_0}^* \cdot \delta) + (R_{i_1}^*)^2 \\ &\quad - ((R_{i_1}^*)^2 + \delta^2 + 2 \cdot (R_{i_1}^*) \cdot \delta) + \sum_{i \neq i_0, i_1} (R_i^*)^2 - (R_i^*)^2 \\ &= 2 \cdot \delta \cdot (\delta + R_{i_0}^* - R_{i_1}^*) \\ &> 0, \end{aligned}$$

where the inequality follows from the fact that $R_{i_0}^* > R_{i_1}^*$. Therefore \mathbf{R}^* is not a minimum, contrary to hypothesis. In exactly the same manner, we see that if $R_{i_0}^* < \hat{R}_{i_0}$ and $R_{i_1}^* < R_{i_0}^*$, then $R_{i_1}^* = \ell_{i_1}$.

Combining this with (15) yields that for all i , if $R_i^* \leq \frac{1}{2} \cdot [\lambda - \Delta]$, then $R_i^* = \min(u_i, \frac{1}{2} \cdot [\lambda - \Delta])$; if $R_i^* \geq \frac{1}{2} \cdot [\lambda + \Delta]$, then $R_i^* = \max(\ell_i, \frac{1}{2} \cdot [\lambda + \Delta])$; and otherwise $R_i^* = \hat{R}_i$. Taking $t^{\text{lwr}} = \frac{1}{2} \cdot [\lambda - \Delta]$ and $t^{\text{upr}} = \frac{1}{2} \cdot [\lambda + \Delta]$ gives the result.

Uniqueness and budget exhaustion. If the budget is not exhausted, then, by complementary slackness, $\nu_k = 0$ for all $k = 1, \dots, 2^k$, whence $\Delta = 0$. This immediately implies that $t^{\text{lwr}} = t^{\text{upr}} = \lambda$.

To see uniqueness, suppose that $\mathbf{R}^{(0)}$ and $\mathbf{R}^{(1)}$ are two distinct minima, both in normal form, with corresponding thresholds t_k^{lwr} and t_k^{upr} for $k = 0, 1$. We note that since $\sum_{i=1}^n R_i^{(0)} = \sum_{i=1}^n R_i^{(1)} = \mu$, we must have, without loss of generality, that $t_0^{\text{lwr}} \leq t_1^{\text{lwr}}$ and $t_0^{\text{upr}} \geq t_1^{\text{upr}}$, with at least one of the two inequalities strict. However, it follows that $\|\mathbf{R}^{(0)} - \hat{\mathbf{R}}\|_1 > \|\mathbf{R}^{(1)} - \hat{\mathbf{R}}\|_1$, which is impossible by the preceding paragraph. \square

Lemma A.6 suggests a natural algorithm for solving the optimization problem in Eq. (7): we can sweep over risk vectors in minimization normal form, drawing the thresholds together until either (1) the L_1 budget is exhausted, or (2) the two thresholds become equal. In doing so, we must be careful to ensure that we only sweep over vectors in minimization normal form with the correct total risk.

To see how to do this, consider an index i to be “active” if one of the thresholds is between ℓ_i and u_i . If k^{lwr} represents the number of indices that are active because of t^{lwr} then a one unit increase in t^{lwr} increases the sum of the risks by k^{lwr} . Therefore, the rates of increase of the two thresholds, r^{lwr} and r^{upr} , must satisfy

$$r^{\text{lwr}} \cdot k^{\text{lwr}} = r^{\text{upr}} \cdot k^{\text{upr}}.$$

If we additionally impose that $r^{\text{lwr}} + r^{\text{upr}} = 1$, so that δ decreases at a constant rate, then we have that¹⁶

$$(17) \quad r^{\text{lwr}} = \frac{k^{\text{upr}}}{k^{\text{lwr}} + k^{\text{upr}}}, \quad r^{\text{upr}} = \frac{k^{\text{lwr}}}{k^{\text{lwr}} + k^{\text{upr}}}$$

Since these rates depend on the number of “active” indices, we must update them every time an index is activated or deactivated as the two thresholds move toward one another.

¹⁶If both k^{lwr} and k^{upr} are zero, we can without loss of generality take both rates to be $\frac{1}{2}$.

Thus, let \mathbf{R}_Δ be the unique solution we can associate with a given gap δ between t^{lwr} and t^{upr} .¹⁷ Then, we can study two functions:

$$\epsilon(\Delta) = \|\mathbf{R}_\Delta - \hat{\mathbf{R}}\|_1, \quad \Sigma(\Delta) = \sum_{i=1}^n R_{\Delta,i}^2 - \hat{R}_i^2.$$

By the preceding discussion we readily derive that $\epsilon(\Delta)$ is a piecewise linear function of Δ , whose slope at a given value of Δ is equal to¹⁸

$$K = 2 \cdot \frac{k^{\text{lwr}} \cdot k^{\text{upr}}}{k^{\text{lwr}} + k^{\text{upr}}},$$

and $\Sigma(\Delta)$ is a piecewise quadratic function of Δ satisfying

$$\begin{aligned} \Sigma(\Delta + t) &= \Sigma(\Delta) + k^{\text{lwr}} \cdot [(t^{\text{lwr}} + r^{\text{lwr}} \cdot t)^2 - (t^{\text{lwr}})^2] + k^{\text{upr}} \cdot [(t^{\text{upr}} - r^{\text{upr}} \cdot t)^2 - (t^{\text{lwr}})^2] \\ &= \Sigma(\Delta) + k^{\text{lwr}} \cdot r^{\text{lwr}} \cdot t \cdot (r^{\text{lwr}} \cdot t + 2t^{\text{lwr}}) + k^{\text{upr}} \cdot r^{\text{upr}} \cdot t \cdot (r^{\text{upr}} \cdot t - 2t^{\text{lwr}}) \\ &= \Sigma(\Delta) - 2 \cdot \frac{k^{\text{lwr}} \cdot k^{\text{upr}}}{k^{\text{lwr}} + k^{\text{upr}}} \cdot (t^{\text{upr}} - t^{\text{lwr}}) \cdot t + \frac{k^{\text{lwr}}(k^{\text{upr}})^2 + k^{\text{upr}}(k^{\text{lwr}})^2}{(k^{\text{lwr}} + k^{\text{upr}})^2} \cdot t^2 \\ &= \Sigma(\Delta) - K\Delta t + \frac{K}{2}t^2, \end{aligned}$$

as long as t is sufficiently small that the number of active indices does not change. It follows that $\epsilon(\Delta)$ and $\Sigma(\Delta)$ are fully determined by the following collections:

$$\mathbf{\Delta} = (\Delta_1, \dots, \Delta_N), \quad 1 = \Delta_1 < \dots < \Delta_N = 0,$$

where Δ_k represents the k -th value of Δ at which k^{lwr} and k^{upr} change; and

$$\mathbf{K} = (K_1, \dots, K_N), \quad K_1 = 0, \quad K_N = 0,$$

where K_k denotes the value of K beginning at $\Delta = \Delta_k$. From $\mathbf{\Delta}$ and \mathbf{K} , the Δ^* such that $\|\mathbf{R}_{\Delta^*} - \hat{\mathbf{R}}\|_1 = \epsilon$ can be calculated in linear time. For completeness, this algorithm is given in Algorithm 2.

Thus, a complete solution for the optimization problem for a single stratum requires only calculating $\mathbf{\Delta}$ and \mathbf{K} . As described above, this can be calculated by moving the thresholds toward each other at the prescribed rates r^{lwr} and r^{upr} , updating the rates each time an index is activated or deactivated, until the gap between the thresholds is zero. Similar to evaluating $\epsilon(\Delta)$ and $\Sigma(\Delta)$, this can be completed in linear time, as described in Algorithm 4. (Note that if $\frac{1}{n} \sum_{i=1}^n \hat{R}_i \neq \mu$, it may be necessary to do a preprocessing step, also in linear time, as shown in Algorithm 3.)

A.2.2. Optimizing over all strata. With a complete solution to the problem of minimizing the sum of squares for a single stratum, the problem of minimizing the sum of squares across all strata is straightforward. We begin by characterizing solutions to the minimization problem in the general case.

LEMMA A.7. *Consider the simplified minimization problem in Eq. (7). Suppose that \mathbf{R}^* is a solution. Then, the restriction of \mathbf{R}^* to any stratum is in minimization normal form. Moreover, \mathbf{R}^* exhausts the L_1 budget, in that either $\frac{1}{n} \|\mathbf{R} - \hat{\mathbf{R}}\|_1 = \epsilon$ or $t^{\text{lwr}} = t^{\text{upr}} = \mu_{a,j}$, where $\mu_{a,j}$ equals either ρ_j or τ_j depending on whether a equals 1 or 0.*

¹⁷We note that by Lemma A.6, for every δ there is a unique \mathbf{R}_δ , but that the reverse is not necessarily true, since there may not be any i such that $R_i = t^{\text{lwr}}$ or t^{upr} , in which case different thresholds can be chosen without altering the underlying risk vector.

¹⁸Or zero if both k^{lwr} and k^{upr} are zero.

PROOF. The proof is virtually identical to the proof of Lemma A.6. The only difference is that the first-order necessary KKT conditions take the form

$$(18) \quad 2 \cdot \mathbf{R}^* - \left(\sum_{j=1}^m \lambda_{0,j} \cdot \mathbf{1}_{0,j} + \lambda_{1,j} \cdot \mathbf{1}_{1,j} \right) + \sum_{i=1}^n \mu_{0,i} \cdot \mathbf{e}_i - \sum_{i=1}^n \mu_{1,i} \cdot \mathbf{e}_i + \sum_{k=1}^{2^n} \nu_k \cdot \mathbf{S}_k = 0,$$

where $\mathbf{1}_{a,j} = \sum_{C_i=c_j, A_i=a} \mathbf{e}_i$. Restricting to a single i gives the following minor variant of Eq. (16):

$$2 \cdot R_i^* - \lambda_{a,j} + \mu_{0,i} - \mu_{1,i} + \sum_{k=1}^n \nu_k \cdot S_{k,i} = 0.$$

The proof then proceeds identically; the only difference is to note that while $\lambda_{a,j}$ varies by stratum, $\Delta = \sum_{k=1}^{2^n} \nu_k$ does not. \square

It follows that solving the minimization problem across all strata can be carried out in almost the same way as across a single stratum:

1. Construct piecewise linear functions $\epsilon_{a,j}(\Delta)$ and piecewise quadratic functions $\Sigma_{a,j}(\Delta)$ for each stratum;
2. Note that the sums $\sum_{a=0}^1 \sum_{j=1}^m \epsilon_{a,j}(\Delta)$ and $\sum_{a=0}^1 \sum_{j=1}^m \Sigma_{a,j}(\Delta)$ are also piecewise linear and quadratic, respectively;
3. Find Δ^* such that $\sum_{a=0}^1 \sum_{j=1}^m \epsilon_{a,j}(\Delta^*) = \epsilon$, and evaluate $\sum_{a=0}^1 \sum_{j=1}^m \Sigma_{a,j}(\Delta^*)$.

Consequently, beyond the machinery established in Algorithms 2, 3, and 4, we only need a way of calculating the sums of these piecewise functions. This can be carried out straightforwardly using a variation on the standard merge-sort algorithm, which we give in Algorithm 5 for completeness.

Putting this all together, we obtain the following lemma summarizing the results of this section.

LEMMA A.8. *If ℓ , $\hat{\mathbf{R}}$, and \mathbf{u} have been sorted, then there exists an $O(\log(m) \cdot n)$ algorithm solving the simplified minimization problem (Eq. (7)).*

PROOF. The proof is straightforward. We begin by noting that Algorithm 2 requires linear time: the index i increases by one in the while loop on lines 6 through 10 on each iteration, and must be less than the length of the input. Similarly, in Algorithm 3, the while loop on lines 11 through 22 increments i^{lwr} by one on each loop, and limits it to three times the length of the input. The same analysis applies to Algorithm 4, in which either i^{upr} or i^{lwr} increases by one on each iteration. We note that the run time of Algorithm 5 is proportional to the sum of the lengths of its inputs. Using a standard divide-and-conquer approach, we can accomplish the combination of the $2m$ strata in m total calls, where each element of the input appears in $1 + \lceil \log_2(m) \rceil$ of the function calls, for a total time complexity of $O(\log(m) \cdot n)$. Putting this all together, the total complexity is $O(\log(m) \cdot n)$. \square

A.3. Maximization. Even in the simplified problem in Eq. (7), maximization is not a convex problem. In general, maximizing a quadratic objective over a linearly constrained convex set is NP-hard (Sahni, 1974). However, the restricted forms of the objective and constraints allow us efficiently to solve the optimization problem exactly in the case of a single stratum and approximately in the case of multiple strata.




Fig 5: *Maximization normal form.* Black dots represent estimated risks \hat{R}_i , the error bars represent the allowable range $[\ell_i, u_i]$, and red dots indicate the corresponding value of the normal form vector R_i . The pivots are indicated on the x-axis.

A.3.1. *Optimizing over a single stratum.* As in the case of minimization, risk vectors that maximize the single-stratum maximization problem have a special normal form.

DEFINITION A.9 (Maximization normal form). We say that a risk vector \mathbf{R} is in *maximization normal form* if there exist indices (“pivots”) i^{lwr} and i^{upr} such that

$$R_i = \begin{cases} \ell_i & i < i^{\text{lwr}}, \\ \hat{R}_i & i^{\text{lwr}} < i < i^{\text{upr}}, \\ u_i & i^{\text{upr}} < i. \end{cases}$$

We additionally require that the risk values at the pivots themselves satisfy $R_{i^{\text{lwr}}} \in [\ell_{i^{\text{lwr}}}, \hat{R}_{i^{\text{lwr}}}]$ and $R_{i^{\text{upr}}} \in [\ell_{i^{\text{upr}}}, \hat{R}_{i^{\text{upr}}}]$.¹⁹

Maximization normal form is similar to minimization normal form, except that instead of begin pushed toward the thresholds, values at indices *below* the lower pivot are pushed down, and values at indices *above* the upper pivot are pushed up. An illustration of maximum normal form is given in Figure 5.

An important difference from the case of minimization is that studying the KKT conditions does not yield the normal form directly; instead, it yields a weaker characterization.

DEFINITION A.10 (Weak normal form). We say that a risk vector is in *weak normal form* if there exist thresholds t^{lwr} and t^{upr} such that

$$R_i = \begin{cases} \ell_i & R_i < t^{\text{lwr}}, \\ t^{\text{lwr}} & R_i = t^{\text{lwr}}, \\ \hat{R}_i & t^{\text{lwr}} < R_i < t^{\text{upr}}, \\ t^{\text{upr}} & R_i = t^{\text{upr}}, \\ u_i & t^{\text{upr}} < R_i, \end{cases}$$

¹⁹Since it involves the ordering of the indices as well as the risk values, the appropriate generalization of maximization normal form is notationally awkward to express when multiple strata are involved. Fortunately, in contrast to the case of minimization, we will not have occasion to use such a generalization, and so do not give its definition.

and there exists at most one lower “pivot” index i^{lwr} such that $R_{i^{\text{lwr}}} = t^{\text{lwr}}$; and similarly, there exists at most one upper “pivot” index i^{upr} such that $R_{i^{\text{upr}}} = t^{\text{upr}}$

LEMMA A.11. *Consider the single-stratum maximization problem—i.e., maximization in Eq. (13). Suppose that \mathbf{R}^* is a solution. Then, \mathbf{R}^* is in weak normal form. Moreover, either $\frac{1}{n}\|\mathbf{R}^* - \hat{\mathbf{R}}\|_1 = \epsilon$ or R_i^* equals ℓ_i or u_i for all but (at most) one i .*

PROOF. The proof is quite similar to the proof of Lemma A.6. We note that the first-order KKT conditions are almost the same, namely that

$$(19) \quad 2 \cdot \mathbf{R}^* - \lambda \cdot \mathbf{1} - \sum_{i=1}^n \mu_{0,i} \cdot \mathbf{e}_i + \sum_{i=1}^n \mu_{1,i} \cdot \mathbf{e}_i - \sum_{k=1}^{2^n} \nu_k \cdot \mathbf{S}_k = 0$$

for some λ and $\mu_{0,i}, \mu_{1,i}, \nu_k \geq 0$ satisfying complementary slackness. Virtually exactly as before, it follows from this that Eq. (15) holds, i.e., that

$$(20) \quad R_i^* \in \left\{ \hat{R}_i, u_i, \ell_i, \frac{1}{2} \cdot [\lambda - \Delta], \frac{1}{2} \cdot [\lambda + \Delta] \right\}.$$

Again, as in the case of minimization, we can strengthen Eq. (20). In particular, exactly the same argument used to strengthen the weak conditions in Lemma A.6 yields that if $R_{i_0}^* > \hat{R}_{i_0}$ and $R_{i_1}^* < \hat{R}_{i_1}$, then $R_{i_1}^* = u_{i_1}$; and, conversely, if $R_{i_0}^* < \hat{R}_{i_0}$ and $R_{i_1}^* > \hat{R}_{i_1}$, then $R_{i_1}^* = \ell_{i_1}$. The proof is then otherwise the same as in Lemma A.6.

To see that there can be at most one i such that $\ell_i < R_i^* < \hat{R}_i$, suppose that $\ell_{i_0} < R_{i_0}^* < \hat{R}_{i_0}$ and $\ell_{i_1} < R_{i_1}^* < \hat{R}_{i_1}$ for $i_0 \neq i_1$. By Eq. (20), $R_{i_0}^* = R_{i_1}^*$. Choose $\delta > 0$ sufficiently small that $R_{i_0}^* - \delta > \ell_{i_0}$ and $R_{i_1}^* + \delta < \hat{R}_{i_1}$. Then, as in the case of minimization, it follows immediately that $\mathbf{R}^* + \delta(\mathbf{e}_{i_1} - \mathbf{e}_{i_0})$ is feasible, but increases the objective by $2\delta^2$, contrary to the hypothesis that \mathbf{R}^* was a maximizer. In the same way, it follows that there exists at most one i such that $u_i > R_i^* > \hat{R}_i$.

Combining this with our previous necessary condition on \mathbf{R}^* and letting $t^{\text{lwr}} = \frac{1}{2} \cdot [\lambda - \Delta]$ and $t^{\text{upr}} = \frac{1}{2} \cdot [\lambda + \Delta]$ yields that for all i such that $R_i^* < t^{\text{lwr}}$, $R_i^* = \ell_i$; that for all i such that $R_i^* > t^{\text{upr}}$, $R_i^* = u_i$; that there is at most one index such that $\hat{R}_i > R_i^* > t^{\text{lwr}}$; and, finally, that there is at most one i' such that $\hat{R}_{i'} < R_{i'}^* < t^{\text{upr}}$. \square

Lemma A.11 is, by itself, insufficient to form the basis of an effective solution algorithm since it does not determine which indices are the pivots, and it is *a priori* unclear how to search over all possible pairs of pivots efficiently.

Therefore, to construct an effective maximization algorithm, we must transform vectors merely in weak maximization normal form into vectors that are in full-fledged maximization normal form.

To do so, we will have need of the following elementary fact.

LEMMA A.12. *Suppose $a \geq b$ and $c \geq d$. Then*

$$|a - c| + |b - d| \leq |a - d| + |b - c|.$$

PROOF. The proof is not complicated, but is simplest to understand if expressed in geometric terms. Consider the points $\mathbf{p}_0 = (a, b)$ and $\mathbf{p}_1 = (c, d)$. Then both \mathbf{p}_0 and \mathbf{p}_1 lie in the lower half-plane

$$H = \{(x, y) : x \geq y\}.$$

Let $T : (x, y) \mapsto (y, x)$ be the linear transformation given by reflection across the line $y = x$. Then, the claim of the lemma can be reframed as follows: for any $\mathbf{p}_0, \mathbf{p}_1 \in H$,

$$(21) \quad \|\mathbf{p}_0 - \mathbf{p}_1\|_1 \leq \|\mathbf{p}_0 - T(\mathbf{p}_1)\|_1.$$

This setup is shown in Figure 6a. We divide into four cases according to whether $a \geq c$ and $b \geq d$, or, equivalently, depending on in which of the four regions shown in Figure 6b the point \mathbf{p}_1 lies.

Case (A): $a \geq c$ and $b \geq d$. In this case, we have that

$$\begin{aligned} |a - c| + |b - d| &= a - c + b - d \\ &= a - d + b - c \\ &\leq |(a - d) + (b - c)| \\ &\leq |a - d| + |b - c|. \end{aligned}$$

Therefore, Eq. (21) holds.

Case (B): $a \leq c$ and $b \geq d$. In this case,

$$\begin{aligned} |a - c| + |b - d| &= c - a + b - d \\ &\leq c - a + b - d + 2(a - b) \\ &= c - b + a - d \\ &\leq |(c - b) + (a - d)| \\ &\leq |c - b| + |a - d|, \end{aligned}$$

so Eq. (21) holds in this case as well.

Before completing the proof, we note that T has the useful property that

$$(22) \quad \|T(\mathbf{p})\|_1 = \|\mathbf{p}\|_1.$$

Case (C): $a \leq c$ and $b \leq d$. We note that this case is the same as Case (A) with the roles of \mathbf{p}_0 and \mathbf{p}_1 reversed. Consequently, by Eq. (21), we have that

$$\|\mathbf{p}_0 - \mathbf{p}_1\|_1 \leq \|T(\mathbf{p}_0) - \mathbf{p}_1\|_1.$$

Applying Eq. (22) and using the fact that T^2 is the identity, we have that

$$\|T(\mathbf{p}_0) - \mathbf{p}_1\|_1 = \|T(T(\mathbf{p}_0) - \mathbf{p}_1)\|_1 = \|\mathbf{p}_0 - T(\mathbf{p}_1)\|_1,$$

which yields Eq. (21).

Case (D): $a \geq c$ and $b \leq d$. We argue as in the previous case, noting that this case is equivalent to Case (B) with the roles of \mathbf{p}_0 and \mathbf{p}_1 reversed, whence

$$\|\mathbf{p}_0 - \mathbf{p}_1\|_1 \leq \|T(\mathbf{p}_0) - \mathbf{p}_1\|_1 = \|\mathbf{p}_0 - T(\mathbf{p}_1)\|_1.$$

□

A further difference from minimization is that an additional hypothesis is required to connect maximizers to maximization normal form.

DEFINITION A.13 (Sortability). We say that the *constraints are sortable* if, possibly after permuting the indices, the lower bounds, upper bounds, and estimates are all ordered in ascending order, i.e., $\ell_1 \leq \dots \leq \ell_n$, $u_1 \leq \dots \leq u_n$, and $\hat{R}_1 \leq \dots \hat{R}_n$




Fig 6: An illustration of the proof of Lemma A.12. The half-plane H is shown in gray, with the dotted lines representing the coordinate axes. Left: The claim of the lemma is equivalent to the claim that the path joining \mathbf{p}_0 to \mathbf{p}_1 is no longer than the path joining \mathbf{p}_0 to $T(\mathbf{p}_1)$. Right: The division into four cases, depending on where \mathbf{p}_1 lies relative to \mathbf{p}_2 .

Sortability is a mild hypothesis which holds in the typical case, $\ell_i = 0$ and $u_i = 1$ for all i . It also holds if, e.g., ℓ_i and u_i differ from \hat{R}_i by Γ on the log odds scale, as is common in many kinds of sensitivity analysis. If sortability fails, there may be intervals $[\ell_i, u_i]$ that can be non-trivially nested for distinct i , which can cause the pivots to “jump between” different indices a potentially exponential number of times in the search for a maximizer. For notational convenience, we assume throughout that sortable risk vectors are, in fact, already sorted.

The first step to deriving an effective maximization algorithm is to remove “trivial” deviations from maximization normal form.

DEFINITION A.14 (Trivial deviations from normality). We say that a risk vector in weak normal form has a *trivial deviation from normality* if there are $i_0 < i_1$ such that $\ell_{i_0} = \ell_{i_1}$, $\ell_{i_0} < R_{i_0}$, and $R_{i_1} = \ell_{i_1}$; or there are $i'_1 > i'_0$ such that $u_{i'_0} = u_{i'_1}$, $R_{i'_1} < u_{i'_1}$, and $R_{i'_0} = u_{i'_0}$.

An illustration of trivial deviations from normality is shown in Figure 7. As their name suggests, trivial deviations from normality are straightforward to eliminate.

LEMMA A.15. *Suppose the constraints are sortable. If a feasible risk vector \mathbf{R} is in weak normal form with trivial deviations from normality, then there exists a feasible \mathbf{R}' in weak normal form with no trivial deviations from normality such that*

$$\frac{1}{n} \sum_{i=1}^n R_i^2 = \frac{1}{n} \sum_{i=1}^n (R'_i)^2.$$

PROOF. If there were any indices i_0 and i_1 representing a trivial deviation from normality, then we could eliminate them by swapping R_{i_0} and R_{i_1} , i.e., by considering

$$R'_i = R_i \quad (i \neq i_0, i_1), \quad R'_{i_0} = R_{i_1}, \quad R'_{i_1} = R_{i_0}.$$




Fig 7: An illustration of trivial deviations from normal form. As in Figures 1 and 5, the black dots represent estimated risks \hat{R}_i , the error bars represent the allowable range $[\ell_i, u_i]$, and the red dots represent the corresponding values of R_i , with indices increasing from left to right on the x -axis. Weak normal form is violated on the left because $\ell_{i_0} = \ell_{i_1}$ and but $R_{i_0} = \ell_{i_0}$, and similarly on the right.

In particular, it follows immediately that \mathbf{R}' is still in weak normal form and that $\frac{1}{n} \sum_{i=1}^n R_i^2 = \frac{1}{n} \sum_{i=1}^n (R'_i)^2$. The only constraint not satisfied trivially is that $\|\mathbf{R}' - \hat{\mathbf{R}}\|_1 \leq \epsilon$. However, this follows from Lemma A.12, since

$$\|\mathbf{R}' - \hat{\mathbf{R}}\|_1 \leq \|\mathbf{R}^* - \hat{\mathbf{R}}\|_1 \leq \epsilon.$$

In the same way, we can arrange that there are no $i_0 < i_1$ such that $u_{i_0} = u_{i_1}$, $R_{i_0} = u_{i_0}$, and $R_{i_1} < u_{i_1}$. \square

With these preliminaries in hand, we are prepared to solve the maximization problem.

LEMMA A.16. *Suppose the constraints are sortable. Then the single stratum maximization problem in Eq. (13) is maximized by a risk vector in maximization normal form.*

PROOF. We adopt a strategy similar to that used in the “strengthening” portion of the proof of Lemma A.6. Namely, we show that any point satisfying the necessary conditions of Lemma A.11 but without pivots can be “improved” to a feasible point that achieves a greater objective.

Let \mathbf{R}^* be a maximizer. By Lemma A.11, \mathbf{R}^* is in weak maximization normal form; by Lemma A.15, we may assume that \mathbf{R}^* has no trivial deviations from normality.

To prove the theorem, we first observe that it suffices to prove the following two claims:

1. There do not exist $i_0 < i_1$ such that $R_{i_0}^* > \ell_{i_0}$ and $R_{i_1}^* < \hat{R}_{i_1}$,
2. There do not exist $i_0 < i_1$ such that $R_{i_1}^* < u_{i_1}$ and $R_{i_0}^* > \hat{R}_{i_0}$.

First, we see how the theorem follows from the two claims. Define i^{LWR} and i^{UPR} as follows:

$$i^{\text{LWR}} = \min\{i : R_i^* > \ell_i\}, \quad i^{\text{UPR}} = \max\{i : R_i^* < u_i\},$$

then, for all $i < i^{\text{LWR}}$, $R_i^* = \ell_i$, and for all $i > i^{\text{UPR}}$, $R_i^* = u_i$. Moreover, for all i satisfying $i^{\text{LWR}} < i < i^{\text{UPR}}$, by Claim 1, $R_i^* \geq \hat{R}_i$; and, by Claim 2, $R_i^* \leq \hat{R}_i$. Consequently, $R_i^* = \hat{R}_i$ for

these indices. Moreover, if $i^{\text{lwr}} \neq i^{\text{upr}}$, by Claim 2, $R_{i^{\text{lwr}}}^* \leq \hat{R}_{i^{\text{lwr}}}$, i.e., $R_{i^{\text{lwr}}}^* \in [\ell_{i^{\text{lwr}}}, \hat{R}_{i^{\text{lwr}}}]$. Similarly, by Claim 1, we conclude that $R_{i^{\text{upr}}}^* \in [\hat{R}_{i^{\text{upr}}}, u_{i^{\text{upr}}}]$.

Therefore it suffices to prove the two claims. Since the proofs are virtually identical, we focus on Claim 1.

The key observation is the following. Suppose there exist indices $R_{i_0}^*$ and $R_{i_1}^*$ with $R_{i_1}^* \geq R_{i_0}^*$, and suppose that there is some δ such that we can decrease $R_{i_0}^*$ and increase $R_{i_1}^*$ by δ . Then \mathbf{R}^* cannot be maximal, since, letting $\mathbf{R}' = \mathbf{R}^* + \delta(\mathbf{e}_{i_1} - \mathbf{e}_{i_0})$, we have that

$$\begin{aligned}
 \left[\sum_{i=1}^n (R'_i)^2 \right] - \left[\sum_{i=1}^n (R_i^*)^2 \right] &= (R_{i_0}^* - \delta)^2 + (R_{i_1}^* + \delta)^2 - (R_{i_0}^*)^2 - (R_{i_1}^*)^2 \\
 (23) \qquad \qquad \qquad &= 2\delta(\delta + [R_{i_1}^* - R_{i_0}^*]) \\
 &\geq 2\delta^2 \\
 &> 0.
 \end{aligned}$$

To actually prove the claim, we proceed by contradiction. In particular, suppose that $i_0 < i_1$ represented an exception to Claim 1, so that $R_{i_0}^* > \ell_{i_0}$ and $R_{i_1}^* < \hat{R}_{i_1}$.

If $R_{i_0}^* \leq R_{i_1}^*$, then Eq. (23) immediately applies, since increasing $R_{i_1}^*$ decreases the L_1 distance between $\hat{\mathbf{R}}$ and \mathbf{R}^* and decreasing $R_{i_0}^*$, at worst, increases it at the same rate. More formally, for $0 < \delta$ sufficiently small,

$$\begin{aligned}
 |(R_{i_0}^* - \delta) - \hat{R}_{i_0}| + |(R_{i_1}^* + \delta) - \hat{R}_{i_1}| &= |(R_{i_0}^* - \delta) - \hat{R}_{i_0}| + \hat{R}_{i_1}^* - (R_{i_0} + \delta) \\
 &\leq |R_{i_0}^* - \hat{R}_{i_0}| + \delta + \hat{R}_{i_1} - R_{i_0}^* - \delta \\
 &= |R_{i_0}^* - \hat{R}_{i_0}| + |R_{i_1}^* - \hat{R}_{i_1}|,
 \end{aligned}$$

and so \mathbf{R}' is feasible and achieves a greater average sum of squares than \mathbf{R}^* , contrary to the assumed maximality of \mathbf{R}^* .

On the other hand, if $R_{i_0}^* \geq R_{i_1}^*$, then, by Lemma A.12, the vector \mathbf{R}^\dagger given by switching the indices i_0 and i_1 , i.e.,

$$R_i^\dagger = R_i^* \quad (i \neq i_0, i_1), \quad R_{i_0}^\dagger = R_{i_1}^*, \quad R_{i_1}^\dagger = R_{i_0}^*,$$

satisfies $\|\mathbf{R}^\dagger - \hat{\mathbf{R}}\|_1 \leq \epsilon$. Moreover, since

$$u_{i_0} \geq R_{i_0}^* > R_{i_1}^* \geq \ell_{i_1} \geq \ell_{i_0},$$

and

$$u_{i_1} \geq u_{i_0} \geq R_{i_0}^* > R_{i_1}^* \geq \ell_{i_1},$$

we have that for all $i = 1, \dots, n$, $\ell_i \leq R_i^\dagger \leq u_i$. Therefore \mathbf{R}^\dagger is feasible. However, since \mathbf{R}^* has no trivial deviations from normality, it follows that

$$R_{i_0}^\dagger = R_{i_1}^* \geq \ell_{i_1} > \ell_{i_0}, \quad R_{i_1}^\dagger = R_{i_0}^* \leq u_{i_0} < u_{i_1}.$$

Since $R_{i_0}^\dagger < R_{i_1}^\dagger$, we can again apply Eq. (23) to derive a contradiction.

Therefore, no such $i_0 < i_1$ exist, and so \mathbf{R}^* is in maximization normal form. \square

Lemma A.16 allows us to efficiently search over the space of risk vectors to find the maximizer. In particular, all that is necessary is to decrease the risk at the lower pivot and increase it at the upper pivot until the L_1 budget is expended. The linear-time algorithm for doing so is given in Algorithm 7. (As in the case of Algorithm 4, there is a potential

preprocessing step, given in Algorithm 6.) As we did in the minimization algorithm, we capture the maximization algorithm’s value at a variety of L_1 budgets, controlled by the step-size γ —which, we assume, satisfies $k\gamma = \epsilon$ for some $k \in \mathbb{N}$ —rather than just the full budget ϵ . We denote this collection with Σ . The reason for doing this is to allow us to maximize across strata.

A.3.2. Optimizing over all strata. Consider the simplified maximization problem in Eq. (7). Since the objective is separable, the problem would be separable in the different strata, except that the L_1 budget constraint ranges over all of the strata simultaneously. Much as we reduced the base problem to the parameterized problem by introducing additional parameters τ_j , $j = 1, \dots, m$, if the optimal budget allocation to each stratum were known, then we could reduce to the single stratum case, applying Algorithm 7. In particular, we note that the following optimization problem (hereafter **“the separable problem”**) parameterized by the stratum-specific budgets $\epsilon_{1,0}, \dots, \epsilon_{m,0}, \epsilon_{1,1}, \dots, \epsilon_{m,1}$ is separable:

$$(24) \quad \begin{aligned} & \underset{\mathbf{R} \in \mathbb{R}^n}{\text{Optimize}} && \frac{1}{n} \sum_{i=1}^n R_i^2 \\ & \text{s.t.} && \frac{1}{|\mathcal{S}_{j,a}|} \sum_{i \in \mathcal{S}_{j,a}} |R_i - \hat{R}_i| \leq \epsilon_{j,a}, \quad (a = 0, 1; j = 1, \dots, m) \\ & && \frac{1}{|\mathcal{S}_{j,1}|} \cdot \sum_{i \in \mathcal{S}_{j,1}} R_i = \rho_j, \quad (j = 1, \dots, m) \\ & && \frac{1}{|\mathcal{S}_{j,0}|} \cdot \sum_{i \in \mathcal{S}_{j,0}} R_i = \tau_j, \quad (j = 1, \dots, m) \\ & && R_i \leq u_i, \quad (i = 1, \dots, n) \\ & && R_i \geq \ell_i. \quad (i = 1, \dots, n) \end{aligned}$$

In particular, the separable problem can be solved stratum by stratum, since the constraints and objective are both separable at the stratum level.

Reducing solving the simplified problem to solving the separable problem, however, is complex; There is no obvious way to determine the optimal L_1 budget allocation other than by sweeping over all possible budgets. However, it is nevertheless possible to efficiently discover an *approximately* optimal budget allocation using a divide-and-conquer approach. In particular, as shown in Algorithm 7, it is only marginally more difficult computationally to output the maximum value of the objective for the entire budget than to output the maximum value of the objective for a range of budgets, i.e., to output Σ , which gives the maximum value of the objective for the budgets $0, \gamma, 2\gamma, \dots, \epsilon$.

Given $\Sigma^{(0)}$ and $\Sigma^{(1)}$ corresponding to two different strata when the allocated budget is $0, \gamma, 2\gamma, \dots, \epsilon$, we can find the approximate maximum value of the objective when the *total* budget allocated to both strata is $0, \gamma, 2\gamma, \dots, \epsilon$ by sweeping over the two dimensional grid of points $\{0, \gamma, \dots, \epsilon\} \times \{0, \gamma, \dots, \epsilon\}$, and, for $k\gamma$, outputting

$$\max \left(\Sigma_0^{(0)} + \Sigma_{k+1}^{(1)}, \Sigma_1^{(0)} + \Sigma_k^{(1)}, \dots, \Sigma_{k+1}^{(0)} + \Sigma_0^{(1)} \right).$$

For completeness, we explicitly give this maximization routine in Algorithm 8.

LEMMA A.17. *If ℓ , $\hat{\mathbf{R}}$ and \mathbf{u} have been sorted, then the approximate solution to the maximization problem in Eq. (7) given by applying Algorithms 7 and 8 is*

$$O \left(m \cdot \left(\frac{\epsilon}{\gamma} \right)^2 + n \right).$$

We note that δ , as used in Section 4 above, equals $2m\gamma$.

PROOF. We begin by analyzing Algorithm 6. We note that the number of iterations of the while loop on lines 7 through 12 is capped by the length of the input, since i^{upr} is decremented on each input. Likewise, Algorithm 7 is almost linear in the size of the input: in the while loop on lines 14 through 30, either i^{lwr} is incremented, or the subsequent multiple of γ is reached, meaning that the algorithm as a whole is linear in the size of the input and ϵ/γ . Therefore, running Algorithm 4 over all $2m$ strata requires $O(n + \frac{m\epsilon}{\gamma})$ time.

All of the strata can be combined using $2m - 1$ applications of Algorithm 5. Reviewing the for loop on lines 3 through 11 and the while loop on lines 5 through 8, we see that exactly

$$\frac{\frac{\epsilon}{\gamma} \cdot (\frac{\epsilon}{\gamma} + 1)}{2}$$

iterations are performed. Adding this to the previous runtime obtained and simplifying gives the desired expression. \square

A.3.3. *Controlling approximation error.* It is useful to understand how much error is introduced by this approximation. We begin with the following simple lemma.

LEMMA A.18. *Suppose $\mathbf{R}^{(0)}$ and $\mathbf{R}^{(1)}$ are risk vectors. Then*

$$(25) \quad \left| \sum_{i=1}^n (R_i^{(1)})^2 - \sum_{i=1}^n (R_i^{(0)})^2 \right| \leq 2 \cdot \|\mathbf{R}^{(1)} - \mathbf{R}^{(0)}\|.$$

PROOF. The claim, which is a version of Hölder’s inequality, follows straightforwardly from considering the following difference:

$$\begin{aligned} \left| \sum_{i=1}^n (R_i^{(1)})^2 - (R_i^{(0)})^2 \right| &\leq \sum_{i=1}^n |(R_i^{(1)})^2 - (R_i^{(0)})^2| \\ &\leq \sum_{i=1}^n (R_i^{(1)} + R_i^{(0)}) \cdot |R_i^{(1)} - R_i^{(0)}| \\ &\leq \sum_{i=1}^n 2 \cdot |R_i^{(1)} - R_i^{(0)}| \\ &= 2 \cdot \|\mathbf{R}^{(1)} - \mathbf{R}^{(0)}\|_1. \end{aligned}$$

The first inequality follows from the triangle inequality, and the second and third from the fact that $0 \leq R_i^{(0)}, R_i^{(1)} \leq 1$. \square

Using Lemma A.18, we obtain the following characterization of the potential error in our maximization algorithm, again recalling that $\delta = 2m\gamma$, where δ is as in Section 4.

LEMMA A.19. *The difference between the true maximum of the simplified problem—i.e., Eq. (7)—and the quantity obtained from applying Algorithms 4 and 5 is at most $2m\gamma$.*

PROOF. Let \mathbf{R}^* be the true maximizing solution. We will show that Algorithm 8 computes the value of the objective at a “close” point at which we can apply Lemma A.18.

For each stratum $\mathcal{S}_{j,a}$, let $\epsilon_{a,j}$ denote

$$\frac{1}{n} \sum_{i \in \mathcal{S}_{j,a}} |R_i^* - \hat{R}_i|.$$

Then, there exist $\tilde{\epsilon}_{a,j} \in \{0, \gamma, 2\gamma, \dots, \epsilon\}$ such that $\tilde{\epsilon}_{a,j} < \epsilon_{a,j}$ and $\epsilon_{a,j} - \tilde{\epsilon}_{a,j}$ is less than γ .

Let

$$\Delta = \sum_{a=0}^1 \sum_{j=1}^m \epsilon_{a,j} - \tilde{\epsilon}_{a,j} = \epsilon - k\gamma$$

for some k . Since ϵ is a multiple of γ , it follows that $\Delta = k'\gamma$ for some $k' \in \mathbb{N}$, where $k' < 2m$, the number of strata.

Set $\epsilon'_{a,j} = \tilde{\epsilon}_{a,j}$ except for k' arbitrarily chosen strata, where we set $\epsilon'_{a,j} = \tilde{\epsilon}_{a,j} + \gamma$ instead. Then, there exists \mathbf{R} that is in (maximization) normal form such that

$$\frac{1}{n} \sum_{i=1}^n |R_i - \hat{R}_i| = \epsilon, \quad \frac{1}{n} \sum_{\substack{A_i=a \\ C_i=c_j}} |R_i - \hat{R}_i| = \epsilon'_{a,j}.$$

The remainder of the proof follows simply by comparing \mathbf{R} and \mathbf{R}^* stratum by stratum. In particular, consider the restrictions of \mathbf{R} and \mathbf{R}^* to a single stratum $\mathcal{S}_{a,j}$. By Eq. (7), we have that

$$\sum_{i \in \mathcal{S}_{a,j}} R_i = \sum_{i \in \mathcal{S}_{a,j}} R_i^*,$$

and, moreover, by Lemma A.16, both $\mathbf{R}_{a,j}$ and $\mathbf{R}_{a,j}^*$ can be assumed to be in maximization normal form. Assume, without loss of generality, that $\epsilon'_{a,j} \leq \epsilon_{a,j}$. (The proof is virtually identical if the inequality is reversed.)

Let i_0 and i_1 denote the pivots of \mathbf{R} , and i_0^* and i_1^* of \mathbf{R}^* , so that $i_0 \leq i_0^* \leq i_1^* \leq i_1$. (To avoid unnecessary notational complication, we write as if all n individuals belong to the stratum $\mathcal{S}_{a,j}$ and the constraints were sorted.) Then, we note that

$$(26) \quad R_i \geq R_i^* \quad (i_0 \leq i \leq i_0^*), \quad R_i \leq R_i^* \quad (i_1^* \leq i \leq i_1), \quad R_i = R_i^* \quad (\text{otherwise}).$$

It follows that

$$(27) \quad \sum_{i=1}^n (R_i^*)^2 - R_i^2 = \left[\sum_{i=i_1^*}^{i_1} (R_i^*)^2 - R_i^2 \right] - \left[\sum_{i=i_0}^{i_0^*} R_i^2 - (R_i^*)^2 \right].$$

Both terms in Eq. (27) are positive by Eq. (26). Moreover, by Lemma A.18, both terms are less than or equal to

$$2 \cdot \left[\sum_{i=i_1}^{i_1^*} R_i^* - R_i \right] = 2 \cdot \left[\sum_{i=i_0}^{i_0^*} R_i - R_i^* \right] = 2n \cdot \frac{\epsilon_{a,j} - \epsilon'_{a,j}}{2} \leq n\gamma.$$

In particular, their difference must also be less than this quantity, and so, summing across strata, we have that

$$\sum_{i=1}^n (R_i^*)^2 - R_i^2 \leq 2nm\gamma.$$

Dividing through by n gives the result. \square

Algorithm 2 Piecewise quadratic function evaluation

Input: The collections Δ and \mathbf{K} , as well as the L_1 “budget” ϵ .

Output: The value of $\Sigma(\Delta^*)$, where $\epsilon(\Delta^*) = \epsilon$.

```

1: Set  $N \leftarrow \text{length}(\delta)$ 
2: Initialize  $i \leftarrow 1$ 
3: Initialize  $\varepsilon \leftarrow 0$ 
4: Initialize  $t \leftarrow 0$ 
5: Initialize  $\Sigma \leftarrow 0$ 
6: while  $\varepsilon + K_i(\Delta_{i+1} - \Delta_i) < \epsilon$  and  $i < N$  do
7:   Set  $\varepsilon \leftarrow \varepsilon + K_i(\Delta_{i+1} - \Delta_i)$ 
8:   Set  $\Sigma \leftarrow \Sigma - K_i \cdot \Delta_i \cdot (\Delta_{i+1} - \Delta_i) + \frac{K}{2} \cdot (\Delta_{i+1} - \Delta_i)^2$ 
9:   Set  $i \leftarrow i + 1$ 
10: end while
11: if  $K_i = 0$  then
12:   return  $\Sigma$       {  $K_i = 0$  if and only if the budget was exactly exhausted on the last iteration of the loop }
13: else if  $i = N$  then
14:   return  $\Sigma$       {  $i = N$  if and only if we have made it to the end of  $\delta$ , i.e., to  $\delta = 0$  }
15: else
16:   Set  $t \leftarrow (\epsilon - \varepsilon)/K_i$ 
17:   Set
18:   Set  $\Sigma \leftarrow \Sigma - K_i \cdot \Delta_i \cdot t + \frac{K}{2} \cdot t^2$ 
19:   return  $\Sigma$ 
20: end if

```

Algorithm 3 Minimization algorithm (sum adjustment)

Input: The bounds ℓ and \mathbf{u} , the estimates $\hat{\mathbf{R}}$, and the sum μ .

Output: A risk vector \mathbf{R} of the form in Lemma A.6 minimizing $\|\mathbf{R} - \hat{\mathbf{R}}\|_1$ and satisfying $\sum_{i=1}^n R_i = \mu$, along with the ending $k^{\text{lwr}}, k^{\text{upr}}, i^{\text{lwr}}, i^{\text{upr}}, t^{\text{upr}}$, and t^{lwr} .

```

1: Set pts to be the concatenation of  $\ell$ ,  $\hat{\mathbf{R}}$ , and  $\mathbf{u}$  in ascending order      {Points at which rates can change}
2: Set  $\mathbf{R} \leftarrow \hat{\mathbf{R}}$                                                                                                {Risk vector}
3: Set  $n \leftarrow \text{length}(\text{pts})$ 
4: Set  $i^{\text{lwr}} = 1, i^{\text{upr}} = n$                                                                                    {Indices of next values at which rates might change}
5: Set  $t^{\text{lwr}} \leftarrow \text{pts}[i^{\text{lwr}}], t^{\text{upr}} \leftarrow \text{pts}[i^{\text{upr}}]$                                          {Thresholds}
6: Set  $k^{\text{lwr}} \leftarrow 0, k^{\text{upr}} \leftarrow 0$                                                                                    {Number of active indices}
7:  $D \leftarrow \mu - \sum_{i=1}^n R_i$                                                                                    {Difference between required sum and actual sum}
8: if  $D > 0$  then
9:    $t^{\text{next}} \leftarrow \text{pts}[i^{\text{lwr}} + 1]$ 
10:   $d \leftarrow 0$                                                                                                        {Change in sum from moving from  $t^{\text{lwr}}$  to  $t^{\text{next}}$ }
11:  while  $D > d$  and  $i^{\text{lwr}} < n$  do
12:     $D \leftarrow D - d$ 
13:     $i^{\text{lwr}} \leftarrow i^{\text{lwr}} + 1$ 
14:     $t^{\text{lwr}} \leftarrow t^{\text{next}}$ 
15:     $t^{\text{next}} \leftarrow \text{pts}[i^{\text{lwr}} + 1]$ 
16:    if  $t^{\text{next}}$  corresponds to an element of  $\hat{\mathbf{R}}$  then
17:       $k^{\text{lwr}} \leftarrow k^{\text{lwr}} + 1$ 
18:    else if  $t^{\text{next}}$  corresponds to an element of  $\mathbf{u}$  then
19:       $k^{\text{lwr}} \leftarrow k^{\text{lwr}} - 1$ 
20:    end if
21:     $d \leftarrow k^{\text{lwr}}(t^{\text{next}} - t^{\text{lwr}})$ 
22:  end while
23:  if  $i^{\text{lwr}} < n$  then
24:     $t^{\text{lwr}} \leftarrow t^{\text{lwr}} + \frac{D}{k^{\text{lwr}}}$ 
25:  end if
26: else if  $D < 0$  then
27:   {Similar steps to the  $D > 0$  case but adapted for the upper threshold}
28: end if
29:
30: return  $\mathbf{R}, k^{\text{lwr}}, k^{\text{upr}}, i^{\text{lwr}}, i^{\text{upr}}, t^{\text{upr}}$ , and  $t^{\text{lwr}}$ 

```

Algorithm 4 Minimization algorithm (single stratum)

Input: The bounds ℓ and \mathbf{u} , the estimates $\hat{\mathbf{R}}$, and the sum μ .

Output: The collections Δ and \mathbf{K} .

```

1: Initialize  $\text{pts}$ ,  $\mathbf{R}$ ,  $t^{\text{lwr}}$ ,  $t^{\text{upr}}$ ,  $k^{\text{lwr}}$ ,  $k^{\text{upr}}$ ,  $i^{\text{lwr}}$ , and  $i^{\text{upr}}$  as in Algorithm 3
2: Set  $n \leftarrow \text{length}(\text{pts})$ 
3: Set  $\epsilon \leftarrow 0$  { Amount of budget expended so far }
4: Set  $\Delta \leftarrow t^{\text{upr}} - t^{\text{lwr}}$  { Gap between thresholds }
5: Using Algorithm 3, update  $\mathbf{R}$ ,  $k^{\text{lwr}}$ ,  $k^{\text{upr}}$ ,  $i^{\text{lwr}}$ ,  $i^{\text{upr}}$ ,  $t^{\text{upr}}$ , and  $t^{\text{lwr}}$  { Ensure that sum is correct }
6: Set  $\epsilon \leftarrow \epsilon - \|\mathbf{R} - \hat{\mathbf{R}}\|_1$ , returning that the problem is infeasible if the result is negative
7: if  $n = 1$  then
8:   return  $\Delta = (\Delta_0)$  and  $\mathbf{K} = 0$ 
9: end if
10:  $D^{\text{lwr}} \leftarrow k^{\text{lwr}}(\text{pts}[i^{\text{lwr}} + 1] - t^{\text{lwr}})$  { Cost of moving from lower threshold to next change point }
11:  $D^{\text{upr}} \leftarrow k^{\text{upr}}(t^{\text{upr}} - \text{pts}[i^{\text{upr}} - 1])$  { Cost of moving from upper threshold and next change point }
12:  $D \leftarrow \min(D^{\text{lwr}}, D^{\text{upr}})$  { Smaller of the costs }
13:  $K \leftarrow 0$ 
14: while  $i^{\text{upr}} - i^{\text{lwr}} > 1$  do
15:   Append  $\Delta$  to  $\Delta$  and  $K$  to  $\mathbf{K}$ 
16:   if  $D^{\text{lwr}} = D$  then
17:      $i^{\text{lwr}} \leftarrow i^{\text{lwr}} + 1$  { Increment lower active index }
18:      $t^{\text{lwr}} \leftarrow \text{dps}[i^{\text{lwr}}]$  { Update the lower threshold }
19:     if  $t^{\text{lwr}}$  came from  $\hat{\mathbf{R}}$  then
20:        $k^{\text{lwr}} \leftarrow k^{\text{lwr}} + 1$  { Increment  $k^{\text{lwr}}$  if an index is activated }
21:     else if  $t^{\text{lwr}}$  came from  $\mathbf{u}$  then
22:        $k^{\text{lwr}} \leftarrow k^{\text{lwr}} - 1$  { Decrement  $k^{\text{lwr}}$  if an index is deactivated }
23:     end if
24:      $D^{\text{upr}} \leftarrow D^{\text{upr}} - D^{\text{lwr}}$  { Calculate new upper active gap }
25:      $t^{\text{upr}} \leftarrow t^{\text{upr}} - \frac{D^{\text{lwr}}}{k^{\text{upr}}}$  { Calculate new upper threshold }
26:      $D^{\text{lwr}} \leftarrow k^{\text{lwr}}(\text{dps}[i^{\text{lwr}} + 1] - t^{\text{lwr}})$  { Calculate new lower active gap }
27:   else
28:     { Similar steps as in previous branch, adapted to the upper threshold }
29:   end if
30:    $D \leftarrow \min(D^{\text{lwr}}, D^{\text{upr}})$ 
31:    $\Delta \leftarrow t^{\text{upr}} - t^{\text{lwr}}$ 
32:    $K \leftarrow \frac{k^{\text{lwr}} \cdot k^{\text{upr}}}{k^{\text{lwr}} + k^{\text{upr}}}$ 
33: end while
34: Append 0 to  $\Delta$  and  $\mathbf{K}$ 
35: return  $\Delta, \mathbf{K}$ 

```

Algorithm 5 Minimization algorithm (combining strata)**Input:** The collections $\Delta^{(0)}$, $\mathbf{K}^{(0)}$, $\Delta^{(1)}$, and $\mathbf{K}^{(1)}$.**Output:** A pair $(\Delta^{(2)}, \mathbf{K}^{(2)})$ representing the sum of functions $\epsilon^{(0)}(\Delta) + \epsilon^{(1)}(\Delta)$ and $\Sigma^{(0)}(\Delta) + \Sigma^{(1)}(\Delta)$.

```

1: Initialize  $i_0 \leftarrow 0, i_1 \leftarrow 1$  { Pointers to smallest indices not yet combined }
2: Initialize  $K_0 \leftarrow 0, K_1 \leftarrow 0, \Delta_0 \leftarrow 1, \Delta_1 \leftarrow 1$  { Current values of the parameters }
3: while  $i_0 \leq \text{length}(\Delta_0)$  and  $i_1 \leq \text{length}(\Delta_1)$  do
4:   if  $\Delta_{i_0}^{(0)} > \Delta_{i_1}^{(1)}$  then
5:     Set  $i_0 \leftarrow i_0 + 1, K_0 \leftarrow K_{i_0}^{(0)}$ , and  $\Delta_0 \leftarrow \Delta_{i_0}^{(0)}$ 
6:     Append  $\Delta_0$  to  $\Delta^{(2)}$  and  $K_0 + K_1$  to  $\mathbf{K}^{(2)}$ 
7:   else
8:     Set  $i_1 \leftarrow i_1 + 1, K_1 \leftarrow K_{i_1}^{(1)}$ , and  $\Delta_1 \leftarrow \Delta_{i_1}^{(1)}$ 
9:     Append  $\Delta_1$  to  $\Delta^{(2)}$  and  $K_0 + K_1$  to  $\mathbf{K}^{(2)}$ 
10:  end if
11: end while
12: return  $\Delta^{(2)}$  and  $\mathbf{K}^{(2)}$ 

```

Algorithm 6 Maximization algorithm (sum adjustment)**Input:** The bounds ℓ and \mathbf{u} , the estimates $\hat{\mathbf{R}}$, and the sum μ .**Output:** A risk vector \mathbf{R} of the form in Lemma A.11 minimizing $\|\mathbf{R} - \hat{\mathbf{R}}\|_1$ and satisfying $\sum_{i=1}^n R_i = \mu$, along with $i^{\text{lwr}}, i^{\text{upr}}$.

```

1: Set  $\mathbf{R} \leftarrow \hat{\mathbf{R}}$  { Risk vector }
2: Set  $n \leftarrow \text{length}(\mathbf{R})$ 
3: Set  $i^{\text{lwr}} = 1, i^{\text{upr}} = n$  { Pivots }
4:  $D \leftarrow \mu - \sum_{i=1}^n R_i$  { Difference between required sum and actual sum }
5: if  $D > 0$  then
6:    $d \leftarrow u_{i^{\text{upr}}} - R_{i^{\text{upr}}}$ 
7:   while  $D > d$  and  $i^{\text{upr}} > 0$  do
8:      $R_{i^{\text{upr}}} \leftarrow u_{i^{\text{upr}}}$ 
9:      $D \leftarrow D - d$ 
10:     $i^{\text{upr}} \leftarrow i^{\text{upr}} - 1$ 
11:     $d \leftarrow u_{i^{\text{upr}}} - R_{i^{\text{upr}}}$ 
12:  end while
13:   $R_{i^{\text{upr}}} \leftarrow R_{i^{\text{upr}}} + D$ 
14: else if  $D < 0$  then
15:   { Similar steps to the  $D > 0$  case but adapted for the lower pivot }
16: end if
17:
18: return  $\mathbf{R}, i^{\text{lwr}}$ , and  $i^{\text{upr}}$ .

```

Algorithm 7 Maximization algorithm (single stratum)

Input: The bounds ℓ and \mathbf{u} , the estimates $\hat{\mathbf{R}}$, the sum μ , and the step size γ .

Output: The collection Σ .

```

1: Initialize  $\mathbf{R}$ ,  $i^{\text{lwr}}$ , and  $i^{\text{upr}}$  as in Algorithm 6
2: Set  $n \leftarrow \text{length}(R)$ 
3: Set  $\varepsilon \leftarrow 0$  { Amount of budget expended so far }
4: Using Algorithm 6, update  $\mathbf{R}$ ,  $i^{\text{lwr}}$ , and  $i^{\text{upr}}$ ,  $t^{\text{upr}}$  { Ensure that sum is correct }
5: Set  $\varepsilon \leftarrow \|\mathbf{R} - \hat{\mathbf{R}}\|_1$  and  $\varepsilon \leftarrow \varepsilon - \varepsilon$ , returning that the problem is infeasible if  $\varepsilon > \epsilon$ 
6: while  $\varepsilon \geq \gamma$  do
7:   Append  $\infty$  to  $\Sigma$  { Sentinel value to indicate infeasibility }
8:    $\varepsilon \leftarrow \varepsilon - \gamma$ 
9: end while
10:  $\Delta^{\text{lwr}} \leftarrow R_{i^{\text{lwr}}} - \ell_{i^{\text{lwr}}}$ 
11:  $\Delta^{\text{upr}} \leftarrow u_{i^{\text{upr}}} - R_{i^{\text{upr}}}$ 
12:  $\Delta \leftarrow \min(\Delta^{\text{lwr}}, \Delta^{\text{upr}})$ 
13:  $\Sigma \leftarrow \sum_{i=1}^n R_i^2$  { Sum of squares }
14: while  $i^{\text{upr}} < i^{\text{lwr}}$  do
15:   if  $\Delta \geq \gamma - \varepsilon$  then
16:      $\Sigma \leftarrow \Sigma + 2\gamma \cdot (\gamma + [R_{i^{\text{upr}}} - R_{i^{\text{lwr}}}]$  { If next step would exceed step size, record change }
17:     Append  $\Sigma$  to  $\Sigma$ 
18:      $R_{i^{\text{lwr}}} \leftarrow R_{i^{\text{lwr}}} + \gamma$ ,  $R_{i^{\text{upr}}} \leftarrow R_{i^{\text{upr}}} + \gamma$  { Update risk vector at pivots }
19:      $\Delta \leftarrow \Delta - \gamma$ ,  $\Delta^{\text{min}} \leftarrow \Delta^{\text{min}} - \gamma$ ,  $\Delta^{\text{max}} \leftarrow \Delta^{\text{max}} - \gamma$  { Update gaps }
20:      $\varepsilon \leftarrow 0$  { Only adjust step size on first iteration }
21:   else
22:     if  $\Delta^{\text{lwr}} = \Delta$  then
23:        $\Sigma \leftarrow \Sigma + 2\Delta \cdot (\Delta + [R_{i^{\text{upr}}} - R_{i^{\text{lwr}}}]$ 
24:        $R_{i^{\text{lwr}}} \leftarrow \ell_{i^{\text{lwr}}}$ ,  $R_{i^{\text{upr}}} \leftarrow R_{i^{\text{upr}}} + \Delta$ 
25:        $i^{\text{lwr}} \leftarrow i^{\text{lwr}} + 1$  { Increment lower active index }
26:     else
27:       { Similar steps as in previous branch, adapted to the upper pivot }
28:     end if
29:   end if
30: end while
31: Append  $\Sigma$  to  $\Sigma$  until it has the appropriate length
32: return  $\Sigma$ 

```

Algorithm 8 Maximization algorithm (combining strata)

Input: The collections $\Sigma^{(0)}$ and $\Sigma^{(1)}$.

Output: A collection $\Sigma^{(2)}$ representing the (approximate) maximum objective across both collections.

```

1: Initialize  $i_0 \leftarrow 0$ ,  $i_1 \leftarrow 1$  { Pointers to indices currently being combined }
2: Initialize  $M \leftarrow -\infty$  { Current maximum }
3: for  $i = 0, \dots, \varepsilon/\gamma$  do
4:    $i_0 \leftarrow 0$ ,  $i_1 \leftarrow i$ 
5:   while  $i_0 < i$  do
6:      $M \leftarrow \Sigma_{i_0}^{(0)} + \Sigma_{i_1}^{(1)}$  { Find the maximum across gridpoints whose total budget is  $i\gamma$  }
7:      $i_0 \leftarrow i_0 + 1$ ,  $i_1 \leftarrow i_1 - 1$ 
8:   end while
9:   Append  $M$  to  $\Sigma^{(2)}$ 
10:   $M \leftarrow -\infty$ 
11: end for
12: return  $\Sigma^{(2)}$ 

```

APPENDIX B: ESTIMATION OF THE NON-PARAMETRIC ESTIMAND BY LINEAR REGRESSION

To assess the quality of risk-adjusted regression as an estimator of the non-parametric estimand in (2), we construct synthetic datasets based on the NYPD data. On this synthetic dataset, we compare the true value of the risk adjusted disparities—as defined in (2)—to estimates from a risk-adjusted regression.

To construct the synthetic datasets, we first use the real NYPD data to estimate the probability of frisk, $\Pr(A = 1 \mid \tilde{X})$, conditional on the same observed pre-frisk covariates \tilde{X} used to estimate risk in the main analysis, excluding race, suspected crime, and precinct.²⁰ We similarly estimate the probability a weapon is recovered on a frisked individual, $\Pr(W = 1 \mid \tilde{X}, A = 1)$, again conditional on the same pre-frisk covariates. In both cases, we use logistic regression to estimate the probabilities. The frisk model is trained on all data from 2008 and 2009, and the weapon-recovery model is trained on the subset of stops of the same in which an individual was frisked.

With the resulting covariate estimates $\hat{\beta}_{\text{weapon}}$ and $\hat{\beta}_{\text{frisk}}$, for the k -th iteration of the $m = 100$ simulation instances, we generate a population of $n = 1,000,000$ observations as follows:

1. Sample noise terms $\epsilon_{\text{weapon}} \sim \mathcal{N}(0, \frac{1}{50} \cdot I)$ and $\epsilon_{\text{frisk}} \sim \mathcal{N}(0, \frac{1}{50} \cdot I)$, resulting in new weapon possession and frisk model coefficients $\beta_{k,\text{weapon}} = \hat{\beta}_{\text{weapon}} + \epsilon_{\text{weapon}}$ and $\beta_{k,\text{frisk}} = \hat{\beta}_{\text{frisk}} + \epsilon_{\text{frisk}}$.
2. Sample, with replacement, n covariate vector and race pairs $(\tilde{X}_i, C_i)_{i=1,\dots,n}$ from the original NYPD dataset covering the years 2010 and 2011.
3. For all $i = 1, \dots, n$:
 - a) Define the probability the i -th individual is frisked: $p_i = \text{logit}^{-1}(\beta_{k,\text{frisk}}^\top \tilde{X}_i)$. Then sample $A_i \sim \text{Bernoulli}(p_i)$.
 - b) Define the probability the i -th individual has a weapon: $R_i = \text{logit}^{-1}(\beta_{k,\text{weapon}}^\top \tilde{X}_i)$. Then, sample $W_i \sim \text{Bernoulli}(R_i)$.

The above procedure produces a set of tuples, $\Omega^j = \{(\tilde{X}_i, C_i, A_i, R_i, W_i)\}_{i=1,\dots,n}$, each of which is a synthetic population. On this population, we compute the ground-truth risk-adjusted disparities via (2).²¹ These disparities are typically non-zero; note, though, that there is no disparate *treatment* in this example since, by construction, the probability p_i of frisking an individual does not explicitly depend on race. The disparate impact we measure in this scenario thus arises from decisions that are simply not appropriately tailored to risk, as in *Griggs*.

We next evaluate the ability of our own risk-adjusted regression to recover the ground-truth disparate impact. To do so, we first divide Ω^j into two random sets of equal size, Ω_1^j and Ω_2^j . On Ω_1^j , we use logistic regression to estimate the probability that frisked individuals are found to have a weapon,²² based on the pre-frisk covariates \tilde{X} . This model is trained on the subset of stops in Ω_1^j in which a frisk was carried out. We then use this fitted model to

²⁰We exclude race to demonstrate that disparate impact can occur in the absence of disparate treatment. We exclude suspected crime and precinct because the large number of strata for these covariates leads to sampling difficulties.

²¹The definition in (2) involves conditioning on risk, which is continuous. This is problematic on a finite population, as there may be only one individual with any given risk. Binning risks avoids this issue, and, in this case, we get nearly identical values for any appropriately small bin size.

²²To avoid risk estimates not being well defined when certain feature levels appear in Ω_2^j but not Ω_1^j , the models are fit using penalized maximum likelihood with the `glmnet` package and an elastic net penalty of $\alpha = \frac{1}{10}$.




Fig 8: *Estimates of disparate impact using risk-adjusted regressions on $m = 100$ synthetic datasets, each consisting of $n = 1,000,000$ simulated stops. Despite model misspecification, the risk-adjusted regressions coefficient is a reasonably robust estimator of the true disparity.*

predict *ex ante* risk \hat{R}_i for every stop in Ω_2^j . Lastly, we fit a linear risk-adjusted regression on Ω_2^j to estimate the Black-white and Hispanic-white disparities.

We note that to simulate real-world settings, we have ensured that, by construction, the linear probability model used to estimate disparities is misspecified. Nevertheless, it is a reasonably robust estimator of the true disparity, as can be seen in Figure 8. For both Black and Hispanic pedestrians in the synthetic data, our risk-adjusted estimates of disparate impact are very close to the ground-truth estimands, with model misspecification leading to slight overestimates in some instances, and slight underestimates in others.

APPENDIX C: ACCOUNTING FOR FACTORS BEYOND RISK

Our analysis in the main text implicitly assumed that risk of weapon possession is the only legitimate consideration for carrying out a frisk—an assumption motivated by the fact that a frisk is legally allowed only to ensure officer safety. Officers, however are not required to frisk every individual who they are legally allowed to frisk (i.e., the law only sets a lower bar), leaving open the possibility that risk is not the only factor officers consider when making frisk decisions.²³ For example, in theory, resource-constrained officers might choose to frisk only the riskiest individuals they stop, effectively setting the bar to frisk individuals higher than the law demands. If resources differ across neighborhoods, which in turn correlate with racial composition, then the risk-adjusted disparities we see may accordingly have a policy-relevant justification.

In Figure 9, we aim to account for this possibility by repeating our analysis on three subsets of stops stratified by geography. Specifically, we split the 76 police precincts in our data into three bins based on the racial composition of stopped individuals. Across strata, we find qualitatively similar estimates of disparate impact, corroborating our main results. These within-strata estimates are also reasonably robust to mismeasurement of risk, as indicated by the black bands, although less so than our main result—a 0.7 p.p. average absolute difference in risk would potentially be sufficient to change the sign of the disparity in the most non-white neighborhoods, or to mask a disparity roughly ten times larger than our estimate.

That factors beyond risk may justifiably inform frisk decisions can be viewed as a form of omitted-variable bias, but one distinct in kind from that typically considered in studies of discrimination. Our disparate impact analysis is predicated on the understanding that risk, appropriately defined and estimated, captures nearly all policy-relevant considerations. Although it can be important in practice to accommodate exceptions when there is a clearly articulated rationale—as we have done above—care must be taken not to blindly adjust for every available factor, lest one re-introduce included-variable bias.

²³In *Floyd*, the court found that officers at times frisked individuals even in the absence of safety concerns, in violation of the Fourth Amendment of the U.S. Constitution (Goel, Rao and Shroff, 2016a).




Fig 9: Estimates of disparate impact across precincts stratified by racial composition of stops, with the black bands showing the range of possible estimates if the odds ratio of estimated risk and true risk differ by at most $\epsilon = 0.7$ p.p.




Fig 10: Distributions of the estimated ex ante risk of probability of carrying a weapon for stopped pedestrians, on a log scale. The vertical lines indicate each group's average risk (i.e., the estimated rate at which stopped members of the group carry weapons): 2.7% for white pedestrians, 1.5% for Black pedestrians, and 1.7% for Hispanic pedestrians.


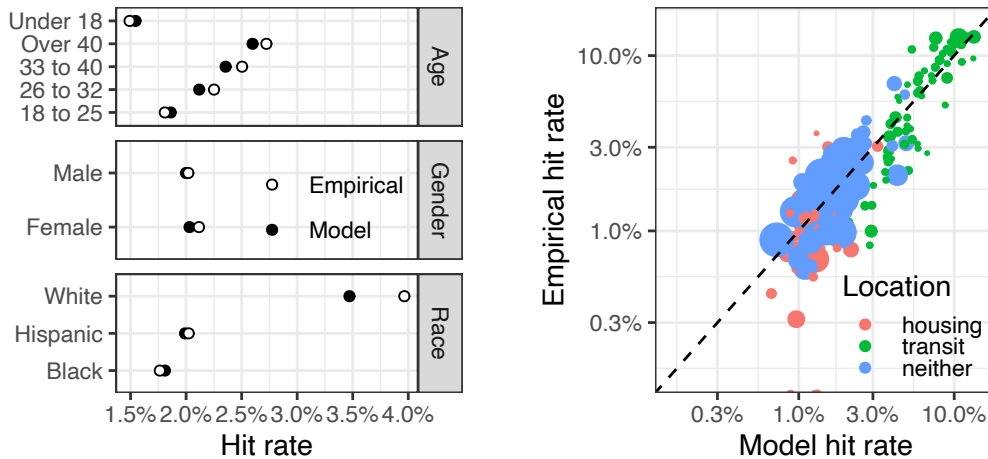


Fig 11: *Distributions of the estimated ex ante risk of probability of carrying a weapon for stopped pedestrians, on a log scale among frisked individuals, faceted by whether or not a weapon was found.*



(a) Demographic groups

(b) Locations

Fig 12: Comparison of model-predicted versus empirical weapon recovery rate (“hit rate”). Figure 12a shows that the model-predicted hit rates are close to their empirical counterparts, conditional on values of age, race, and gender. In Figure 12b, stops are binned by precinct and stop location. Points are plotted for each bin with more than 100 stops, sized by the number of stops, with colors representing the stop location type: transit, housing, or other. The plotted points are near the diagonal, suggesting that the outcome model is well-calibrated and predicts well over the full range of hit rates and frisk rates, respectively. The model itself achieves an AUC of 81% on the second half of the data.




Fig 13: Frisk rates vs. risk (where risk has been estimated without race as a covariate), as estimated via logistic regression curves fit separately for each race group. Across risk levels, stopped Black and Hispanic pedestrians are frisked substantially more frequently than comparably risk white individuals, indicative of disparate impact.




Fig 14: Racial gaps in frisk rates adjusting for different sets of covariates, where the y-axis shows the percentage point difference relative to stopped white individuals. The left panel shows the raw disparities in frisk rates. As a measure of discrimination, raw disparities suffer from omitted-variable bias: there may, in theory, be legitimate reasons why Black and Hispanic pedestrians are more likely to be frisked. The middle panel shows the estimated race effects in a kitchen-sink regression, adjusting for all pre-frisk covariates. These estimates suffer from included-variable bias because they adjust for features that are correlated with race but unrelated to risk. The right panel shows the results of our risk-adjusted regression—where risk has been estimated without race as a covariate—adjusting exclusively for estimated risk of weapon possession. In all cases, estimated standard errors are less than 0.2 percentage points, and so are not visible in the plot.




Fig 15: Sensitivity of the risk-adjusted disparities in frisk decisions to mismeasurement of risk, where risk has been estimated without race as a covariate. The blue bands bound our estimates of disparate impact as a function of the average absolute difference between the true and estimated risks ϵ , relative to the base rate (1.7%). The dotted line at 44% ($\epsilon = 0.7$ p.p.) corresponds to a simulated situation with severe confounding. The grey bands represent 95% percentile bootstrapped confidence intervals ($N = 1000$; Zhao, Small and Bhattacharya, 2019). The step size for the grid search over group-level total risks was 1 p.p. for all groups, and $\gamma = 0.01$ p.p. was the approximation parameter for the maximization routine.




Fig 16: Estimates of disparate impact across precincts stratified by racial composition of stops, with the black bands showing the range of possible estimates if the odds ratio of estimated risk and true risk differ by at most $\epsilon = 0.7$ p.p. Here, risk is estimated without race as a covariate.




Fig 17: *Distributions of the estimated ex ante risk of carrying a weapon for stopped pedestrians, on a log scale and estimated without race as a covariate. The vertical lines indicate each group's average risk (i.e., the estimated rate at which stopped members of the group carry weapons): 2.7% for white pedestrians, 1.5% for Black pedestrians, and 1.7% for Hispanic pedestrians.*

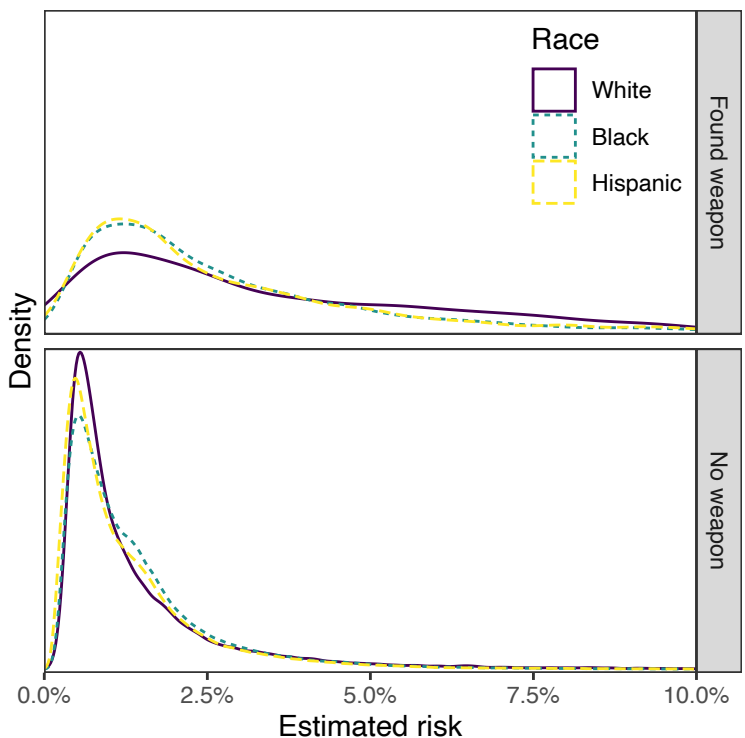
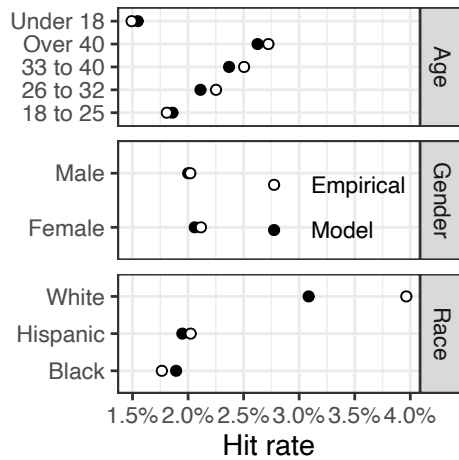
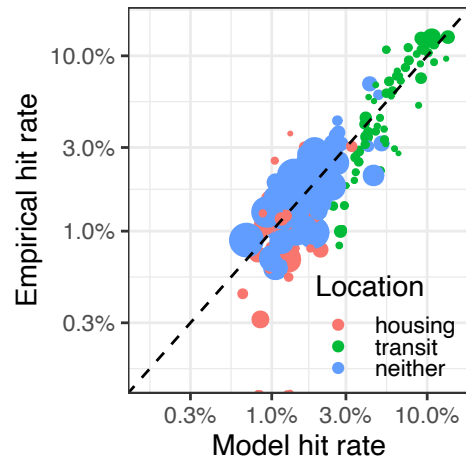


Fig 18: Distributions of the estimated ex ante risk of probability of carrying a weapon—estimated without race as a covariate—for stopped pedestrians, on a log scale among frisked individuals, faceted by whether or not a weapon was found.



(a) Demographic groups



(b) Locations

Fig 19: Comparison of model-predicted versus empirical weapon recovery rate (“hit rate”), where the risk model does not include race as a covariate. Figure 19a shows that the model-predicted hit rates are close to their empirical counterparts, conditional on values of age, race, and gender. In Figure 19b, stops are binned by precinct and stop location. Points are plotted for each bin with more than 100 stops, sized by the number of stops, with colors representing the stop location type: transit, housing, or other. The plotted points are near the diagonal, suggesting that the outcome model is well-calibrated and predicts well over the full range of hit rates and frisk rates, respectively. The model itself achieves an AUC of 81% on the second half of the data.

Acknowledgments. We thank Josh Grossman and Julian Nyarko for helpful conversations.

REFERENCES

- ANGRIST, J. D. and PISCHKE, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- ARNOLD, D., DOBBIE, W. S. and HULL, P. (2020). Measuring racial discrimination in bail decisions Technical Report, National Bureau of Economic Research.
- ARNOLD, D., DOBBIE, W. and HULL, P. (2021). Measuring racial discrimination in algorithms. In *AEA Papers and Proceedings* **111** 49–54.
- ARROW, K. (1973). The theory of discrimination. *Discrimination in labor markets* **3** 3–33.
- AYRES, I. (2002). Outcome tests of racial disparities in police practices. *Justice Research and Policy* **4** 131–142.
- AYRES, I. (2005). Three tests for measuring unjustified disparate impacts in organ transplantation: The problem of “included variable” bias. *Perspectives in Biology and Medicine* **48** 68–S87.
- AYRES, I. (2010). Testing for Discrimination and the Problem of “Included Variable Bias”. *Working paper*. Available at <http://islandia.law.yale.edu/ayers/ayresincludedvariablebias.pdf>.
- BALSA, A. I., MCGUIRE, T. G. and MEREDITH, L. S. (2005). Testing for statistical discrimination in health care. *Health Services Research* **40** 227–252.
- BECKER, G. S. (1993). Nobel lecture: The economic way of looking at behavior. *Journal of Political Economy* **101** 385–409.
- BERKOVEC, J. A., CANNER, G. B., GABRIEL, S. A. and HANNAN, T. H. (2018). Mortgage discrimination and FHA loan performance. In *Mortgage Lending, Racial Discrimination, and Federal Policy* 289–305. Routledge.
- BOHREN, J. A., HULL, P. and IMAS, A. (2022). Systemic discrimination: Theory and measurement Technical Report, National Bureau of Economic Research.
- BOYD, S. P. and VANDENBERGHE, L. (2004). *Convex optimization*. Cambridge university press.
- CHAUDHURI, S. and SETHI, R. (2008). Statistical Discrimination with Peer Effects: Can Integration Eliminate Negative Stereotypes? *The Review of Economic Studies* **75** 579–596. <https://doi.org/10.1111/j.1467-937X.2008.00468.x>
- CHOULDECHOVA, A., BENAVIDES-PRADO, D., FIALKO, O. and VAITHIANATHAN, R. (2018). A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency* 134–148.
- CINELLI, C. and HAZLETT, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society Series B-Statistical Methodology* **82** 39–67.
- CORBETT-DAVIES, S., GAEBLER, J., NILFOROSHAN, H., SHROFF, R. and GOEL, S. (2023). The Measure and Mismeasure of Fairness. *arXiv preprint arXiv:1808.00023*.
- COSTON, A., MISHLER, A., KENNEDY, E. H. and CHOULDECHOVA, A. (2020). Counterfactual risk assessments, evaluation, and fairness. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* 582–593.
- DOBBIE, W., GOLDIN, J. and YANG, C. S. (2018). The effects of pre-trial detention on conviction, future crime, and employment: Evidence from randomly assigned judges. *American Economic Review* **108** 201–240.
- DORN, J. and GUO, K. (2022). Sharp sensitivity analysis for inverse propensity weighting via quantile balancing. *Journal of the American Statistical Association* 1–13.
- ESPENSHADE, T. J., CHUNG, C. Y. and WALLING, J. L. (2004). Admission preferences for minority students, athletes, and legacies at elite universities. *Social Science Quarterly* **85** 1422–1446.
- FAGAN, J. (2010). Report of Jeffrey Fagan in the case of Floyd v. the City of New York. Available at <https://www.law.columbia.edu/sites/default/files/microsites/policing-litigation-conference/files/Fagan%20Report%20with%20Technical%20Appendices.pdf>.
- GAEBLER, J., CAI, W., BASSE, G., SHROFF, R., GOEL, S. and HILL, J. (2022). A causal framework for observational studies of discrimination. *Statistics and Public Policy* **9** 26–48.
- GELMAN, A., FAGAN, J. and KISS, A. (2007). An analysis of the New York City police department’s “stop-and-frisk” policy in the context of claims of racial bias. *Journal of the American Statistical Association* **102** 813–823.
- GOEL, S., RAO, J. M. and SHROFF, R. (2016a). Precinct or Prejudice? Understanding Racial Disparities in New York City’s Stop-and-Frisk Policy. *Annals of Applied Statistics* **10**.
- GOEL, S., RAO, J. M. and SHROFF, R. (2016b). Personalized Risk Assessments in the Criminal Justice System. *The American Economic Review* **106** 119–123.
- GOEL, S., PERELMAN, M., SHROFF, R. and SKLANSKY, D. (2017). Combatting Police Discrimination in the Age of Big Data. *New Criminal Law Review* **20**.

- GOEL, S., SHROFF, R., SKEEM, J. and SLOBOGIN, C. (2021). The accuracy, equity, and jurisprudence of criminal risk assessment. In *Research Handbook on Big Data Law* Edward Elgar Publishing.
- GREINER, D. J. and RUBIN, D. B. (2011). Causal effects of perceived immutable characteristics. *Review of Economics and Statistics* **93** 775–785.
- GROSSMAN, J., TOMKINS, S., PAGE, L. and GOEL, S. (2023). The Disparate Impacts of College Admissions Policies on Asian American Applicants. *Working paper*.
- HARDT, M., PRICE, E. and SREBRO, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems* **29**.
- HUANG, M. and PIMENTEL, S. D. (2022). Variance-based sensitivity analysis for weighting estimators result in more informative bounds. *arXiv preprint arXiv:2208.01691*.
- JUNG, J., CONCANNON, C., SHROFF, R., GOEL, S. and GOLDSTEIN, D. G. (2020). Simple rules to guide expert classifications. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **183** 771–800.
- KARUSH, W. (1939). Minima of functions of several variables with inequalities as side constraints. *M. Sc. Dissertation. Dept. of Mathematics, Univ. of Chicago*.
- KLEINBERG, J., LAKKARAJU, H., LESKOVEC, J., LUDWIG, J. and MULLAINATHAN, S. (2018). Human decisions and machine predictions. *The quarterly journal of economics* **133** 237–293.
- KNOWLES, J., PERSICO, N. and TODD, P. (2001). Racial Bias in Motor Vehicle Searches: Theory and Evidence. *Journal of Political Economy* **109**.
- KUHN, H. W. and TUCKER, A. W. (1951). Nonlinear programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950* pp 481–492. University of California Press, Berkeley-Los Angeles, Calif. [MR47303](#)
- MACDONALD, J. and RAPHAEL, S. (2019). The effect of scaling back punishment in racial disparities in criminal case outcomes. *University of California Berkeley Working Paper*.
- MAYSON, S. G. (2018). Bias in, bias out. *Yale Law Journal* **128** 2218.
- MONAHAN, J. and SKEEM, J. L. (2016). Risk assessment in criminal sentencing. *Annual review of clinical psychology* **12** 489–513.
- NILFOROSHAN, H., GAEBLER, J. D., SHROFF, R. and GOEL, S. (2022). Causal conceptions of fairness and their consequences. In *International Conference on Machine Learning* 16848–16887. PMLR.
- PHELPS, E. S. (1972). The statistical theory of racism and sexism. *The American Economic Review* **62** 659–661.
- PIERSON, E., CORBETT-DAVIES, S. and GOEL, S. (2018). Fast threshold tests for detecting discrimination. In *The 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- PIERSON, E., SIMOIU, C., OVERGOOR, J., CORBETT-DAVIES, S., JENSON, D., SHOEMAKER, A., RAMACHANDRAN, V., BARGHOUTY, P., PHILLIPS, C., SHROFF, R. et al. (2020). A large-scale analysis of racial disparities in police stops across the United States. *Nature Human Behaviour* **4** 736–745.
- POLACHEK, S. W. (2008). Earnings over the life cycle: The Mincer earnings function and its applications. *Foundations and Trends in Microeconomics* **4** 165–272.
- ROSENBAUM, P. R. (2002). *Observational studies*. Springer.
- ROSENBAUM, P. R. and RUBIN, D. B. (1983). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society. Series B (Methodological)* **45** 212–218.
- SAHNI, S. (1974). Computationally related problems. *SIAM Journal on computing* **3** 262–279.
- SHROFF, R. (2017). Predictive Analytics for City Agencies: Lessons from Children’s Services. *Big Data* **5** 189–196.
- SIMOIU, C., CORBETT-DAVIES, S. and GOEL, S. (2017). The problem of infra-marginality in outcome tests for discrimination. *The Annals of Applied Statistics* **11** 1193–1216.
- SMART, S. and WALDFOGEL, J. (1996). A citation-based test for discrimination at economics and finance journals Technical Report, National Bureau of Economic Research.
- WASHINGTON v. DAVIS (1976). 426 U.S. 229.
- SFFA v. HARVARD (2023). Students for Fair Admissions, Inc., Petitioner, v. President and Fellows of Harvard College. Students for Fair Admissions, Inc., Petitioner, v. University of North Carolina, et al. https://www.supremecourt.gov/opinions/22pdf/20-1199_l6gn.pdf.
- VANDERWEELE, T. J. and ROBINSON, W. R. (2014). Confounding and mediating variables. *Epidemiology* **25** 473–484.
- FLOYD v. CITY OF NEW YORK (2013). 959 F. Supp. 2d 540, S.D.N.Y.
- ZHANG, Y. and ZHAO, Q. (2022). Bounds and semiparametric inference in L^∞ - and L^2 -sensitivity analysis for observational studies. *arXiv preprint arXiv:2211.04697*.
- ZHAO, Q., SMALL, D. S. and BHATTACHARYA, B. B. (2019). Sensitivity analysis for inverse probability weighting estimators via the percentile bootstrap. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **81** 735–761.