

Omitted and Included Variable Bias in Tests for Disparate Impact*

Jongbin Jung
Stanford University

Sam Corbett-Davies
Stanford University

Ravi Shroff
New York University

Sharad Goel
Stanford University

August 29, 2019

Abstract

Policymakers often seek to gauge discrimination against groups defined by race, gender, and other protected attributes. One popular strategy is to estimate disparities after controlling for observed covariates, typically with a regression model. This approach, however, suffers from two statistical challenges. First, omitted-variable bias can skew results if the model does not control for all relevant factors; second, and conversely, included-variable bias can skew results if the set of controls includes irrelevant factors. Here we introduce a simple three-step strategy—which we call risk-adjusted regression—that addresses both concerns in settings where decision makers have clearly measurable objectives. In the first step, we use all available covariates to estimate the utility of possible decisions. In the second step, we measure disparities after controlling for these utility estimates alone, mitigating the problem of included-variable bias. Finally, in the third step, we examine the sensitivity of results to unmeasured confounding, addressing concerns about omitted-variable bias. We demonstrate this method 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 yield misleading results. We conclude by discussing implications of our statistical approach for questions of law and policy.

Keywords— Discrimination, risk-adjusted regression, sensitivity analysis

*We thank Guillaume Basse, Avi Feller, and Aziz Huq for extended discussions and comments. Code and data to reproduce our analysis are available at <https://github.com/stanford-policylab/risk-adjusted-regression>.

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 white loan applicants receive credit extensions more often than 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 differences in group composition. 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 composition versus discrimination—the most popular statistical approach is ordinary linear or logistic regression. In the banking context, for example, one could examine race-contingent lending rates after controlling for relevant factors, such as income and credit history. Disparities that persist after accounting for such differences in group composition are often interpreted as evidence of discrimination. This basic statistical strategy has been used in hundreds of studies to test for bias in dozens of domains, including education [Espenshade et al., 2004], employment [Polachek et al., 2008], policing [Gelman et al., 2007], and medicine [Balsa et al., 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. 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 protected attributes, skewing results.

The 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 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.

In this paper, we develop a statistically principled and 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 protected attributes and their potential proxies, to estimate the expected utility of taking a particular action. For example, in the lending context, we might estimate an applicant’s risk of default if granted a loan (or equivalently, likelihood of repayment if granted a loan) conditional on all available covariates. In the second step, we assess disparities via a regression model where we control only for individual-level risk scores, and accordingly measure the extent to which similarly qualified individuals are treated similarly. 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 included-variable bias. Finally, we adapt the classical method of Rosenbaum and Rubin [1983] to estimate the sensitivity of our estimates to unmeasured confounding, addressing the problem of omitted-variable bias.

To demonstrate this technique, we examine 2.2 million stops of pedestrians conducted by the New York Police Department between 2008 and 2011. Controlling for a stopped individual’s statistical risk of carrying a weapon—based in part on detailed behavioral indicators recorded by officers—we find that blacks and Hispanics are searched for weapons more often than whites. These risk-adjusted disparities are substantially larger than disparities suggested by a standard regression that controls for all available

covariates, underscoring the importance of accounting for included-variable bias. Our results persist after allowing for search policies that may differ by location, and appear robust to the possibility of substantial unmeasured confounding.

2 Background and related work

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, certain affirmative action programs for college admissions are legally permissible to further the government’s interest in promoting diversity [Fisher v. University of Texas, 2016].

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 protected traits. In this approach, the investigator controls 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; and in the absence of a compelling justification, such evidence is suggestive of a discriminatory purpose.¹ We note that it is difficult—and perhaps impossible—to rigorously define the *influence*, or causal effect, of largely immutable traits like race and gender on decisions [Greiner and Rubin, 2011, VanderWeele and Robinson, 2014]. 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]. However, for an equal protection claim to succeed in court, 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 predictive 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

¹The regression-based approach to discrimination is closely related to *benchmark analysis*. In benchmark analysis, one compares the demographic composition of those receiving favorable decisions to the composition of a “qualified pool” [Ridgeway and MacDonald, 2010]. For example, one might compare the proportion of loan recipients who are minorities to the proportion of minority applicants having a certain minimum credit rating. Deciding which individuals are “qualified” is analogous to deciding which variables to control for in a regression model.

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.”

By including educational status in the regression, one would mask the policy’s unjustified disparate impact. We note that such included-variable bias can also conceal discriminatory intent, as a decision maker may have purposely adopted a sub-optimal policy to further a prejudicial goal.² (In reality, the Court found no evidence that the Duke Power Company’s policy was enacted with discriminatory purpose; it held only that the education requirement was unrelated to job performance and so had an illegal disparate impact.)

In general, to assess claims of disparate impact, 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) control for sufficiently many covariates, omitted-variable bias may skew results; conversely, if one does control for a rich set of covariates, included-variable bias may skew results. Such 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 minorities are less creditworthy than whites, minorities who are granted loans should still be found to repay their loans at the same rate as whites who are granted loans. If loans to minorities have a higher repayment rate than loans to whites, it suggests that lenders are effectively applying a double standard (intentionally or not), granting loans only to exceptionally qualified minorities. Such a finding would be evidence of a disparate impact violation. The outcome test has been applied almost as broadly as simple regression, to study discrimination in policing [Ayres, 2002, Goel et al., 2016b, 2017, Knowles et al., 2001], lending [Berkovec et al., 1996], and scientific publication [Smart and Waldfogel, 1996].

Outcome tests, however, have their own statistical shortcomings. To see this, suppose that there are two, easily distinguishable types of white loan applicants: those who have a 95% chance of repayment, and those who have a 50% chance of repayment. Similarly assume that black loan applicants have either a 99% or 50% chance of repayment. If bank officers, in a race-neutral manner, approve loans to all applicants at least 90% likely to repay their loans, then loans to whites will be repaid 95% of the time whereas loans to blacks will be repaid 99% of the time. In this stylized example, the outcome test would (incorrectly) suggest a double standard, with only exceptionally qualified minorities granted loans. This limitation of outcome tests is known as the problem of infra-marginality, and it stems from the fact that a group’s aggregate repayment rate is an average over individuals with different risk levels [Ayres, 2002, Corbett-Davies and Goel, 2018, Corbett-Davies et al., 2017, Simoiu et al., 2017].

The problem of infra-marginality in outcome tests is more than a hypothetical possibility. Analyzing policing patterns in North Carolina, Simoiu et al. [2017] found that infra-marginality likely caused the outcome test to yield misleading results. To address this issue, they introduced the threshold test, which uses a Bayesian strategy to jointly infer group-specific risk distributions and decision thresholds. The test has since been applied in several studies, including a large-scale analysis of traffic stops across the country [Pierson et al., 2018a,b]. While a significant step forward, the threshold test has two notable limitations. First, the test is identified in part by the prior distributions on the parameters of interest and by the specific structural form of the model—not by the data alone. The authors suggest several diagnostics to assess the robustness of conclusions, but they also note this non-identifiability as an important caveat of the approach. Second, and perhaps more importantly, the threshold test requires considerable statistical knowledge to understand and to appropriately carry out, hindering adoption by practitioners with less technical training.

²One example is the historical practice of *redlining*—where borrowers from minority neighborhoods were denied mortgages and federal loan insurance. A disparate treatment analysis of redlining that controlled for the neighborhood of the borrower would find no evidence of racial discrimination, concealing the fact that this is exactly the mechanism by which an intentionally discriminatory outcome is being achieved.

3 A new statistical test of disparate impact

Our proposed test of disparate impact, which we call risk-adjusted regression, measures the extent to which individuals of similar risk (or equivalently, similar qualifications) from different groups are treated similarly. For example, in the banking context, the test can assess whether black and white loan applicants with similar risk of default are approved at similar rates; and in the policing context, the test can assess whether black and white individuals with similar likelihood of carrying a weapon are searched for weapons at similar rates. By estimating decision rates conditional on risk—rather than on a specific set of covariates—the test avoids the problem of included-variable bias. The test also avoids the problem of infra-marginality, since it computes decision rates conditional on a specific risk level rather than computing an average over a distribution of risk levels. The test cannot completely circumvent the problem of omitted-variable bias, but we present an approach to assess the sensitivity of estimates to potential unmeasured confounding.

3.1 Estimating risk-adjusted disparities

The first two steps of our risk-adjusted regression address potential included-variable bias. First, for each individual, we estimate risk as a function of all available covariates, including membership in protected classes such as race and gender.³ We then control for this risk estimate in a regression of the (binary) decision a on the protected class c to compute *risk-adjusted disparities*, the difference in average decision rates between groups after adjusting for individual-level risk. (The third step of our strategy, in which we account for omitted-variable bias, is described in Section 3.2.)

To carry out this procedure, suppose we have data of the form $\Omega = \{(c_i, x_i, a_i, r_i)\}$, where for each observation i , $c_i \in \mathcal{C}$ indicates membership in a protected class, x_i denotes all other information available to the decision maker prior to selecting a course of action, $a_i \in \{0, 1\}$ is the selected action, and $r_i \in \{0, 1\}$ indicates the outcome of interest. As a concrete example, consider estimating risk-adjusted racial disparities in police searches of pedestrians for weapons—the application we discuss in Section 4. In this case, c_i , x_i , a_i , and r_i would, respectively, indicate a stopped pedestrian’s race, all information available prior to the search, whether the individual was searched, and whether a weapon was found. Now, let $r_i(0)$ and $r_i(1)$ denote potential outcomes under the two possible actions a one could take, where only $r_i = r_i(a_i)$ is observed for each case. For example, $r_i(1)$ might indicate whether a weapon would be discovered if pedestrian i were searched, and $r_i(0)$ whether a weapon would be discovered if that pedestrian were not searched. (In this particular example, $r_i(0) \equiv 0$, since a weapon cannot be found if a search is not conducted, though that need not be the case in all applications.⁴)

Next, for each individual, we define the ex-ante *risk* $\mu_i = \Pr(r_i(1) = 1 \mid x_i, c_i)$. In our policing example, μ_i is the probability of finding a weapon if an individual with characteristics $\{x_i, c_i\}$ is searched. As outlined above, the estimand of interest is a risk-adjusted measure of average differences in group decision rates. Specifically, in the case of two protected groups $\mathcal{C} = \{c_0, c_1\}$, we define our estimand non-parametrically as follows:

$$\int_0^1 \left[\log \frac{\Pr(a = 1 \mid c = c_1, \mu)}{1 - \Pr(a = 1 \mid c = c_1, \mu)} - \log \frac{\Pr(a = 1 \mid c = c_0, \mu)}{1 - \Pr(a = 1 \mid c = c_0, \mu)} \right] f(\mu) d\mu, \quad (3.1)$$

³In some cases, protected covariates can add substantial predictive power, and thus excluding them can lead to poor estimates of risk. For example, in statistical estimates of recidivism, women have been found to reoffend less often than men with similar criminal profiles [Corbett-Davies and Goel, 2018, Skeem et al., 2016]. Consequently, by excluding gender from such models, one would over-estimate the recidivism risk of women, potentially skewing estimates of risk-adjusted disparities. This phenomenon is closely related to what Arrow et al. [1973] call *statistical discrimination*.

⁴Suppose, for example, that we seek to estimate risk-adjusted disparities in judicial bail decision, where the outcome r indicates whether a defendant fails to show up at required court proceedings [Jung et al., 2017, 2018]. Then $r_i(0)$ might indicate failure to appear if the judge requires bail as a condition of release, and $r_i(1)$ might indicate failure to appear if the judge releases a defendant on his or her own recognizance. In this case, both $r_i(0)$ and $r_i(1)$ could be positive, since requiring bail decreases—but does not eliminate—flight risk.

where $f(\mu)$ is the population density function of risk (i.e., f gives the distribution of μ in the population). Risk-adjusted disparities are hence defined to be the difference in log-odds of taking an action for one group relative to another, averaged over the distribution of risk.⁵

To facilitate computation, we approximate the estimand in (3.1) by a logistic regression model—what we call a risk-adjusted regression:

$$\Pr(a_i = 1 \mid c_i, \mu_i) = \text{logit}^{-1}(\theta_{c_i} + g_{\theta}(\mu_i)), \quad (3.2)$$

where θ_c is the coefficient for membership in group c , and g_{θ} is an appropriately selected transformation of risk to allow for complex relationships between risk and decisions (as discussed below). We set $\theta_c = 0$ for the base group (typically the majority group); accordingly, in the case of two groups— c_0 and c_1 , where c_0 is the base group— $\hat{\theta}_{c_1}$ yields an estimate of (3.1). Under this model, positive values of $\hat{\theta}_{c_1}$ indicate that members of group c_1 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 groups with positive $\hat{\theta}_{c_1}$ are searched more often than members of the base group who were equally likely to be carrying a weapon, and we would say that such elevated search rates are *unjustified* by risk. Importantly, positive values of $\hat{\theta}_{c_1}$ do not imply intentional discrimination—as in *Griggs*, unjustified disparate impact is possible even under a facially neutral policy undertaken without animus.

In such a risk-adjusted regression, one should choose the transformation g_{θ} to suit the application. In some cases, we may believe the log-odds of taking action a is approximately proportional to the log-odds of risk, suggesting a logit-linear model:

$$g_{\theta}(\mu_i) = \theta_0 + \theta_1 \text{logit}(\mu_i). \quad (3.3)$$

If, however, there is reason to believe that decision makers are applying a *threshold rule* (where the probability of taking a certain action rapidly increases from near-zero to near-one at some risk threshold), a linear model would not be appropriate. In that case, binning the risk (into K fixed bins B_k) may better capture the decision-making process:

$$g_{\theta}(\mu_i) = \sum_{k=1}^K \theta_k \mathbf{1}\{\mu_i \in B_k\}. \quad (3.4)$$

We might alternatively use splines, or other functional forms, to express more complicated relationships between risk and decisions. No matter what form of g is ultimately chosen, one must ensure that any risk-adjusted disparities are robust to reasonable perturbations, as we illustrate in Section 4.

Our formulation implicitly assumes that the value of a decision depends only on the risk μ_i . While reasonable in many cases, this may not hold in general. For example, if confiscating a weapon in a school zone has higher value than recovering one elsewhere, then the probability of weapon recovery would be an imperfect proxy for utility. To address this concern, one could replace risk in Eq. (3.2) with a measure of context-specific utility—though that approach can be difficult in practice. Alternatively, one could control for the factors by which utility varies, such as proximity to a school. Specifically, we can fit the model:

$$\Pr(a_i = 1 \mid c_i, \mu_i, x_i^{\ell}) = \text{logit}^{-1}(\theta_{c_i} + g_{\theta}(\mu_i) + \theta_{\ell}^{\top} x_i^{\ell}), \quad (3.5)$$

where $x_i^{\ell} \subseteq x_i$ are the relevant non-risk factors, and θ_{ℓ} is a vector of coefficients for the covariates x^{ℓ} . As before, the fitted coefficients $\hat{\theta}_c$ provide a measure of unjustified disparities. While easier to carry out, this strategy can re-introduce included-variable bias, since the non-risk factors are only partially related to utility;

⁵As we discuss further in Section 5, different methods of averaging may result in different substantive interpretations, which points to an inherent ambiguity in legal and colloquial understandings of “disparate impact”. For ease of exposition, we define risk only in terms of the potential outcome $r(1)$, and our risk-adjusted regression accordingly controls only for that risk. This setup is sufficient in situations where $r(0)$ is naturally zero, including many employment, lending, and policing applications. But in some settings—as described in Footnote 4—we require more generality. It is straightforward to extend the approach we outline here, including the sensitivity analysis of Section 3.2, to such scenarios.

as such, one might underestimate the actual unjustified disparities. These difficulties highlight the inherent challenges of measuring disparities when utilities are imprecisely specified.

Finally, we note that μ_i cannot typically be computed without additional assumptions, as it involves the potentially unobserved outcome $r(1)$. We thus further define the *risk function* $f(c, x)$:

$$f(c, x) = \Pr(r(1) = 1 \mid a = 1, c, x). \tag{3.6}$$

Because the probability in Eq. (3.6) is conditioned on $a = 1$, $f(c, x)$ can be estimated from the observed features and decisions. For example, an accurate estimate \hat{f} of f can be constructed by first limiting to cases with $a = 1$, since $r(1)$ is observed in those instances. If we assume that a is ignorable given c and x , then $\hat{\mu}_i = \hat{f}(c_i, x_i)$ is an accurate estimate of μ_i for all individuals, including those for whom the action $a = 0$ was taken. We address concerns regarding possible confounding in the next section.

3.2 Assessing sensitivity to unmeasured confounding

We estimated risk μ_i above by assuming that decisions were ignorable given the observed covariates c and x . Formally, this ignorability assumption means that

$$r(1) \perp\!\!\!\perp a \mid c, x. \tag{3.7}$$

In practice, however, it is likely that decision makers observe unrecorded information which is predictive of the potential outcome $r(1)$ and is therefore used to inform their actions, violating ignorability. As an example, a stopped pedestrian’s response to police questioning may legitimately alter an officer’s estimate of risk, and thus the officer’s decision to search. In this case, it may be that searched individuals were, on average, more likely to be carrying a weapon compared to those with the same observed covariates who were not searched, since the former group is more likely to have provided suspicious answers to an officer’s questions. As such, estimates $\hat{f}(c, x)$ fit on searched pedestrians would systematically overstate risk for those who were not searched, in turn corrupting estimates of risk-adjusted disparities. Unfortunately, it is typically impossible to directly account for every factor that may plausibly affect risk—unlike legitimate *non-risk* variables x^ℓ which often can be explicitly enumerated.

We address this issue by adapting the method of Rosenbaum and Rubin [1983], as recently extended by Jung et al. [2017], for assessing the sensitivity of estimated causal effects to an unobserved binary covariate. We stress, though, that we are not attempting to infer the causal effect of race on decisions, in contrast to studies of disparate treatment. Indeed, as discussed above, it is not entirely clear even how one can rigorously define such a race effect. We are instead measuring group-specific disparities—as formalized in Eq. (3.1)—and the Rosenbaum and Rubin approach allows us to gauge the robustness of our results to estimation error in the risk μ .

To start, we assume there exists an unobserved covariate $u \in \{0, 1\}$ that affects both the decision a (e.g., whether or not to carry out a search) and also the potential outcome $r(1)$ (e.g., recovery of a weapon if a search were conducted). Our key assumption is that the observed action is ignorable given the observed covariates c, x , and the unobserved covariate u :

$$r(1) \perp\!\!\!\perp a \mid c, x, u. \tag{3.8}$$

This model of confounding has three important parameters, each which may depend on the observed covariates c and x : (1) the effect of u on the action a ; (2) the effect of u on the potential outcome $r(1)$; and (3) the probability that $u = 1$. As we show below, once these parameters are specified, one can compute risk-adjusted disparities that account for the unmeasured confounding. Because the confounding is, by definition, unobserved, we cannot infer these parameters from the data. We can, however, specify plausible ranges for them and thus gauge the sensitivity of estimated disparities, as described in Section 4.3.

Now, without loss of generality, we can write

$$\Pr(a = 1 \mid c, x, u) = \text{logit}^{-1}(\gamma_{c,x} + u\alpha_{c,x}) \tag{3.9}$$

for appropriately chosen parameters $\gamma_{c,x}$ and $\alpha_{c,x}$ that depend on the observed covariates c and x . Here $\alpha_{c,x}$ is the change in log-odds of taking action $a = 1$ when $u = 0$ versus when $u = 1$. We can similarly write

$$\Pr(r(1) = 1 \mid c, x, u) = \text{logit}^{-1}(\beta_{c,x} + u\delta_{c,x}) \tag{3.10}$$

for parameters $\beta_{c,x}$ and $\delta_{c,x}$. In the case of police searches, $\delta_{c,x}$ is the change in log-odds of recovering a weapon if searched when $u = 0$ versus when $u = 1$.

For any posited values of the three parameters $\Pr(u = 1 | c, x)$, $\alpha_{c,x}$, and $\delta_{c,x}$, we can use the observed data to estimate $\gamma_{c,x}$ and $\beta_{c,x}$, as described in Rosenbaum and Rubin [1983] and Jung et al. [2017]. To see this, first note that $\Pr(a = 1 | c, x)$ can be decomposed into two components conditioned on $u = 0$ and $u = 1$, respectively:

$$\Pr(a = 1 | c, x) = \Pr(u = 0 | c, x) \text{logit}^{-1}(\gamma_{c,x}) + \Pr(u = 1 | c, x) \text{logit}^{-1}(\gamma_{c,x} + \alpha_{c,x}). \quad (3.11)$$

The left-hand side of Eq. (3.11) depends only on the observed quantities a , c , and x , and so can be directly estimated from the data (e.g., via a logistic regression). The right-hand side is a continuous, increasing function of $\gamma_{c,x}$, which takes values from 0 to 1 as $\gamma_{c,x}$ goes from $-\infty$ to ∞ . Thus, there is a unique value of $\gamma_{c,x}$ that ensures the equality in Eq. (3.11) is satisfied.

To estimate $\beta_{c,x}$, we start by estimating the conditional distribution of u , given that the action $a = 1$ was taken. By Bayes' rule:

$$\Pr(u = 1 | a = 1, c, x) = \frac{\Pr(a = 1 | u = 1, c, x) \Pr(u = 1 | c, x)}{\sum_{u \in \{0,1\}} \Pr(a = 1 | u, c, x) \Pr(u | c, x)}. \quad (3.12)$$

The right-hand side can be estimated using Eq. (3.9) and our estimate of $\gamma_{c,x}$ computed above, yielding an approximation of the left-hand side.

We now write:

$$\begin{aligned} \Pr(r(1) = 1 | a = 1, c, x) &= \sum_{u \in \{0,1\}} \Pr(r(1) = 1 | a = 1, u, c, x) \Pr(u | a = 1, c, x) \\ &= \sum_{u \in \{0,1\}} \Pr(r(1) = 1 | u, c, x) \Pr(u | a = 1, c, x) \\ &= \sum_{u \in \{0,1\}} \text{logit}^{-1}(\beta_{c,x} + u\delta_{c,x}) \Pr(u | a = 1, c, x). \end{aligned} \quad (3.13)$$

The second equality above follows from the ignorability assumption stated in Eq. (3.8), and the third equality follows from Eq. (3.10). The left-hand side is the risk function $f(c, x)$ defined in Eq. (3.6), which can be estimated from the observed data. Given the estimate of $\Pr(u = 1 | a = 1, c, x)$ from above, and our assumed value of $\delta_{c,x}$, the only unknown on the right-hand side is $\beta_{c,x}$. As before, there is a unique value of $\beta_{c,x}$ that ensures the equality is satisfied.

To summarize, for any given values of $\Pr(u = 1 | c, x)$, $\alpha_{c,x}$, and $\delta_{c,x}$, we can derive estimates for $\gamma_{c,x}$ and $\beta_{c,x}$. In particular, these posited and derived parameter values together yield estimates of $\Pr(a | c, x, u)$ and $\Pr(r(1) | c, x, u)$ by Eqs. (3.9) and (3.10), respectively. As a result, we can also estimate the full joint distribution of $\{r(1), a, c, x, u\}$, since

$$\Pr(r(1), a, c, x, u) = \Pr(r(1) | c, x, u) \Pr(a | c, x, u) \Pr(u | c, x) \Pr(c, x), \quad (3.14)$$

where we have used ignorability to write the first two terms in the product. (Note that the final term in the product, $\Pr(c, x)$, can be estimated from the empirical joint distribution of c and x in the observed data.)

We now have nearly all the necessary ingredients to assess the sensitivity of our risk-adjusted disparity estimate, as given by Eq. (3.5). To conclude, imagine drawing a new, synthetic dataset of n samples $\{(\tilde{r}_i(1), \tilde{a}_i, \tilde{c}_i, \tilde{x}_i, \tilde{u}_i)\}_{i=1}^n$ according to the (estimated) joint distribution in Eq. (3.14). For each datapoint, we then use Eq. (3.10) to compute the ex-ante risk $\tilde{\mu}_i = \Pr(\tilde{r}_i(1) | \tilde{c}_i, \tilde{x}_i, \tilde{u}_i)$, yielding the augmented set of tuples $\{(\tilde{r}_i(1), \tilde{a}_i, \tilde{c}_i, \tilde{x}_i, \tilde{u}_i, \tilde{\mu}_i)\}_{i=1}^n$. Finally, we compute $\hat{\theta}_c^n$ —the estimated risk-adjusted disparity for group c —on this sampled dataset via our usual logistic regression estimator:

$$\Pr(\tilde{a}_i = 1 | \tilde{c}_i, \tilde{\mu}_i, \tilde{x}_i^\ell) = \text{logit}^{-1}(\theta_{\tilde{c}_i} + g_\theta(\tilde{\mu}_i) + \theta_\ell^\top \tilde{x}_i^\ell). \quad (3.15)$$

The fitted coefficient $\hat{\theta}_c^n$ is a random variable that depends on n , and its limiting value as n goes to infinity yields an estimate of risk-adjusted disparities that takes potential confounding into account.

One could approximate the limiting value of $\hat{\theta}_c^n$ by simply sampling a large number of datapoints and fitting the regression above. However, in Appendix A, we present a more computationally efficient approach for computing this limit. At a high level, we first construct an expanded dataset Ω' by combining two copies of the observed data Ω —one with u set to 0, and another with u set to 1 (recall that u is not present in the original dataset). We then fit a fractional-response logistic regression [Papke and Wooldridge, 1996] on the doubled dataset Ω' , weighting each datapoint by either $\Pr(u = 0 \mid c_i, x_i)$ or $\Pr(u = 1 \mid c_i, x_i)$, and using $\Pr(a_i = 1 \mid c_i, x_i, u)$ (computed from Eq. (3.9)) as the response variable.

4 Testing for disparate impact in police stops

We apply our risk-adjusted regression developed above to investigate the “stop-and-frisk” practices of the New York Police Department (NYPD). 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. Though a policy of stopping and frisking individuals is not inherently illegal, a federal district court ruled that the NYPD carried out such stops with racial animus, violating the equal protection clause of the Fourteenth Amendment [Floyd v. City of New York, 2013].

The court in *Floyd* was interested in assessing claims of disparate treatment; here we re-analyze the data to test directly for unjustified disparate impact. We specifically focus on 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. Frisk decisions are thus particularly amenable to our statistical approach. 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).⁶

To start, we note that white pedestrians are frisked in 44% of police stops, whereas black and Hispanic pedestrians are frisked in 57% and 58% of stops, respectively. This corresponds to stopped minority pedestrians having about 1.7 times greater odds of being frisked than whites. These raw disparities are computed without controlling for any potentially explanatory variables, and so represent an extreme case of omitted-variable bias. At the other extreme is the “kitchen sink regression”, which controls for all pre-frisk covariates in a standard logistic regression model. In this case, stopped blacks and Hispanics have about 1.2 times the odds of being frisked relative to whites. Though 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), they may understate the extent to which the policy imposes an unjustified disparate impact on minorities, due to included-variable bias.

In Section 4.2, we apply our risk-adjusted regression to the New York dataset to measure unjustified racial disparities in frisk decisions. Then, in Section 4.3, we gauge the sensitivity of our estimated risk-adjusted disparities to potential omitted-variable bias. First, though, we carry out a simulation study on synthetic data to assess the quality of our estimator.

4.1 A simulation study

To conduct our simulation study, we first create a synthetic dataset based on the NYPD data. On this synthetic dataset, we compare the true value of the risk adjusted disparities—as defined in Eq. (3.1)—to estimates from both a kitchen-sink and a risk-adjusted regression.

We create the synthetic dataset in several steps. First, we use the real NYPD data to estimate the probability of frisk, $\hat{p}_f(x)$, conditional on all observed pre-frisk covariates x , excluding race. We similarly estimate the probability a weapon is recovered on a frisked individual, $\hat{p}_w(x)$, again conditional on all observed pre-frisk covariates x , excluding race. In both cases, we use gradient-boosted decision trees, a

⁶This 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 at: <https://www.prisonlegalnews.org/news/publications/blank-uf-250-form-stop-question-and-frisk-report-worksheet-nypd-2016/>.

popular, non-linear machine-learning model for estimating the conditional probability of binary outcomes.⁷ The frisk model is trained on the entire NYPD dataset, and the weapon-recovery model is trained on the subset of stops in which an individual was frisked.

With these fitted models $\hat{p}_f(x)$ and $\hat{p}_w(x)$ in hand, we generate a synthetic dataset by repeating the following four-step procedure $M = 10^7$ times.

1. Sample, with replacement, a covariate vector and race pair (x_i, c_i) from the original NYPD dataset.
2. Sample a perturbation parameter $\varepsilon_i \sim N(0, \sigma^2)$, where $\sigma = 0.1$.
3. Define the probability the i -th individual is frisked: $p_i = \text{logit}^{-1}(\text{logit}(\hat{p}_f(x_i)) + \varepsilon_i)$. Then sample $a_i \sim \text{Bernoulli}(p_i)$.
4. Define the probability the i -th individual possesses a weapon: $\mu_i = \text{logit}^{-1}(\text{logit}(\hat{p}_w(x_i)) + \varepsilon_i)$. Then, if $a_i = 1$, sample $r_i \sim \text{Bernoulli}(\mu_i)$; otherwise set $r_i = 0$.

The above procedure produces a set of tuples, $\Omega = \{(x_i, c_i, a_i, r_i, \mu_i)\}_{i=1}^M$, which is our synthetic population. Now, on this population, we compute the ground-truth risk-adjusted disparities via Eq. (3.1).⁸ Setting the reference group to be white pedestrians, we find the risk-adjusted disparity is 1.6 for black pedestrians and 1.7 for Hispanic pedestrians, as indicated by the “ground truth” X’s in Figure 1. 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 a kitchen-sink regression and our own risk-adjusted regression to recover the ground-truth disparate impact. Starting from the synthetic population Ω , we first create a restricted dataset $\bar{\Omega} \subsetneq \Omega$ of 100,000 observations: $\bar{\Omega} = \{(x_i, c_i, a_i, r_i)\}_{i=1}^{100,000}$. This smaller dataset mimics the fact that researchers applying these methods typically do not have data on the full population, nor do they directly observe the risk μ_i . The blue points in Figure 1 show the results of fitting a kitchen-sink regression on $\bar{\Omega}$. After adjusting for all the covariates x , frisk rates are comparable across white, black, and Hispanic pedestrians. The kitchen-sink regression thus accurately captures the fact that there is no disparate treatment in this synthetic example—as discussed above—but it dramatically underestimates the true level of disparate impact.

Finally, we compute risk-adjusted disparities via our estimator in Eq. (3.2). To do so, we first divide $\bar{\Omega}$ into two random sets of equal size, $\bar{\Omega}_1$ and $\bar{\Omega}_2$. On $\bar{\Omega}_1$, we use gradient-boosted decision trees to estimate the likelihood that frisked individuals are found to have a weapon, based on all available pre-frisk covariates. This model is trained on the subset of stops in $\bar{\Omega}_1$ in which a frisk was carried out. We then use this fitted model to infer ex-ante risk μ_i for every stop in $\bar{\Omega}_2$. Lastly, we fit two separate risk-adjusted regressions, one to estimate the black-white disparity (fit only on stops of black and white pedestrians) and another to estimate the Hispanic-white disparity (fit only on stops of Hispanic and white pedestrians). For simplicity, we use a logit-linear transformation of risk g in the risk-adjusted regressions. Because estimated risk $\hat{\mu}$ is a data-dependant covariate, we estimate error bars for the risk-adjusted disparities via bootstrapping. Specifically, we estimate the standard error of estimates by repeating our entire inference procedure—including estimation of risk—on 100 bootstrap samples of $\bar{\Omega}$; 95% confidence intervals are computed around the point estimates by adding and subtracting twice the bootstrapped standard error. The resulting point estimates and error bars are displayed in Figure 1. For both black and Hispanic pedestrians in the synthetic data, our risk-adjusted estimates of disparate impact are nearly identical to the ground-truth estimands.

4.2 Estimating risk-adjusted disparities in the NYPD data

We now analyze the real stop-and-frisk data for disparate impact. We first divide the original stops into two sets, a subset of about one million stops that occurred between 2008 and 2009, and a subset of about 1.2 million stops that occurred between 2010 and 2011. Similar to our analysis above of the synthetic data, we begin by fitting a model to estimate the likelihood of weapon recovery, using only stops in the first subset

⁷The gradient-boosted decision tree model was fit using the `gbm` package in R [Greenwell et al., 2019].

⁸The definition in Eq. (3.1) involves an integral but our population is finite, so we must still select a bin size to approximate the integral. In this case, though, we get nearly identical values for any appropriately small bin size.

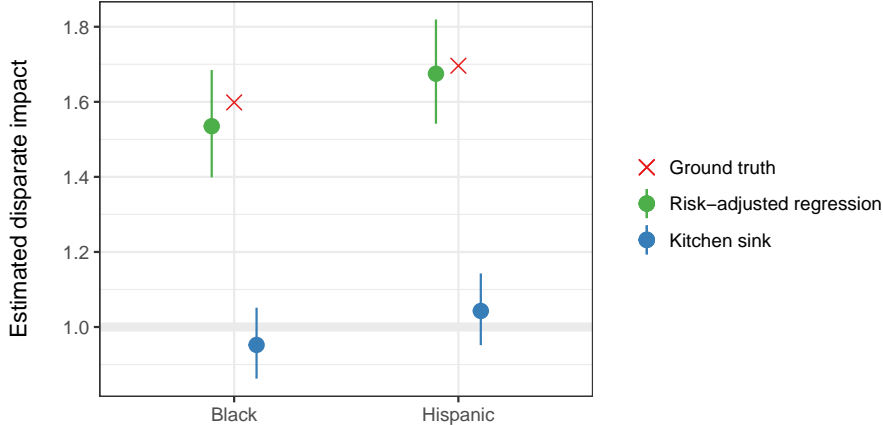


Figure 1: *Estimates of disparate impact derived from kitchen-sink and risk-adjusted regressions on a synthetic dataset. The risk-adjusted regressions recover the ground-truth values but, due to included-variable bias, the kitchen-sink regressions yield underestimates.*

of data that resulted in a frisk. Predictive performance and model checks presented in Appendix B indicate that the model yields predictions with reasonable performance that are well-calibrated across groups. We then use the fitted model to infer ex-ante risk μ_i for every pedestrian stopped in the second split of the data, including those who were not frisked. These risk estimates are implicitly based on an assumption of ignorability—an assumption we return to in Section 4.3.

The distributions of inferred risk, disaggregated by race, are shown in Figure 2. The plot illustrates several interesting patterns. First, the absolute level of risk is quite low, with even the riskiest pedestrians estimated—based on the recorded evidence—to be carrying weapons only about 10% of the time. This observation is consistent with the court’s ruling in *Floyd* that stops were often conducted without sufficient legal justification, in violation of the Fourth Amendment requirement that stops be based on “reasonable suspicion” [Goel et al., 2017].⁹ Further, stopped white pedestrians are, on average, more likely to be carrying a weapon than stopped minority pedestrians. In particular, 2.7% of stopped whites are estimated to carry weapons, compared to 1.5% of stopped blacks and 1.7% of stopped Hispanics. Following the logic of Becker’s outcome test, these differences provide some evidence of discrimination in the *stop decision*, since they suggest that officers stopped minorities on the basis of less evidence than whites.¹⁰

Risk-adjusted disparities are now computed as in Section 4.1; namely, we fit the risk-adjusted regression in Eq. (3.2) on the second half of the NYPD data, separately computing estimates for the black-white disparity θ_b and Hispanic-white disparity θ_h . 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 had about twice the odds of being frisked than white pedestrians who were equally likely to be carrying a weapon. We also find that these estimated disparities are robust to several different specifications of the risk transformation g : logit-linear (as in Eq. (3.3)), binned by decile (as in Eq. (3.4)), and thin-plate spline [Wood, 2003].¹¹ Further, the risk-adjusted disparities are in fact *greater* than the raw disparities in frisk rates, a finding consistent with the fact that stopped whites are, on average, more likely to be carrying a weapon than stopped minorities, as discussed above. 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 controls for a variety of

⁹The court ruled that the NYPD violated both the Fourth Amendment demand for reasonable suspicion when carrying out stops, and the Fourteenth Amendment prohibition against racial animus.

¹⁰In and of themselves, these differences provide only weak evidence of bias, as outcome tests suffer from the problem of infra-marginality (as discussed in Section 2), and officers may stop individuals for legitimate reasons other than suspected weapon possession, potentially producing the observed patterns in the absence of discrimination. We note, however, that more thorough investigations of the NYPD’s practices have indeed found compelling statistical evidence of bias in stop decisions [Gelman et al., 2007, Goel et al., 2016b, Pierson et al., 2018a].

¹¹The thin-plate regression spline was fit using the `mgcv` package in R.




Figure 2: *Distributions of the estimated 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 whites, 1.5% for blacks, and 1.7% for Hispanics.*

features—including whether the suspect made “furtive movements”—that are correlated with race but are poor predictors of whether a pedestrian is carrying a weapon, skewing estimates of disparate impact.¹²

Our analysis above assumes that risk of weapon possession is the only legitimate consideration for carrying out a frisk. It is possible, however, that there is a justifiable reason for frisking certain low-risk pedestrians more often than other, higher-risk individuals, and such a policy might in turn justify the risk-adjusted disparities we find between race groups. One possible policy justification is that different neighborhoods benefit from different enforcement standards. For example, frisks of low-risk pedestrians might be particularly effective at deterring criminal activity in high-crime neighborhoods by raising the perceived chance of getting caught; additionally, greater police presence in such areas might lower the effective costs of carrying out stops, making it feasible to frisk lower-risk individuals. While the merits of such policies are debated, we can still examine whether risk-adjusted disparities persist after controlling for location, via Eq. (3.5). We specifically control for the police precinct in which the stop occurred, as well as the location type (e.g., public housing or public transit). The right-most panel in Figure 3 shows that the disparities cannot be explained by policing practices that may differ by location. We find the estimated disparity for Hispanics decreases (from 1.9 to 1.6) but is still large, and the disparity for blacks is approximately unchanged (at 2.0).

4.3 Sensitivity to omitted-variable bias

The disparities computed above account for included-variable bias by adjusting for each individual’s estimated risk. But, as discussed earlier, our estimates of risk may be skewed if officers observe factors that are predictive of risk but are not recorded in our data. We now apply the method of Section 3.2 to gauge the sensitivity of our estimated risk-adjusted disparities to such potential omitted-variable bias.

First we estimate the left-hand side of Eq. (3.11), $\Pr(a = 1 \mid c, x)$, the probability of being frisked given the observed covariates. As when estimating our risk model, we use gradient-boosted decision trees, fit on the first half of the NYPD dataset. Predictive performance on the second half and model checks in Appendix B indicate well-calibrated predictions across various subgroups of the data. Now recall that our sensitivity analysis is based on three key parameters: (1) $\alpha_{c,x}$, the effect of the confounder on being frisked; (2) $\delta_{c,x}$, the

¹²We 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 which 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.

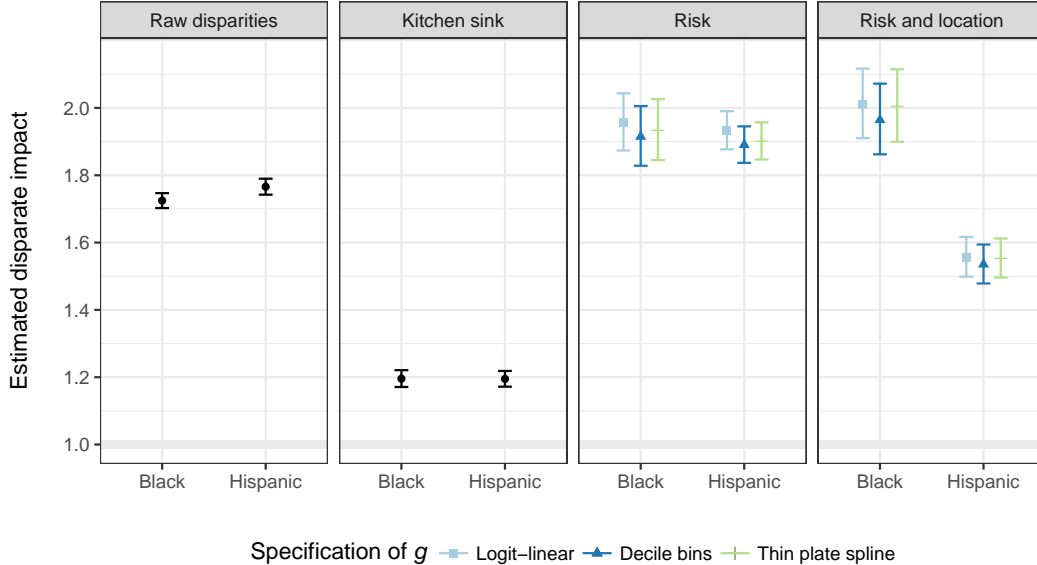


Figure 3: *Different estimates of racially disparate impacts of frisk decisions, where the y-axis is the odds of being frisked for each race group relative to white individuals. The left-most panel shows the raw disparities in frisk rates. As a measure of discrimination, raw disparities suffer from omitted-variable bias: there may be legitimate reasons why black and Hispanic pedestrians are more likely to be frisked. The second panel shows the estimated race effects in a kitchen-sink regression, controlling for all pre-frisk covariates. These estimates suffer from included-variable bias because they control for features that are correlated with race but unrelated to risk. The right two panels show the results of our risk-adjusted regression, controlling for risk only, and for risk and location. The measured disparities are robust to different specifications of the risk transformation g , and 95% bootstrap confidence intervals are shown for each estimate.*

effect of the confounder on recovering a weapon if a frisk is carried out; and (3) $\Pr(u = 1 \mid c, x)$, the prevalence of the confounder. In principle, these parameters could vary arbitrarily with both c and x . However, for simplicity, and following standard practice [Jung et al., 2017, Rosenbaum and Rubin, 1983], we limit the dependence of these parameters on the covariates. We specifically assume that $\alpha_{c,x} = \alpha$ and $\delta_{c,x} = \delta$ are constant across all individuals. We further assume that the prevalence of the confounder $\Pr(u = 1 \mid c, x) = q_c$ depends only on an individual’s race. This means that while the confounder may be present at different rates for different groups, its effects on risk and frisk rates are the same for all those individuals for whom it is present.¹³

To assess the sensitivity of estimates, we now posit ranges for the parameters above, and then compute the minimum and maximum values of θ_c —our disparate impact estimate—across this range via grid search. In constructing these ranges, we assume that the effect parameters (α and δ) are non-negative, which ensures that decision makers respond to the confounder rationally—that is, we assume a confounder cannot increase an individual’s risk while also *decreasing* that individual’s probability of being frisked. In particular, we assume $\alpha, \delta \in [0, \Theta]$ (for a parameter Θ that we select below), and we assume that $q_c \in [0, 1]$ is unconstrained. What remains is to determine an appropriate upper bound Θ on the effect sizes. There are several reasonable approaches outlined in the sensitivity literature. In some problems, there are “known unknowns,” variables that are not present in the data but whose effects can be estimated from theory or from additional investigation. Another approach is to fit a linear model estimating $\Pr(r(1) \mid c, x, a = 1)$, and assume that

¹³In preliminary numerical experiments, we found that our estimates did not change substantially if we allowed α and δ to vary by group. Thus, for computational efficiency, we assume constant values for these parameters.

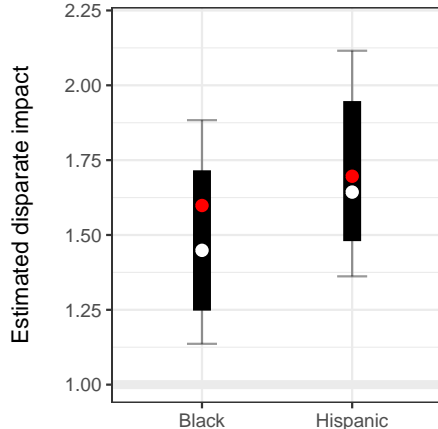


Figure 4: *Grounding the sensitivity parameters using synthetic data.* The red points show the true disparate impact, as computed from the synthetic population of 10^7 cases, while the white points are the estimates from the censored data of 100,000 samples assuming no confounding. The black bands show the plausible range of estimates under our model of confounding with $\Theta = \log(3)$. The error bars show the upper and lower halves of the 95% confidence intervals for the maximum and minimum estimates in the sensitivity analysis, computed via bootstrapping. The imputed intervals are sufficient to cover the true values.

any confounder will have a marginal effect no bigger than the largest effect associated with a binary feature in that model. We take a third approach, which is loosely related to cross-validation, as described next.

We start with the synthetic population produced in Section 4.1, where the true racial disparities are known. Recall that both risk and frisk decisions in the synthetic population are a function of observed covariates, and thus, by construction, ignorability holds in $\bar{\Omega}$. Now, to simulate confounding, we censor a set of pre-frisk covariates in $\bar{\Omega}$. Specifically, we remove variables listed in two sections of the UF-250 stop forms that describe the “circumstances” prompting the encounter. These sections consist of 20 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. By censoring these variables, we are simulating severe unobserved confounding.

We now run the sensitivity analysis described above on the censored dataset with $\Theta = \log(3)$, corresponding to u tripling the odds of frisking and of finding weapons on stopped individuals. By assuming this degree of confounding, the inferred sensitivity bands around our estimates easily cover the true risk-adjusted disparities in the synthetic dataset $\bar{\Omega}$, as shown in Figure 4.¹⁴ Accordingly, it appears that $\log(3)$ is a reasonable upper bound on the effect of confounders in the real data. It might seem surprising that the censored estimates (in white) are so close to the true values (in red), even with 20 key variables removed. But we note that for a confounder to substantially change our estimates, it has to be relatively prevalent in the data, and also correlated with risk, frisk rate, and race. Some of the censored variables (such as “furtive movements”) have little relation to risk, others (such as “suspicious object”) are predictive but rare, and most are only moderately correlated with race. As a result, censoring has only a limited effect on our disparity estimates.

Finally, we compute sensitivity bands for our estimates of risk-adjusted disparities in the real NYPD stop-and-frisk data, assuming $\Theta = \log(3)$. After accounting for such possible confounding, Figure 5 shows that black and Hispanic pedestrians were still more likely to be frisked than whites who were deemed equally

¹⁴These results are based on a logit-linear risk transformation g . As shown in Figure 3, our results are largely robust to the precise transformation applied, and so for simplicity and speed we default to a logit-linear model. As before, confidence intervals are computed via bootstrapping. Specifically, we compute the bootstrap standard error of estimates assuming no confounding, as in Figure 1, and then add and subtract twice that computed quantity to the endpoints of the sensitivity bands.




Figure 5: *Sensitivity of the risk-adjusted disparities in frisk decisions to unmeasured confounding. The black bands show the plausible range of estimates under our model of confounding (with $\Theta = \log(3)$). The 95% confidence intervals are computed via bootstrapping. Even under substantial confounding, we find unjustified disparities affecting black and Hispanic individuals.*

likely to be carrying a weapon. The figure also shows that this finding holds if we allow frisk policies to vary by location.

5 Discussion

With our risk-adjusted regression, we have sought to develop a simple, intuitive test that addresses the most serious concerns of omitted- and included-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 control for all available covariates—in line with common legal and statistical convention—can substantially skew estimates of discrimination.

Throughout our analysis, we have estimated “disparate impact” by a regression coefficient on protected-group identity in a model that controls for estimated risk. This procedure is consistent with current practice, 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. (3.1), 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 6, where we plot frisk rates as a function of risk, disaggregated by race, estimated by separate logistic regression curves. Stopped black and Hispanic pedestrians are frisked more often than stopped whites at every level of risk. As a result, one would find that 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 blacks to Hispanics, the direction of the disparity depends on risk.¹⁵ Low-risk Hispanics are frisked more often than low-risk blacks, but high-risk blacks are frisked more often than their high-risk Hispanic counterparts. Consequently, a conclusion of disparate impact between blacks and Hispanics would depend heavily on the precise definition applied. The analysis is further complicated if the risk distributions differ substantially between groups. If, hypothetically, Hispanics were mostly low-risk and blacks 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.¹⁶

¹⁵Such a comparison between minority groups is unusual in disparate impact cases, but it illustrates the underlying theoretical issue.

¹⁶Some scholars have similarly investigated interactions between race and other decision-making criteria. For example, Espenshade et al. [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




Figure 6: *Risk vs. frisk rates, as estimated via logistic regression curves fit separately for each race group, with risk plotted on a log scale. Low-risk Hispanics are more likely to be frisked than equally low-risk blacks, while high-risk blacks are frisked more often than similarly high-risk Hispanics, complicating definitions of disparate impact.*

The crossing of risk curves that we see in Figure 6 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 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 6 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 minorities—while recovering the same number of weapons [Goel et al., 2016a]. A more efficient frisk strategy could thus reduce the burdens of policing on 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 and Goel, 2018, Goel et al., 2018, 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. Though there are still important unresolved questions, we believe that our general statistical approach provides practitioners with a tractable way to assess disparities in many domains while avoiding some of the pitfalls with traditional methods. Looking forward, we hope this work spurs further theoretical and empirical research at the intersection of statistics, law, and public policy.

analysis, however, found no score ranges where minority applicants faced an absolute disadvantage relative to whites 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 blacks and Hispanics in Figure 6.

References

- Joshua D Angrist and Jörn-Steffen Pischke. *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press, 2008.
- Kenneth Arrow et al. The theory of discrimination. *Discrimination in labor markets*, 3(10):3–33, 1973.
- Ian Ayres. Outcome tests of racial disparities in police practices. *Justice Research and Policy*, 4(1-2):131–142, 2002.
- Ian Ayres. Three tests for measuring unjustified disparate impacts in organ transplantation: The problem of “included variable” bias. *Perspectives in biology and medicine*, 48(1):68–S87, 2005.
- Ian Ayres. Testing for discrimination and the problem of “included variable bias”. *Working paper*, 2010. Available at <http://islandia.law.yale.edu/ayers/ayresincludedvariablebias.pdf>.
- Ana I Balsa, Thomas G McGuire, and Lisa S Meredith. Testing for statistical discrimination in health care. *Health Services Research*, 40(1):227–252, 2005.
- Gary S Becker. Nobel lecture: The economic way of looking at behavior. *Journal of Political Economy*, pages 385–409, 1993.
- James A Berkovec, Glenn B Canner, Stuart A Gabriel, and Timothy H Hannan. Mortgage discrimination and FHA loan performance. *Cityscape*, pages 9–24, 1996.
- Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*, pages 134–148, 2018.
- Sam Corbett-Davies and Sharad Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*, 2018.
- Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 797–806. ACM, 2017.
- Thomas J Espenshade, Chang Y Chung, and Joan L Walling. Admission preferences for minority students, athletes, and legacies at elite universities. *Social Science Quarterly*, 85(5):1422–1446, 2004.
- Jeffrey Fagan. Report of jeffrey fagan in the case of floyd v. the city of new york, 2010. Available at <https://www.law.columbia.edu/sites/default/files/microsites/policing-litigation-conference/files/Fagan%20Report%20with%20Technical%20Appendices.pdf>.
- Fisher v. University of Texas. 579 U.S., 2016.
- Floyd v. City of New York, 2013. 959 F. Supp. 2d 540, S.D.N.Y.
- Andrew Gelman, Jeffrey Fagan, and Alex Kiss. 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(479):813–823, 2007.
- Sharad Goel, Justin M Rao, and Ravi Shroff. Personalized risk assessments in the criminal justice system. *The American Economic Review*, 106(5):119–123, 2016a.
- Sharad Goel, Justin M Rao, and Ravi Shroff. Precinct or prejudice? Understanding racial disparities in New York City’s stop-and-frisk policy. *Annals of Applied Statistics*, 2016b.
- Sharad Goel, Maya Perelman, Ravi Shroff, and David Sklansky. Combatting police discrimination in the age of big data. *New Criminal Law Review*, 20, 2017.

- Sharad Goel, Ravi Shroff, Jennifer L Skeem, and Christopher Slobogin. The accuracy, equity, and jurisprudence of criminal risk assessment, 2018.
- Brandon Greenwell, Bradley Boehmke, Jay Cunningham, and GBM Developers. *gbm: Generalized Boosted Regression Models*, 2019. URL <https://CRAN.R-project.org/package=gbm>. R package version 2.1.5.
- D James Greiner and Donald B Rubin. Causal effects of perceived immutable characteristics. *Review of Economics and Statistics*, 93(3):775–785, 2011.
- Griggs v. Duke Power Co. 401 U.S. 424, 1971.
- Jongbin Jung, Connor Concannon, Ravi Shroff, Sharad Goel, and Daniel G. Goldstein. Simple rules for complex decisions. *arXiv preprint arXiv:1702.04690*, 2017.
- Jongbin Jung, Ravi Shroff, Avi Feller, and Sharad Goel. Algorithmic decision making in the presence of unmeasured confounding. *arXiv preprint arXiv:1805.01868*, 2018.
- John Knowles, Nicola Persico, and Petra Todd. Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy*, 109(1), 2001.
- John Monahan and Jennifer L Skeem. Risk assessment in criminal sentencing. *Annual review of clinical psychology*, 12:489–513, 2016.
- Leslie E. Papke and Jeffrey M. Wooldridge. Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics*, 11(6):619–632, 1996. doi: 10.1002/(SICI)1099-1255(199611)11:6<619::AID-JAE418>3.0.CO;2-1.
- Emma Pierson, Sam Corbett-Davies, and Sharad Goel. Fast threshold tests for detecting discrimination. In *The 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018a.
- Emma Pierson, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Vignesh Ramachandran, Cheryl Phillips, and Sharad Goel. A large-scale analysis of racial disparities in police stops across the United States. *arXiv preprint arXiv:1706.05678*, 2018b.
- Solomon W Polachek et al. Earnings over the life cycle: The mincer earnings function and its applications. *Foundations and Trends® in Microeconomics*, 4(3):165–272, 2008.
- Greg Ridgeway and John MacDonald. Methods for assessing racially biased policing. *Race, ethnicity, and policing: New and essential readings*, pages 180–204, 2010.
- Paul R Rosenbaum and Donald B Rubin. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 212–218, 1983.
- Ravi Shroff. Predictive analytics for city agencies: Lessons from children’s services. *Big Data*, 5(3):189–196, 2017.
- Camelia Simoiu, Sam Corbett-Davies, Sharad Goel, et al. The problem of infra-marginality in outcome tests for discrimination. *The Annals of Applied Statistics*, 11(3):1193–1216, 2017.
- Jennifer Skeem, John Monahan, and Christopher Lowenkamp. Gender, risk assessment, and sanctioning: The cost of treating women like men. *Law and human behavior*, 40(5):580, 2016.
- Scott Smart and Joel Waldfogel. A citation-based test for discrimination at economics and finance journals. Technical report, National Bureau of Economic Research, 1996.
- Tyler J VanderWeele and Whitney R Robinson. Confounding and mediating variables. *Epidemiology*, 25(4): 473–484, 2014.

Washington v. Davis. 426 U.S. 229, 1976.

Simon N Wood. Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1):95–114, 2003.

Appendix A Efficient computation of sensitivity estimates

Using the setup of Section 3.2, we show that the limiting value of $\hat{\theta}^n$ can be efficiently computed via a weighted fractional-response regression. To do so, we first compute the limit of the log-likelihood function of the expression in Eq. (3.15). We start with some notation. Let

$$p(c, x, u; \boldsymbol{\theta}) = \text{logit}^{-1}(\theta_c + g_{\boldsymbol{\theta}}(\mu) + \boldsymbol{\theta}_{\ell}^{\top} x_{\ell}),$$

where $\mu = \text{logit}^{-1}(\beta_{c,x} + u\delta_{c,x})$ and $\boldsymbol{\theta}$ denotes the full vector of parameters on the right-hand side. Then, define the random variables,

$$\begin{aligned} n_{c,x} &= \sum_{i=1}^n \mathbf{1}\{\tilde{c}_i = c \wedge \tilde{x}_i = x\}, \\ n_{c,x,u} &= \sum_{i=1}^n \mathbf{1}\{\tilde{c}_i = c \wedge \tilde{x}_i = x \wedge \tilde{u}_i = u\}, \text{ and} \\ a_{c,x,u} &= \sum_{i=1}^n \tilde{a}_i \mathbf{1}\{\tilde{c}_i = c \wedge \tilde{x}_i = x \wedge \tilde{u}_i = u\}. \end{aligned}$$

Now, we can write the (normalized) log-likelihood of $\boldsymbol{\theta}$:

$$\begin{aligned} \ell_n(\boldsymbol{\theta}) &= \frac{1}{n} \sum_{i=1}^n [\tilde{a}_i \log p(\tilde{c}_i, \tilde{x}_i, \tilde{u}_i; \boldsymbol{\theta}) + (1 - \tilde{a}_i) \log(1 - p(\tilde{c}_i, \tilde{x}_i, \tilde{u}_i; \boldsymbol{\theta}))] \\ &= \frac{1}{n} \sum_{c,x,u} [a_{c,x,u} \log p(c, x, u; \boldsymbol{\theta}) + (n_{c,x,u} - a_{c,x,u}) \log(1 - p(c, x, u; \boldsymbol{\theta}))] \\ &= \sum_{c,x} \frac{n_{c,x}}{n} \sum_u \frac{n_{c,x,u}}{n_{c,x}} \left[\frac{a_{c,x,u}}{n_{c,x,u}} \log p(c, x, u; \boldsymbol{\theta}) + \frac{n_{c,x,u} - a_{c,x,u}}{n_{c,x,u}} \log(1 - p(c, x, u; \boldsymbol{\theta})) \right]. \end{aligned}$$

In the limit, we have:

$$\lim_{n \rightarrow \infty} \ell_n(\boldsymbol{\theta}) = \sum_{c,x} \Pr(c, x) \sum_u \Pr(u | c, x) [\log p(c, x, u; \boldsymbol{\theta}) \Pr(a = 1 | c, x, u) + \log(1 - p(c, x, u; \boldsymbol{\theta})) \Pr(a = 0 | c, x, u)].$$

Further, letting N denote the size of the original dataset Ω , we have that the limit of $\ell_n(\boldsymbol{\theta})$ is equivalent to,

$$\ell(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N \sum_u \Pr(u | c_i, x_i) [\log p(c_i, x_i, u; \boldsymbol{\theta}) \Pr(a_i = 1 | c_i, x_i, u) + \log(1 - p(c_i, x_i, u; \boldsymbol{\theta})) \Pr(a_i = 0 | c_i, x_i, u)].$$

Finally, note that the expression above is identical to the log-likelihood that is optimized when fitting a weighted fractional-response regression to an expanded dataset Ω' comprised of two copies of the observed data Ω —one with u set to 0 and another with u set to 1—with fractional response $\Pr(a_i = 1 | c_i, x_i, u)$ and weights $\Pr(u | c_i, x_i)$. As a result, one can infer the limiting value of $\hat{\theta}^n$ by simply fitting this fractional-response regression on Ω' .

Appendix B Model checks

The outcome model from Section 4.2 for estimating $\Pr(r(1) | c, x, a = 1)$, the probability a weapon is found conditional on all observable pre-frisk variables, achieves an AUC of 81% on the second half of the data. The decision model from Section 4.3 for estimating $\Pr(a = 1 | c, x)$, the probability of being frisked given the observed covariates, achieves an AUC of 83% on the second half of the data. Figure B1 shows calibration between model-predicted rates and empirical rates by various subgroups and locations.

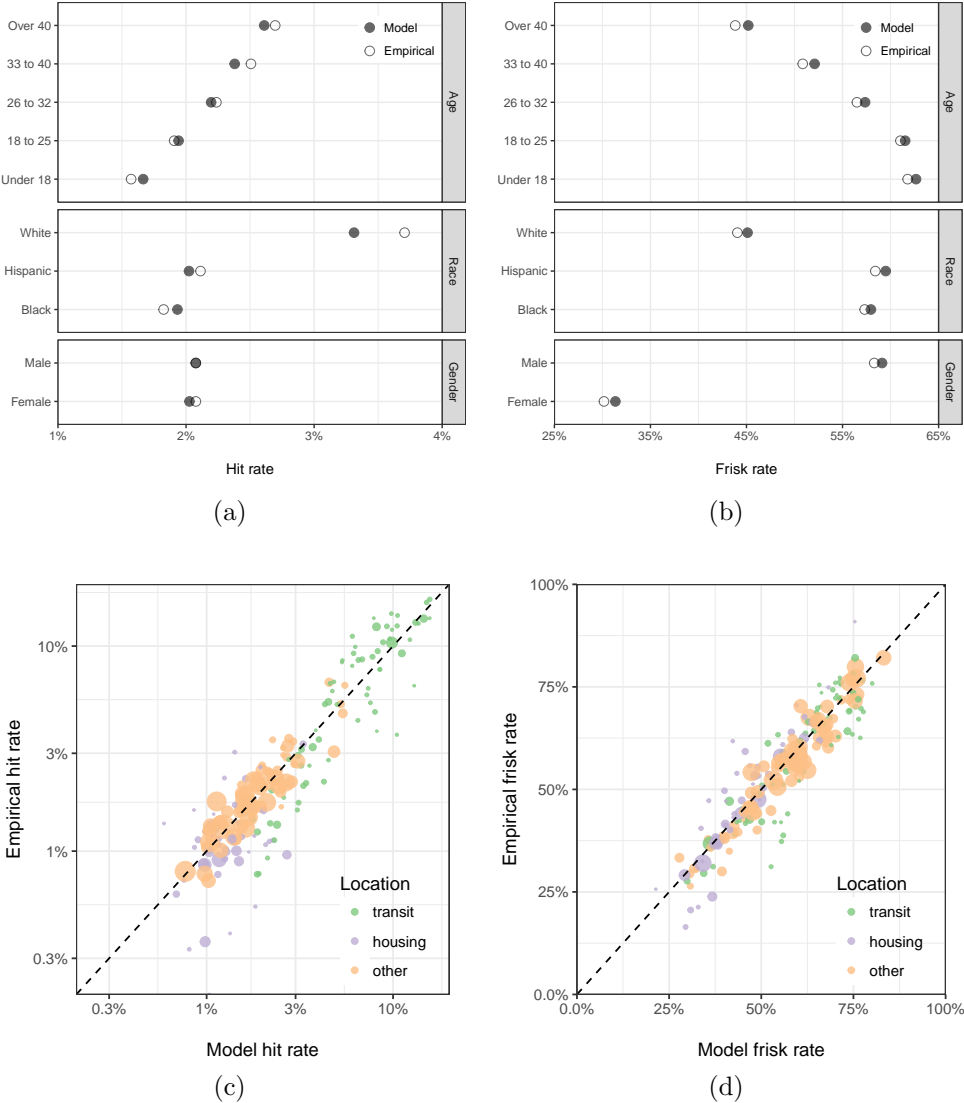


Figure B1: Comparison of model-predicted versus empirical weapon recovery rate (“hit rate”) (left column, used to estimate risk-adjusted disparities) and frisk rate (right column, used to assess the sensitivity of estimates). Panels (a) and (b) show that the model-predicted hit rates and frisk rates are close to their empirical counterparts, conditional on values of age, race, and gender. In panels (c) and (d), 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 and decision models predict well over the full range of hit rates and frisk rates, respectively.