Judicial Bias Against Minority and Female Attorneys

Priyanka Goonetilleke*†

Northwestern University

May 26, 2025

Abstract

I study judicial bias against attorneys. To do so, I exploit the random assignment of judges and public defenders to bail cases in Miami-Dade County. I find that the clients of Black attorneys are 5.1 percentage points less likely to be released on bail if the case is heard by the judge least favorable to Black attorneys rather than the most favorable judge. However, I find no evidence that judges are biased in the treatment of Hispanic or female attorneys. The findings suggest that racial bias exists in courtrooms beyond that against criminal defendants.

*Marti Family Fellowship Visiting Assistant Professor of Law, Northwestern University Pritzker School of Law. email: priyanka.goonetilleke@law.northwestern.edu. An earlier version of this paper served as the first chapter of my dissertation submitted in partial fulfillment of the requirements for the PhD in Economics.

†I am deeply indebted to David Abrams, Juan Pablo Atal, Hanming Fang, and Aurélie Ouss for their continued support and mentorship. Rupalee Rashatwar and Chad Sinkler provided valuable insight into the operation of the Miami-Dade Public Defender's Office. Special thanks also to Francesco Agostinelli, Aislinn Bohren, Stephanie Holmes Didwania, Guy Grossman, Marit Rehavi, Petra Todd, Conor Walsh, João Ritto, Ornella Darova, Connor Gregory, Jorge Hansen, Cung Truong Hoang, Kathleen Hui, Artem Kuriksha, Andrew Lilley, Alfonso Maselli, and seminar participants at the University of Pennsylvania Empirical Microeconomics Seminar and the Washington University Economics Graduate Student Conference. George Zhang provided excellent research assistance. This research has been supported by the Horowitz Foundation for Social Policy, the University of Pennsylvania SASGov, and the Center for the Study of Ethnicity, Race, and Immigration (CSERI) at the University of Pennsylvania. All mistakes are my own.

1 Introduction

A significant body of research has analyzed the presence and effects of judicial bias. Judges' biases against criminal defendants affect the conditions of bail (Arnold et al. 2018), whether a defendant is found guilty (Abrams et al. 2012), and the length of prison sentences (Fischman & Schanzenbach 2012). Given the prevalence of documented bias against criminal defendants, one may expect judges to be similarly biased against the attorneys who litigate before them. If judges are biased against attorneys, it could impact the attorneys' careers and partially explain the lower diversity in senior positions in the legal profession. Moreover, criminal defendants are more likely to select attorneys of their demographic group (Agan et al. 2021), so judicial bias against attorneys may also have a disparate impact on minority defendants.

In this article, I study whether judges are biased in the treatment of attorneys. To do so, I exploit a setting in Miami-Dade County where all criminal defendants arrested for a felony are assigned a public defender. Specifically, I use case data from 2006-2017, which includes 170 public defenders and 169 judges. To measure the impact of judicial bias, I examine whether a public defender secured pretrial release for their client. To identify the causal effect of bias, I exploit the fact that criminal defendants are randomly assigned both a public defender and a judge at their bail hearing. The random assignment of cases to public defenders removes the concern of case selection, and the random assignment of cases to judges means that the distribution of unobserved defendant characteristics should be the same across judges.

My primary specification identifies bias based upon variation across judges in the gap between the release rate of the clients of public defenders from specific demographic groups—Black, Hispanic, and female—and the release rate of the clients of White male public defenders. While differences in unobservables between attorney groups may justify a racial or gender gap in release rates across attorney demographic groups, these gaps should be stable across judges absent bias. Hence, testing for differences in these gaps across judges allows me to estimate whether judicial bias in the treatment of different attorney demographic groups exists. In my primary specification, statistical inference is based on comparing the distribution of judge-level estimates to a simulated distribution of judge-level estimates under the null hypothesis of no bias.

Employing this design, I find statistically significant variation in the Black vs. White racial gap in public defender release rates across judges. For example, I find that Black attorneys' clients are 5.1 percentage points more likely to be released on bail if the case is heard by the judge most favorable to Black attorneys, compared to if the least favorable judge hears it; 9.1% of the mean release rate of 56% for Black attorneys. However, I find no evidence of bias in the treatment of either Hispanic or female attorneys.

One concern with these estimates is that they may be driven by differences in the behavior of Black attorneys when arguing before certain judges. Given that attorneys in my setting have limited experience with the judges they are appearing before in these bail hearings, behavioral changes would likely be in response to observable characteristics of the judges. Hence, to explore this concern, I rerun my analysis examining variation only across White male judges, thereby reducing the likelihood that variation is driven by Black attorneys adjusting their behavior based on judge identity. The variation between this restricted set of judges is also statistically significant, providing suggestive evidence that behavioral differences alone are not driving the results.

In a final analysis, I quantify the potential effect of the observed bias in the treatment of Black attorneys on racial disparities in defendant outcomes. This analysis is motivated by evidence that defendants are more likely to hire an attorney of their demographic group (Agan et al. 2021).To quantify the downstream effects in settings where defendants choose their representation, I combine my estimates with estimates of client-attorney matching patterns from the setting of Agan et al. (2021). The results suggest that judicial bias against Black attorneys could contribute to racial disparities among defendants. For instance, if all judges were replaced with the judge most favorable to Black attorneys, the Black-White gap in pretrial detention would reverse: from Black defendants being detained at a rate 3.1 percentage points higher than White defendants to 0.6 percentage points lower.

This article contributes to the extensive literature on judicial bias (See, e.g., Abrams et al. 2012, Albright 2019, Fischman & Schanzenbach 2012, Grossman et al. 2016, Harris & Sen 2019, Rachlinski et al. 2008, Shayo & Zussman 2011). The fact that these biases have an impact on case outcomes has been shown in numerous studies (e.g., Alesina & La Ferrara 2014, Arnold et al. 2022, Didwania 2020*a*, Fishman et al. 2006, Sen 2015, amongst others). However, only a handful of prior papers study bias against attorneys. In an experiment, Hodgson & Pryor (1984) studied gender bias by hiring actors to play the role of attorneys, with participants assessing the defendant's guilt and the attorney's credibility. Several non-causal papers have looked at this question using real-world case outcomes. Szmer et al. (2010) and Szmer et al. (2013) investigate the relationship between attorney gender and appellate court outcomes, and Chen et al. (2017) examines the effect of attorney vocal characteristics on Supreme Court outcomes. This paper is the first to causally identify judicial bias in the treatment of attorneys in a real-world setting with case outcomes. This paper is also the first to exploit double random assignment in the criminal justice setting.[1]

The remainder of the paper is organized as follows. Section 2 provides an overview of the Miami-Dade pretrial system and describes the data. Section 3 outlines the identification and presents the empirical tests of covariate balance. Section 4 provides the main results of my test for judicial bias, and Section 5 estimates the potential impact on defendant outcomes and tests of robustness. Section 6 concludes the paper.

[1]For papers using random assignment of judges to cases to test for judicial bias against defendants based on demographics, see, e.g., Abrams et al. (2012), Arnold et al. (2018), Grossman et al. (2016), Kastellec (2021), Shayo & Zussman (2011). The random assignment of cases to prosecutors has been leveraged by Sloan (2019) to test for prosecutorial bias. Meanwhile, the random assignment of defense attorneys has been utilized by Mikdash & Oh (2024).

2 Empirical Setting and Data

2.1 Empirical Setting

I apply my test for bias against attorneys in the context of first appearance hearings in Miami-Dade County, Florida. The first appearance is the first stage of criminal court proceedings in Florida. This hearing must occur within 24 hours of the defendant's arrest if they have not already obtained release pretrial. While most defendants are eligible for immediate release upon paying a scheduled bail amount, many fail to pay this and will have a first appearance hearing (Goldkamp & Gottfredson 1988). The hearings are short, lasting only a few minutes per defendant. An assistant state attorney represents the state in these proceedings. As the hearing happens shortly after the arrest, all defendants are represented by a public defender, regardless of indigency status. The judge can decide to raise or lower the predetermined bail amount. They can also impose non-monetary conditions such as electronic monitoring. If the defendant cannot pay the set bail, they must stay in jail until the conclusion of their case. The judge directly chooses a bail amount and should set the least restrictive measures possible to prevent pretrial misconduct. However, in the literature, it is generally considered that the judge is, in effect, deciding whether or not to release the defendant (e.g., Arnold et al. (2018) and Kleinberg et al. (2018)). Hence, pretrial release will be defined as the measure of a public defender's "success" in my analysis.

The first appearance hearing serves as an ideal setting for several reasons. First, the hearing occurs at the start of a defendant's legal process. As a result, no other attorneys or judges have been involved in the case when this decision is made, thus reducing potential confounding factors. Second, the short nature of the hearing means there is more scope for heuristics to play a role in the outcome.[2] In addition, while determining guilt is the paramount decision in a criminal case, the prevalence of plea bargains in the U.S. means that often the judge does not play an early and active role in determining guilt. Meanwhile,

[2]Bertrand et al. (2005) amongst others, find that implicit bias has a greater impact in settings where the decision-maker is rushed.

first appearance hearings occur in a large proportion of cases, and despite their short nature, they substantially impact defendant outcomes. Dobbie et al. (2018) show that for marginal defendants, pretrial release decreases the probability of conviction by 14 percentage points and increases employment prospects 3-4 years after conviction by 5.1 percentage points.[3]

My identification strategy relies upon random assignment of cases to judges and public defenders. On weekends, trial judges cover the bail shift on a rotating basis. Each Saturday, Sunday, and federal holiday, a single judge will work the felony first appearance shift, hearing all felony first appearances.[4] The public defender's office would separately assign a public defender to handle the first appearance hearings on these days. Importantly, the public defender is assigned without knowing who the judge for the relevant shift will be. In addition, as the hearing must occur within 24 hours of arrest, the assigned judge and public defender handle the hearings of all defendants scheduled for that day. I explain my identification strategy in more detail in Section 3.

2.2 Data

I use administrative data from all felony pretrial hearings from the Miami-Dade County Clerk of Courts between 2006 and 2017.[5] This consists of $286,975$ distinct hearings. I restrict the data in several ways. First, I restrict the period of my estimation sample to cases from January 2008 onwards so that I can calculate pretrial misconduct in the two years before the defendant's first appearance hearing. Next, I restrict the data to cases where the first appearance hearing occurred on a weekend or federal holiday because the quasi-random assignment of both bail judges and public defenders only applies on those

[3]Gupta et al. (2016), Heaton et al. (2017), Leslie & Pope (2017), Stevenson (2018) similarly find that pretrial detention increases the probability a defendant is convicted in state courts. However, Didwania (2020*b*) does not find an effect on conviction probability for federal felony defendants.

[4]Meanwhile, on weekdays, felony first appearance hearings are handled by a specialized bail judge and a dedicated team of public defenders.

[5]I end the sample in 2017 because of a change in the Miami-Dade Public Defenders Office process for handling pretrial release. In 2017, the Public Defenders Office established an Early Representation Unit, which would attempt to secure pretrial release for a defendant after the first appearance hearing. This would confound the analysis as a defendant may thus obtain release other than at the first appearance hearing.

days. Finally, I drop 2,567 hearings with public defenders who are not "White," "Black," or "Hispanic."

Following these exclusions, $42,202$ cases remain. These cases involve 170 unique public defenders and 169 unique judges. The data contains information about the defendant (race/ethnicity, gender, date of birth, and residential zip code) and the case (charge, offense type, case disposition, the identity of the judge and attorneys, and whether the defendant was released pretrial). To code for the gender and race of the judges and attorneys, I match their names to self-reported race and gender from the Florida Electoral Rolls for Miami-Dade County. For the 20% of attorneys and 50% judges who could not be matched to the electoral rolls, I used machine learning tools to predict race, ethnicity, and gender from their names.

Although machine learning tools are commonly used and generally reliable for imputing gender and ethnicity, they are less accurate for race. To validate the race classifications of attorneys and judges, a research assistant manually searched public sources such as LinkedIn and law firm websites. For 5% of public defenders and 3% of judges, manual verification was not possible, and the machine learning classification was relied upon. Conversely, race was determined based only on the manual searches for 5% of attorneys and 8% of judges where the algorithm was inconclusive or incorrect.

Finally, I use the defendant's home 5-digit zip code to control for a proxy for income (using Food Stamps/SNAP eligibility), local unemployment rates, and education from the 2017 American Community Survey.[6]

This paper defines " White " as non-Hispanic White, "Hispanic" as Hispanic White, and "Black" as all Black individuals regardless of ethnicity. I separate Black and White Hispanic individuals to allow for potential differences in their treatment, which may be especially important in my context, given the unique position of the Cuban-American population

[6]For the 12.7% of observations where the defendant's zipcode is missing, multiple imputation via the "mice" package in R is used to impute these control variables. For this imputation, I use the race and sex of the defendant, the crime type committed, and the date of the first appearance.

in Florida.[7]

2.3 Descriptive Statistics

Table 1 reports summary statistics for the public defenders in my sample separated by race and gender, and Table 2 reports the analogous summary statistics for the judges.

Black attorneys make up 15% of public defenders and 8% of judges, while Hispanic attorneys make up 23% of public defenders and 29% of judges. Female attorneys make up 55% of public defenders and 50% of judges. Pretrial release rates are similar across judge and public defender race and gender, ranging from 56% to 58% of cases.

Next, I assess variation in leniency between judges. Figure 1 reports the distribution of the pretrial release rate of judges who presided over at least 30 weekend hearings in my dataset. The average release rate across judges is 57.4%. Consistent with other studies, there is substantial variation in leniency across judges in my setting. The most lenient judge releases 79% of defendants, while the strictest releases just 41%.

Finally, I assess gaps in the release rates between White male public defenders and the other demographic groups. For each judge, I calculate the proportion of defendants they release pretrial, grouped by the demographic group of the public defender. I then calculate the difference between the release rate for each demographic group of public defenders and White male public defenders. Table 3 reports summary statistics, and Figures 2a, 2b and 2c report the distribution in these gaps across individual judges. These figures show substantial variation across judges in these gaps for the three demographic groups. For example, sorting judges by their Black vs. White attorney release rate gap, the release rate of Black public defenders is 8.8 percentage points less than that of White attorneys in front of the 25$^{\text{th}}$ percentile judge. Meanwhile, it is 1.8 percentage points higher in front of the 75$^{\text{th}}$ percentile judge. Comparing the variation across groups, the standard deviation of these gaps for

[7] Eckstein (2014) cites that in 2000, Cuban Americans held one-third of the top appointed positions in Miami-Dade and that three-quarters of Miami-Dade residents perceived them as the ethnic group with the greatest political power. Ojito (2000) expresses a similar view of White Cuban males as representing a majority group within Miami-Dade. Given this, my results for the Hispanic ethnic group may not represent the experience of the broader Hispanic community.

Black and Hispanic attorneys is 10.1 and 10.3 percentage points, respectively. Meanwhile, the standard deviation for female attorneys is just 5.6 percentage points. Using a Bartlett Test of homogeneity of variances, the estimated test statistic is 13.04 with a p-value of 0.001. Thus, the difference in the standard deviation across the demographic groups is statistically significant.

Table 1: Case Characteristics By Public Defender Demographics

	Black	Hispanic	White	Male	Female
A. Sample					
Public Defenders	25	39	105	76	93
Cases	4599	9468	28113	21197	20983
B. Case Outcome					
Pretrial Release	0.56	0.58	0.57	0.57	0.57

This table provides summary statistics for the Public Defenders who conducted first appearances in Miami-Dade between January 2008 and December 2016, separated by Public Defender race and gender.

Table 2: Case Characteristics By Judge Demographics

	Black	Hispanic	White	Other Race	Male	Female
A. Sample						
Judges	14	49	105	1	84	85
Cases	3989	11471	26304	416	21094	21086
B. Case Outcome						
Pretrial Release	0.58	0.58	0.56	0.55	0.57	0.57

This table provides summary statistics for the judges who conducted first appearances in Miami-Dade between January 2008 and December 2016, broken out by judge race and gender.

Figure 1: Release Rates by Judge

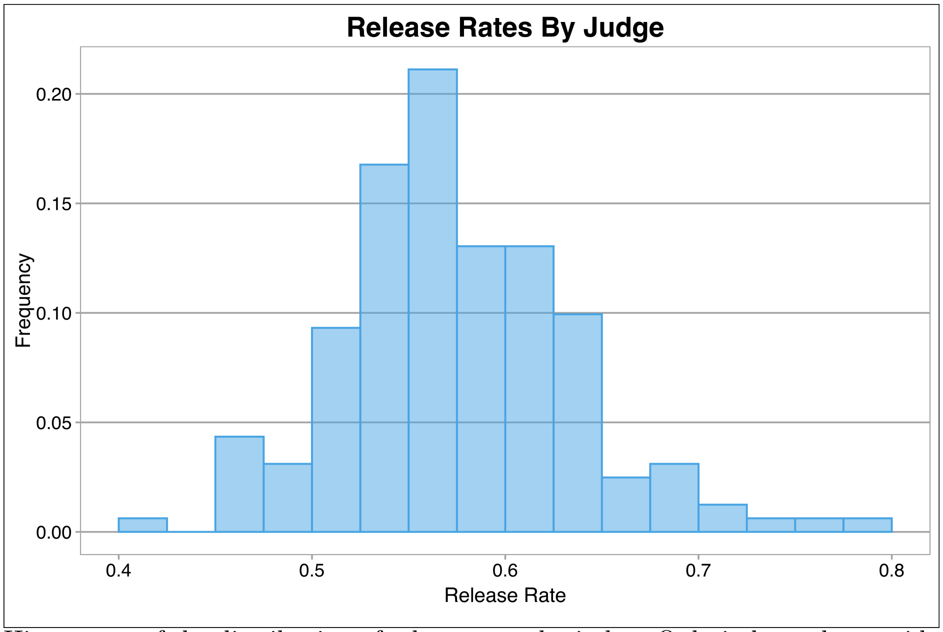

Histograms of the distribution of release rates by judge. Only judges who preside over at least 30 pretrial hearings are included.

Figure 2: Attorney Demographic Group Gaps by Judge

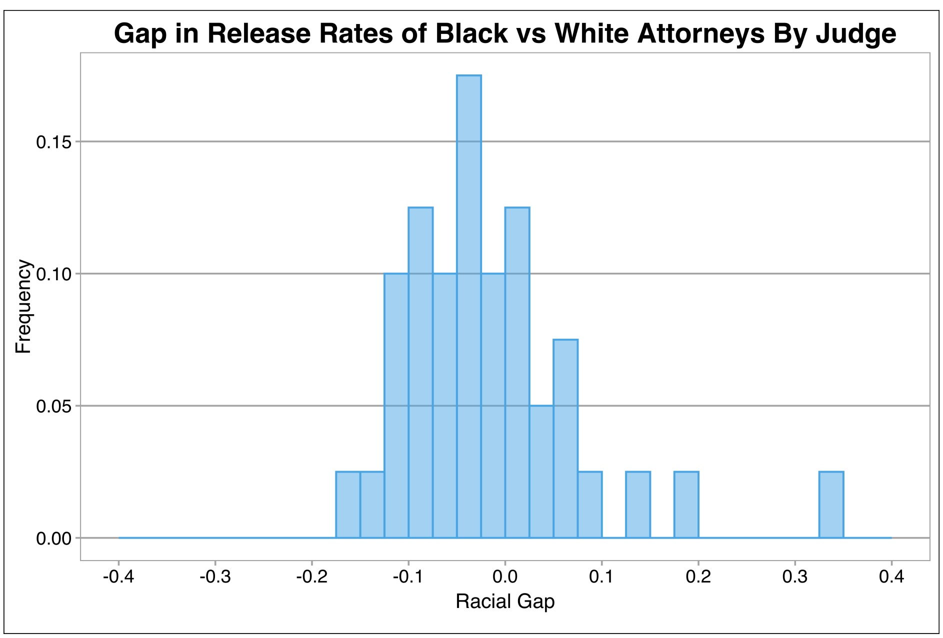

(a) Black vs. White Public Defenders

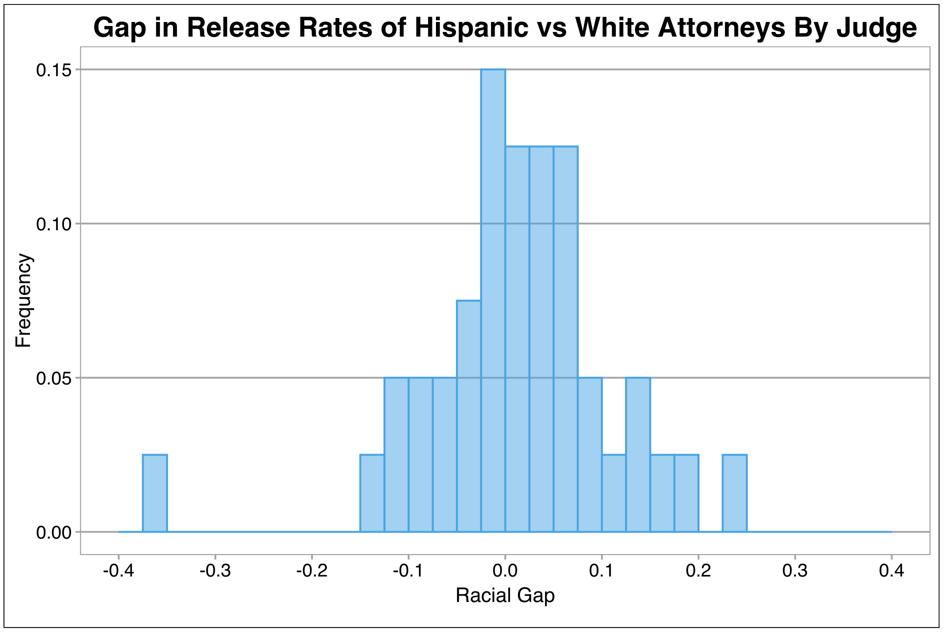

(b) Hispanic vs. White Public Defenders

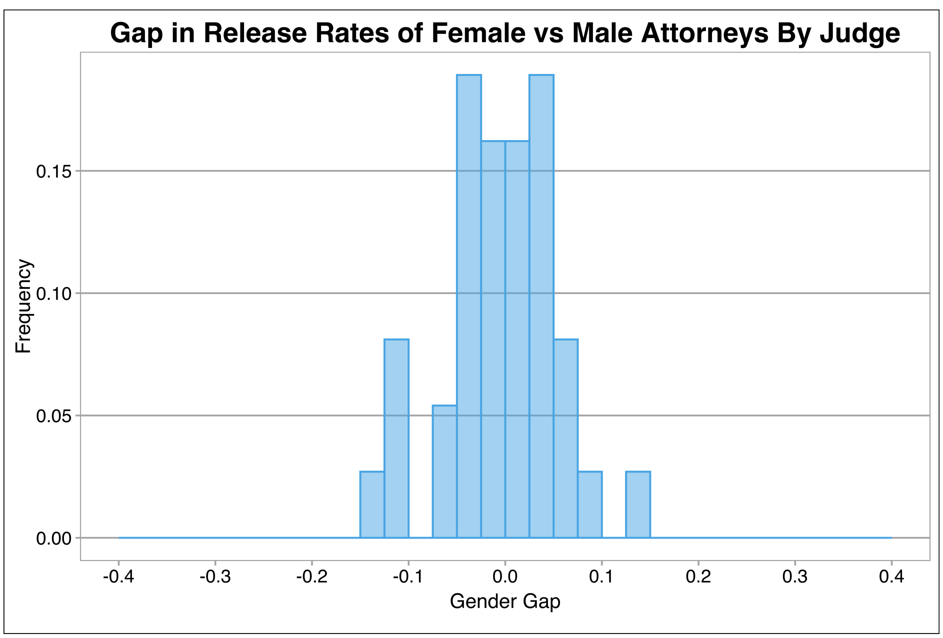

(c) Female vs. Male Public Defenders

Histograms of the distribution across judges of the gap in the rate at which the relevant demographic group secures the release of their defendants relative to the rate at which White male attorneys do. Only judges who preside over at least 30 pretrial hearings involving public defenders of each demographic group and 30 involving White male public defenders are included.

Table 3: Variation Across Judges in PD Success Rates

	Black	Hispanic	Female
Median	-0.028	0.016	-0.007
IQR	0.106	0.095	0.064
sigma	0.094	0.103	0.056
Range	0.497	0.602	0.264

This table provides summary statistics for the distribution across judges of the gap in the rate at which the relevant demographic group secures the release of their defendants relative to the rate at which White male attorneys do. Only judges who preside over at least 30 pretrial hearings involving public defenders of each marginalized demographic group and 30 involving White male public defenders are included.

3 Research Design

My research design leverages the quasi-random assignment of judges and public defenders to first appearance shifts. This assignment mechanism ensures that case characteristics are uncorrelated with either public defender demographics or judge identity. To test for differential treatment, I examine how the release rates of minority and female public defenders differ from those of White male defenders across judges. I refer to these differences as "relative bias. While some performance gaps may be attributable to attorney characteristics correlated with demographics, systematic variation in these demographic gaps across judges would suggest that some judges treat attorneys differently based on race or gender.

Note that this does not account for the potential for public defenders to change the quality of their representation depending on the judge's identity. Thus, my identification method assumes that attorneys do not change their behavior based on the judge presiding over the hearing. Section 5.3 discusses this assumption in detail.

In my main specification, I estimate judge-specific differences in the rate of pretrial release by public defender race and gender using the following probit model:

$$\begin{aligned} R_i =& \alpha + \sum_{j=1}^{J} \beta_j^B (J_{ij} \times B_i) + \sum_{j=1}^{J} \beta_j^H (J_{ij} \times H_{ik}) + \sum_{j=1}^{J} \beta_j^F (J_{ij} \times F_k) \\ & + \delta_j + \gamma_d + P_i' \upsilon + X_i \kappa + \theta_m + \eta_y + h_i + \varepsilon_i, \end{aligned} \tag{1}$$

where R_i is a dummy for whether pretrial release was granted in hearing i. J_{ij} is an indicator for whether judge j presided over the pretrial hearing. B_k (H_k, F_k) is a dummy for whether the public defender k was Black (Hispanic, female). I also include fixed effects for each judge (δ_j) and public defender (γ_k) to account for variation in leniency and ability, respectively. P_i and X_i are vectors of controls for prosecutor and defendant/case characteristics. I also include dummies for the month (θ_m) and year (η_y) of the pretrial hearing, and an indicator for whether it occurred on a holiday (h_i). The defendant/case controls include gender, race, home zip code characteristics, and criminal history. The prosecutor controls include race, gender, experience, and law school rank. In my primary specification, I include year by defendant race and ethnicity fixed effects to account for potential variation over time in the treatment of minority defendants. The results are consistent if I do not include these interactions.

I use this as my main specification rather than simply regressing defendant outcomes on interactions between public defender demographics and judge demographic groups, because offsetting bias by different judges of the same race or gender could lead to a finding of no bias for a given demographic group in such a test. Therefore, while judge race and gender are predefined groupings that facilitate comparisons across judges, a null finding when comparing between judge demographic groups is insufficient to conclude that there is no bias.

I use a probit because a linear probability model can conflate bias with differences in overall judge leniency and attorney ability. For example, an unbiased but strict judge may show larger racial gaps in release rates than an equally unbiased but more lenient judge, simply because the strict judge operates in a part of the distribution where ability differences translate into larger outcome gaps. A non-linear model accounts for this as the marginal effect of any estimated bias depends on the baseline probability of release. Within this functional form, I assume that the impact of bias is additive, as motivated by my theoretical model in Appendix B.

However, even with random assignment, judges may receive different mixes of cases along unobserved dimensions, potentially generating differences in racial and gender gaps without any behavioral differences. To assess whether the observed variation across judges could arise by chance, I use randomization inference, comparing the estimated coefficients from Equation 1 to counterfactual estimates simulated under a null hypothesis of no bias. I use Monte Carlo simulation to construct the appropriate counterfactual, randomly reassigning public defender race and gender in each iteration.[8] Because public defender demographic

[8]See Appendix A for simulation details.

characteristics are randomized, assigned public defender demographics could not have impacted outcomes. Hence, the estimated racial and gender gaps reflect sampling variation.

I use the variance of the estimated judge-level coefficients as my test statistic. By comparing this observed variance to its simulated distribution, I compute the probability that the observed heterogeneity in release rates by public defender demographics could occur by chance.

Standard asymptotic inference is inappropriate in this context due to the limited number of observations at the judge-by-public defender demographic level.[9] The conventional approach—a Likelihood Ratio Test—performs poorly in small samples and leads to an over-rejection of the null in my setting. In Appendix A, I use simulated data to show the potential for substantial over-rejection when using the Likelihood Ratio Test.

The estimated individual judge and public defender dummies are only separately identified within a "connected set" (Abowd et al. 2002). In this context, two judges are connected if an attorney has appeared before both. A connected set contains all of the judges connected to at least one other judge in the set and all of the attorneys who have appeared before any of the judges in the set. To illustrate the concept of a connected set in this context, consider an example with five judges $\{j_1, j_2, j_3, j_4, j_5\}$ and four public defenders $\{d_1, d_2, d_3, d_4\}$. Suppose d_1 has only appeared before judges j_1 and j_2, while d_2 has appeared only before judges j_2 and j_3. Meanwhile, d_3 has appeared before j_4, and d_4 has appeared before j_4 and j_5. In this example, we would have two connected sets. The largest set contains judges j_1, j_2, and j_3, as well as public defenders d_1 and d_2. Meanwhile, the second set contains judges j_4 and j_5 alongside defendants d_3 and d_4.

For each connected set, the estimation of judge and public defender fixed effects requires the omission of one public defender fixed effect, which then serves as the reference point against which all other judge and public defender fixed effects within that set are

[9] See Abrams et al. (2012) for a discussion of the potential for over-rejection of the null hypothesis when using an F-test in a similar setting.

measured.

Due to the addition of judge-by-public defender demographic fixed effects, judges must be connected by both a White male attorney and an attorney from the relevant demographic group. In addition, in each connected set it is also necessary to drop a public defender fixed effect of the relevant demographic group. Hence, two are dropped: one for a White male attorney and an additional one for an attorney from the relevant demographic group.

In determining the judges and attorneys within the largest connected set, I restrict "connections" between judges and attorneys to those with at least 30 hearings involving the relevant judge-attorney pair.[10]

The largest connected sets consist of 16 connected judges for Black attorneys, 29 for Hispanic attorneys, and 66 for female attorneys. These judges, respectively, account for 15.8%, 26.5%, and 54.9% of hearings in the sample.

The identifying assumption is that there was, in fact, quasi-random assignment of cases to judges and public defenders. I first regress defendant and case characteristics on dummies for public defender demographics to investigate whether the public defender assignment process was quasi-random. Specifically, I examine defendant characteristics (gender, race, and prior pretrial misconduct); case characteristics (the number of first-degree and life felonies, as well as whether the case involves a violent crime or property crime); and prosecutor characteristics (race and gender).

Table 4 reports the results and reveals that public defender demographics are jointly insignificant in each regression and thus do not predict defendant, case, or prosecutor characteristics. This suggests covariate balance and therefore provides evidence consistent with the identifying assumption.

[10]The 'fixest' package in R automatically calculates connected sets and omits fixed effects as required. The manual calculation of connected sets is only used to determine which judge fixed effects can be directly compared.

Table 4: Test of Covariate Balance Across Public Defenders

	F	Pr(>F)
Defendant Male	0.869	0.456
Defendant Black	1.071	0.360
Prior Misconduct	0.145	0.933
# of First Degree Felonies	0.691	0.557
# of Life Felonies	0.264	0.852
Incl Violent Crime	0.432	0.730
Incl Property Crime	1.376	0.248
Probation Violation	1.690	0.167
Male ASA	1.482	0.218
White ASA	0.112	0.953

This table reports the results of the test of covariate balance across public defender demographic groups.

Because my primary analysis compares demographic gaps in release rates across individual judges, I also investigate the extent to which each judge has a similar mix of cases. To do so, I regress the same characteristics on dummies for judge identity to assess whether certain judges systematically see different types of cases. I use Monte Carlo simulations to generate a distribution of my test statistic—the variance of judge coefficients—under random assignment. I randomly reassign judges to bail shifts for each simulation while keeping the underlying case data fixed, ensuring the null is true by construction, and re-estimate the regression for each defendant and case characteristic. I then compare the observed variance to the simulated distribution to compute a p-value. Table 5 reports the results. I fail to reject the null hypothesis of covariate balance for any characteristic at conventional significance levels. These results support that cases are quasi-randomly assigned to judges.

Table 5: Test of Covariate Balance Across Judges

	Variance	p-value
σ^2_{Black}	0.025	0.779
$\sigma^2_{Hispanic}$	0.047	0.920
σ^2_{Male}	0.080	0.569
$\sigma^2_{DefendantMale}$	0.001	0.917
$\sigma^2_{DefendantBlack}$	0.002	0.983
$\sigma^2_{LifeFelony}$	0.001	0.342
$\sigma^2_{FirstDegreeFelony}$	0.007	0.636
$\sigma^2_{PriorMisconduct}$	0.003	0.695
$\sigma^2_{Violent}$	0.002	0.986
$\sigma^2_{Property}$	0.003	0.977
$\sigma^2_{Probation}$	0.001	0.905
$\sigma^2_{ASAMale}$	0.068	0.942
$\sigma^2_{ASAWhite}$	0.057	0.996

This table reports the variance of the estimates of the judge dummies from estimating whether judge identity predicts case characteristics.This empirical variance is compared to the distribution of variances estimated from 10,000 simulated datasets constructed under the null hypothesis that there is no correlation between judge identity and defendant/case characteristics. The reported p-value is the proportion of simulated variances larger than that estimated in the true data.

4 Results

Although my primary specification uses the randomization inference described above, I begin by estimating a modified version of Equation 1 where the interactions between the judge-specific indicators and the public defender demographics are replaced with interactions between public defender demographics and judge demographics. As this specification does not rely on connected sets, I replace prosecutor characteristics with prosecutor fixed effects.

In addition, as Table 2 shows minimal variations in leniency across judge demographic groups, I estimate this as a linear probability model rather than a probit model.

$$\begin{aligned} R_i = & \alpha + \beta_1(B_i^J \cdot B_i) + \beta_2(B_i^J \cdot H_i) + \beta_3(H_i^J \cdot B_i) + \beta_4(H_i^J \cdot H_i) + \beta_5(F_i^J \cdot F_i) \\ & + \delta_j + \gamma_d + \psi_p + \kappa X_i + \theta_m + \eta_y + h_i + \varepsilon_i, \end{aligned} \tag{2}$$

where $B_i^J(H_i^J, F_i^J)$ is a dummy for whether bail in hearing i was decided by a Black (Hispanic, female) judge. Table 6 reports the results (Column 1). The coefficient on Black Judge × Black PD suggests that a Black judge is 6.1 percentage points more likely to release a defendant if represented by a Black public defender rather than a White public defender, but the estimates are not statistically significant at conventional levels. Column (2) presents p-values derived from randomization inference based on 5,000 simulations in which public defender race is randomly assigned. These p-values indicate the proportion of simulated estimates with coefficients more extreme than those reported in column (1). Overall, none of the estimates are statistically significant under either approach.

As mentioned above, offsetting bias by different judges of the same race or gender could lead to a finding of no bias for a given demographic group in the above test. My primary specification, therefore, assesses variation in the treatment of public defenders across specific judges. Figures 3a , Figure 3b, and 3c report the primary results from the randomization inference.[11]

Figure 3a displays the distribution of the variance in judge-specific treatment effects for Black public defenders (β_j^B) across simulations, with the dashed line indicating the variance from the actual data. Only 3.8% of simulations produced more extreme variation than the observed data, suggesting the variation in the treatment of Black attorneys exceeds that expected from sampling noise alone. I interpret this as evidence of relative bias: some judges

[11]See Appendix Figure A1 for the distribution of the raw probit coefficients within each connected set.

Table 6: Test of Homophily

	Dependent variable:	
	Pretrial Release	p-value
	(1)	(2)
Black Judge * Black PD	0.061	0.32
	(0.046)	
Hispanic Judge * Black PD	−0.021	0.52
	(0.023)	
Black Judge * Hispanic PD	0.015	0.74
	(0.036)	
Hispanic Judge * Hispanic PD	0.035	0.17
	(0.022)	
Female Judge * Female PD	−0.018	0.15
	(0.011)	
Mean	0.57	
Observations	42,202	
Adjusted R^2	0.143	
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$	

Column (1) reports the results of the regression in equation 2. Robust standard errors clustered at the Judge and Public Defender level.

treat Black attorneys less favorably than others.

To aid interpretation, I convert probit estimates into average marginal effects.[12] Randomization inference also provides an estimate of the variation expected from sampling noise. I use the simulation with the median variance as a benchmark and compare the observed variation in the real data to this reference point.[13] The standard deviation of marginal effects in the real data exceeds the median simulated value by 2 percentage points (a 22.7% increase). The range of marginal effects is 5.1 percentage points wider, equivalent to 9.1% of the average release rate for Black attorneys (56%).[14] For context, Arnold et al. (2022) find a 3.1 percentage point standard deviation in judge-level bias against Black defendants in NYC's bail system.

In contrast, for Hispanic attorneys (Figure 3b), 26.9% of simulations show greater variance than the actual data, and for female attorneys (Figure 3c), 83.2% of simulations had greater variance. I cannot reject the null hypothesis of equal treatment across judges for either group. Given that statistically significant variation appears only in the treatment of Black attorneys, I focus the remainder of the analysis on this group.

---

[12]Note that as discussed in Ai & Norton (2003), this is not equivalent to the "interaction effect." However, under the theoretical model outlined in Appendi B, it is the interaction term itself that represents bias. Thus, it is the marginal effect of this term in isolation that measures the impact of relative bias.

[13]See Chilton et al. (2023) for a similar use of randomization inference to isolate the impact of a potential confounder on point estimates.

[14]Appendix Figure A2a shows the demeaned distribution of the marginal effects, interpreted as each judge's relative bias toward Black attorneys. Appendix Figures A2b and A2c show analogous results for Hispanic and female attorneys.

Figure 3: Finite Sample Distribution of Variance of Judge Specific Coefficients

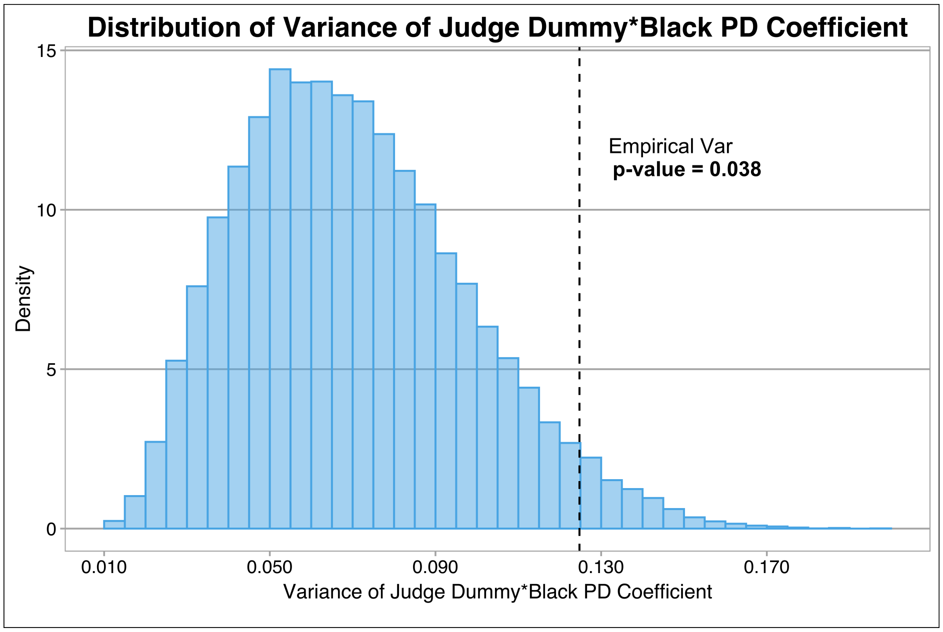

(a) Black vs. White Public Defenders

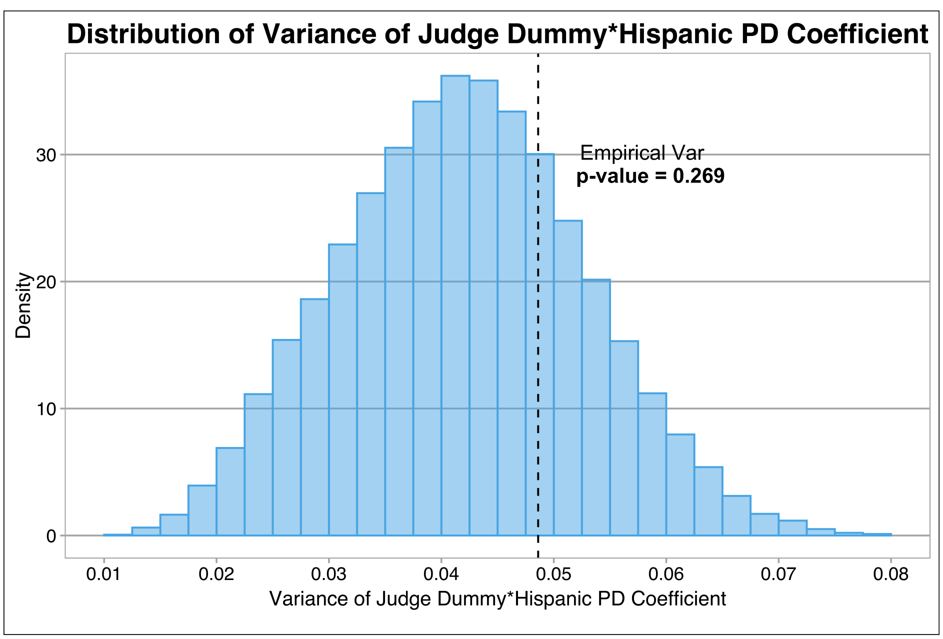

(b) Hispanic vs. White Public Defenders

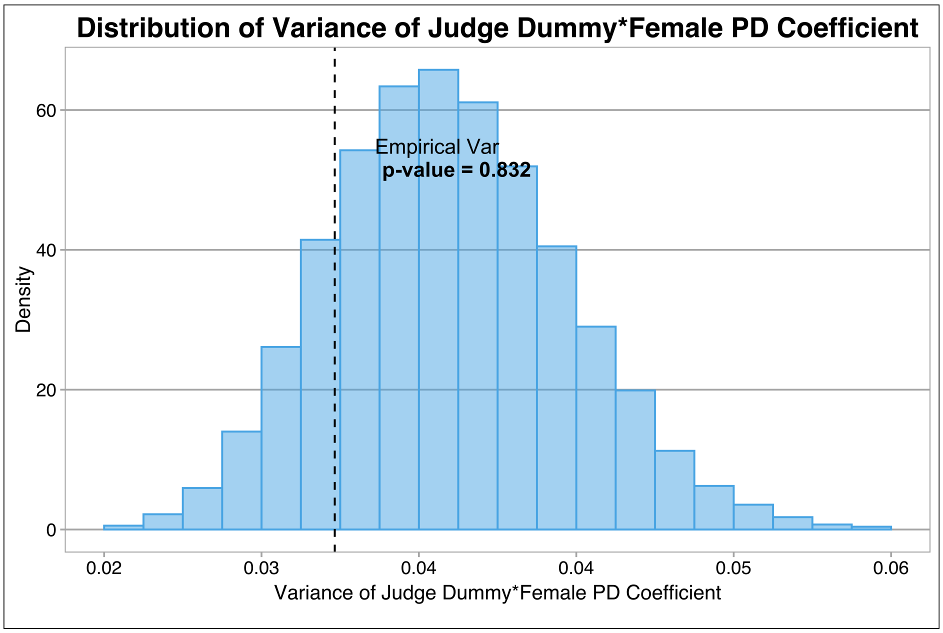

(c) Female vs Male Public Defenders

Density histograms of the distribution of the variance of the estimated coefficients on the judge dummy × public defender demographic group dummy from Equation 1. Variances are calculated from 10,000 simulated datasets constructed under the null hypothesis that there is no bias against public defenders of the relevant demographic group. The dashed line marks the variance of the estimated coefficients from the true dataset. (a) shows the distribution of variances of β_j^B, (b) the distribution of variances of β_j^H, and (c) the distribution of variances of β_j^F.

5 Extensions and Robustness

5.1 Impact on Defendant Outcomes

Agan et al. (2021) find that minority defendants are significantly more likely to choose an attorney of their race when given a choice, with Black defendants being 2.15 times more likely to retain a Black attorney than under random assignment.[15] While the majority of felony defendants rely on indigent defense and would generally be unable to select their attorney, the BJS estimates that 35% of felony defendants in U.S. District Courts have privately retained attorneys (Harlow 2001).[16] Thus, Black defendants would disproportionately bear the impact of bias against Black attorneys.

I estimate the potential impact of judicial bias against Black attorneys on the racial disparities in pretrial release rates in settings where defendants can choose their attorneys. I use data from Bexar County District Courts between 2005 and 2013 to estimate the preferences of Black, Hispanic, and White defendants for attorneys of different races and genders.[17] Using these preferences, I simulate matches between defendants and attorneys in my dataset of Miami-Dade cases.[18]

Under these simulated matches, the predicted pretrial release rate gap between Black and White defendants is 3.1 percentage points.[19] I then forecast the release gap between Black and White defendants if all judges shared the same level of relative bias. As I cannot determine which judge, if any, is unbiased in their treatment of Black attorneys,[20] I estimate

[15]See Hoag (2021) for a discussion of the potential reasons why a Black defendant may prefer a Black attorney, including improved quality of communication and trust.

[16]Agan et al. (2021) find a similar number in their Bexar County data at 33%.

[17]This data is the same data used by Agan et al. (2021).

[18]I also reallocate cases to judges so that all cases involve judges from the largest connected set.

[19]Note that this does not account for any potential match effects once defendants can choose their preferred attorney. However, this does not affect the analysis, as the purpose is to isolate the potential effect of judicial bias.

[20]A judge would be unbiased if their relative bias coefficient is precisely equal to the negative of the mean level of judicial bias. However, as the mean level of bias is unidentified, the unbiased judge is similarly unidentified.

the defendant racial gap at several quantiles of the relative bias coefficient. As the simulation introduces correlation between defendant and public defender race, in this analysis I need to assume that the interaction between defendant and public defender race does not impact judicial bias.

Table 7 shows the simulation results. If all judges had the same relative bias coefficient as the 25$^{\text{th}}$ percentile judge, Black defendants would have a 4.2 percentage point lower release rate than White defendants. This gap would decrease to 3.2 percentage points if all judges had the same level of bias as the 75$^{\text{th}}$ percentile judge, a decline of 23.8%. If all judges had the relative bias coefficient of the judge most favorable to Black attorneys, the Black-White release gap would flip, with Black defendants detained 0.6 percentage points less often. However, this likely overstates the impact of relative bias as it includes the variation due to sampling variability. An increase in the favorability of all judges toward Black attorneys by one unit of my estimate of the excess variance of the relative bias coefficients would decrease the racial gap in release rates by just 12.2% amongst defendants who choose their attorneys. This would be equivalent to a reduction of 4.2% in the overall racial gap between Black and White defendants.[21] While I cannot estimate the average level of bias held by all judges, it is unlikely that it would be substantially larger than the variation in bias across judges. Thus, while judicial bias against minority attorneys may predominantly affect minority defendants, it does not appear to be a major driver of the racial disparities seen in defendant outcomes.

[21] Assuming that 35% of defendants choose their attorney.

Table 7: Impact on Defendant Racial Gap

Bias Quantile	B-W Release Gap
Baseline	−0.031
Min	−0.052
10th	−0.045
25th	−0.042
50th	−0.036
75th	−0.032
90th	−0.003
Max	0.006

This table provides the forecast gap in the release rate of Black defendants and White defendants under the simulation in section 5.1. A negative value means that a higher proportion of White defendants are released than Black defendants.

5.2 Intersectionality

In analyzing the potential role of bias in Section 4, I estimate separate coefficients for racial and gender bias. This specification implicitly assumes that the role of gender bias is the same for a White female attorney and a Black female attorney. However, this ignores the potential role of intersectionality.[22] Under the concept of "multiple jeopardy," King (1988) asserts that the disadvantages of race and gender often compound. This theory suggests, for example, that in many contexts, Black women face greater disadvantages than the addition of the two statuses would imply. In the current context, it could imply that Black women may be particularly affected by biased perceptions of their performance in a male-dominated profession such as law.

I cannot incorporate intersectionality into my analysis of individual judge heterogeneity due to data limitations. The requirement of a connected set, discussed in Section 3, means that the number of judges across whom I could compare judge-specific coefficients

[22] The term intersectionality was originally coined by Crenshaw in Crenshaw (1989). See Browne & Misra (2003) for an overview of various theories on the operation of intersectionality in the labor market context.

would be too small to draw any inference. Given this, I again group judges by demographic characteristics and test for variation across these groups. I estimate a modified version of Equation 2.

$$\begin{aligned} R_i =& \alpha + \beta_1(B_i^J \cdot B_i \cdot F_i) + \beta_2(F_i^J \cdot B_i \cdot F_i) + \beta_3(B_i^J \cdot B_i) \\ & + \beta_4(B_i^J \cdot F_i) + \beta_5(F_i^J \cdot B_i) + \beta_6(F_i^J \cdot F_i) \\ & + \delta_j + \gamma_d + \psi_p + \kappa X_i + \theta_m + \eta_y + h_i + \varepsilon_i, \end{aligned} \tag{3}$$

Judges and attorneys are separated into two racial groups, "Black" and "non-Black," as well as two gender groups, "male" and "female." The coefficients relating to intersectionality are β_1 and β_2, which are the coefficients on the triple interaction between judge demographic group, attorney race, and attorney gender.

The results from estimating Equation 3 are shown in Table 8. I once again include p-values derived from randomization inference. The coefficient on the interaction term Female $\times$ Black Female PD is negative and statistically significant at the 5% level. Specifically, Black female public defenders are 8.3 percentage points less likely to secure the release of their defendants when a female judge presides over the case compared to when a male judge presides. This suggests that the excess variation in the treatment of Black attorneys may be driven by the treatment of Black female attorneys specifically. This highlights the need for further analysis in a setting where a larger sample size makes it possible to incorporate intersectionality into the interjudge analysis.

5.3 Robustness

A potential concern with the above analysis is that the results may be driven by differences in the behavior of Black attorneys when arguing before certain judges. Similar concerns are present in many observational studies testing for bias. For example, Anwar & Fang (2012) must assume that physicians can perfectly account for differences in behavior across patient demographic groups when forming their prior beliefs of the severity of

Table 8: Test of Homophily (Intersectionality)

	Dependent variable:	
	Pretrial Release	
		p-value
	(1)	(2)
Black Judge * Black Female PD	−0.149	0.27
	(0.076)	
Female Judge * Black Female PD	−0.083	0.04
	(0.034)	
Black Judge * Black PD	0.094	0.30
	(0.050)	
Female Judge * Black PD	−0.007	0.81
	(0.033)	
Black Judge * Female PD	0.101	0.01
	(0.023)	
Female Judge * Female PD	0.006	0.64
	(0.013)	
Mean	0.57	
Observations	42,202	
Adjusted R^2	0.129	

This table reports the results of the regression in Equation 3. Robust standard errors clustered at the Judge ID and Public Defender ID level.

a patient's illness.[23] In my context, there are many potential reasons why a Black attorney may behave differently depending on the race of the judge. For example, "stereotype threat," the pressure to avoid conforming to a negative stereotype, has been shown to lead to underperformance in various settings (Spencer et al. 2016). To address this concern, I rerun the above analysis, measuring the inter-judge variation in the racial gap across White male judges. While it seems reasonable to consider that Black attorneys may act differently before a White judge than before a Black judge due to discrimination concerns, in the current setting it seems less likely that behavior should vary substantially before judges of the same race and gender. During the week, attorneys are assigned to a specific courtroom and thus will only appear before particular judges. Therefore, in most bail hearings, the public defender will have had minimal experience arguing before the assigned judge. The results of this analysis are presented below in Figure 4. Looking at the racial gap between Black and White attorneys, I continue to find statistically significant variation, with only 0.3% of the simulated datasets having a larger variance in the coefficients of White male judges. This is consistent with the null result from Equation 2, which indicates that judge demographics do not account for the variation across judges in the gap between Black and White attorneys.

While this does not rule out the possibility of changes in behavior when appearing before specific judges, the fact that the excess variation is more extreme amongst these observationally similar judges, coupled with the fact that these attorneys have had limited prior experience with each judge, suggests that behavioral changes do not drive the results.

[23]This study tests for racial bias by emergency department physicians in the diagnosis and treatment of minority patients.

Figure 4: Black vs White Public Defenders - White, male judges

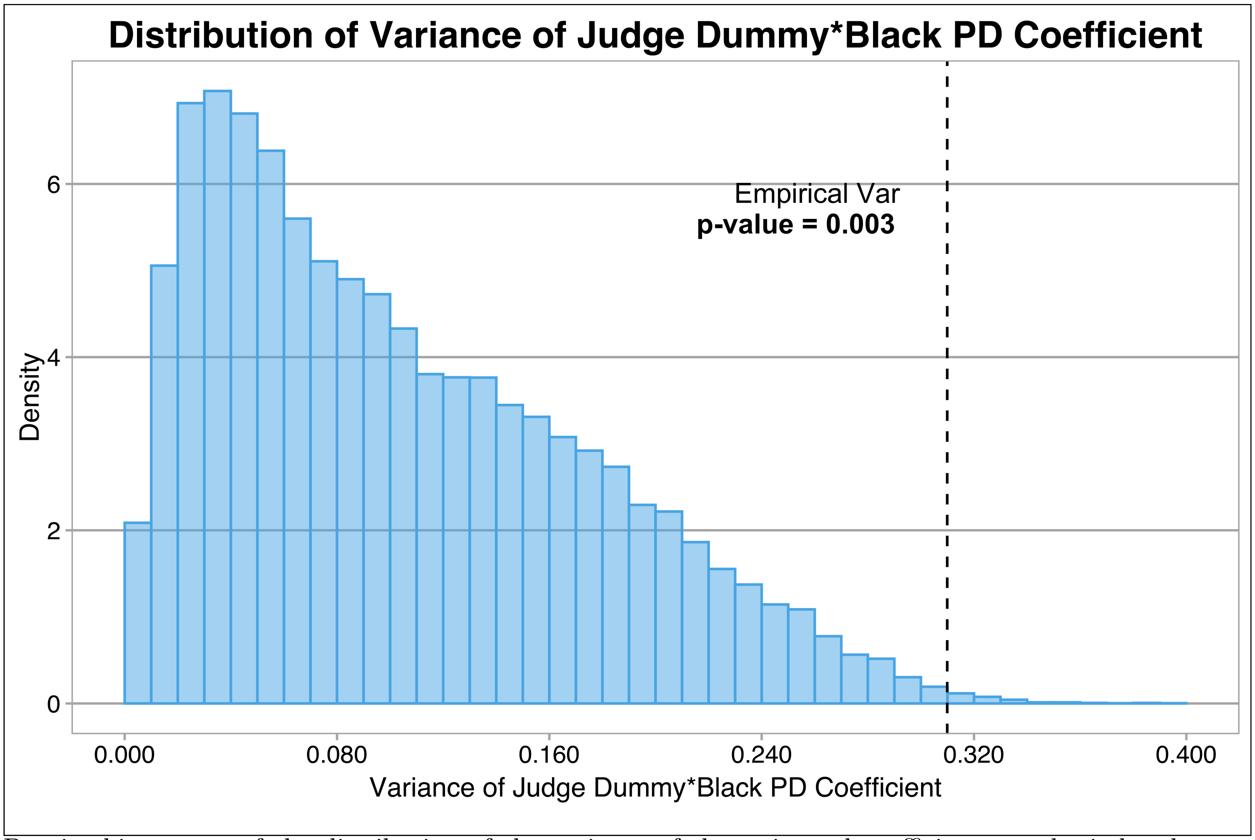

Density histograms of the distribution of the variance of the estimated coefficients on the judge dummy × Black attorney dummy for White male judges. These estimate come from Equation 1. Variances are calculated from 10,000 simulated datasets constructed under the null hypothesis that there is no bias in the treatment of Black attorneys. The dashed line marks the variance of the estimated coefficients for White male judges from the true dataset.

6 Conclusion

This paper investigates whether bias against different attorney groups influences case outcomes after controlling for potential unobservable differences across racial and gender groups. By exploiting the random assignment of judges and attorneys to bail shifts, I test for bias by analyzing variation across judges in the rates at which different demographic groups of attorneys secure the pretrial release of their defendants. In the setting of Miami-Dade County first appearance hearings, I find evidence that the clients of Black attorneys are less likely to be released on bail if the hearing occurs before judges who are unfavorable to Black attorneys, but no evidence that judges are biased in the treatment of Hispanic or female attorneys.

Before concluding, it is important to note several limitations. First, the requirement of connected sets in the main specification limited the analysis to a subset of judges, which may have affected the null results in particular. Second, the analysis is restricted to a single county—Miami-Dade—and may not reflect the experiences of minority and female attorneys in other regions. Further study is needed to assess whether these findings generalize more broadly across the United States.

Most importantly, while the methodology provides evidence of differential treatment, it does not identify the direction of bias, as the counterfactual gap in release rates absent bias cannot be recovered. That said, any systematic bias in the legal system is concerning and may undermine public confidence in the fairness of judicial proceedings.

References

Abowd, J. M., Creecy, R. H., Kramarz, F. et al. (2002), Computing person and firm effects using linked longitudinal employer-employee data, Technical report, Center for Economic Studies, US Census Bureau.

Abrams, D. S., Bertrand, M. & Mullainathan, S. (2012), 'Do judges vary in their treatment of race?', *The Journal of Legal Studies* **41**(2), 347–383.

Agan, A., Freedman, M. & Owens, E. (2021), 'Is your lawyer a lemon? Incentives and selection in the public provision of criminal defense', *The Review of Economics and Statistics* **103**(2), 294–309.

Ai, C. & Norton, E. C. (2003), 'Interaction terms in logit and probit models', *Economics Letters* **80**(1), 123–129.

Albright, A. (2019), 'If you give a judge a risk score: evidence from Kentucky bail decisions', *Law, Economics, and Business Fellows' Discussion Paper Series* **85**.

Alesina, A. & La Ferrara, E. (2014), 'A test of racial bias in capital sentencing', *American Economic Review* **104**(11), 3397–3433.

Anwar, S. & Fang, H. (2012), 'Testing for the role of prejudice in emergency departments using bounceback rates', *The BE Journal of Economic Analysis & Policy* **13**(3).

Arnold, D., Dobbie, W. & Hull, P. (2022), 'Measuring racial discrimination in bail decisions', *American Economic Review* **112**(9), 2992–3038.

Arnold, D., Dobbie, W. & Yang, C. S. (2018), 'Racial bias in bail decisions', *The Quarterly Journal of Economics* **133**(4), 1885–1932.

Bertrand, M., Chugh, D. & Mullainathan, S. (2005), 'Implicit discrimination', *American Economic Review* **95**(2), 94–98.

Browne, I. & Misra, J. (2003), 'The intersection of gender and race in the labor market', *Annual Review of Sociology* **29**(1), 487–513.

Chen, D. L., Halberstam, Y. & Yu, A. (2017), 'Covering: mutable characteristics and perceptions of voice in the US Supreme Court', *SSRN 2816241* .

Chilton, A., Cotropia, C., Rozema, K. & Schwartz, D. (2023), 'Political ideology and judicial administration: evidence from the COVID-19 pandemic', *Journal of Law, Economics, and Organization* **41**(1), 91–126.

Crenshaw, K. (1989), 'Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics', *University of Chicago Legal Forum* **1989**(1), 139–167.

Didwania, S. H. (2020*a*), 'Discretion and disparity in federal detention', *Nw. U. L. Rev.* **115**, 1261.

Didwania, S. H. (2020*b*), 'The immediate consequences of federal pretrial detention', *American Law and Economics Review* **22**(1), 24–74.

Dobbie, W., Goldin, J. & 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**(2), 201–240.

Eckstein, S. (2014), How Cubans transformed Florida politics and leveraged local for national political influence., *in* V. D. Baldera & R. May, eds, 'La Florida: 500 Years of Hispanic Presence', University of Florida Press, Gainesville, chapter 12, pp. 263–283.

Fischman, J. B. & Schanzenbach, M. M. (2012), 'Racial disparities under the federal sentencing guidelines: The role of judicial discretion and mandatory minimums', *Journal of Empirical Legal Studies* **9**(4), 729–764.

Fishman, G., Rattner, A. & Turjeman, H. (2006), 'Sentencing outcomes in a multinational society: When judges, defendants and victims can be either Arabs or Jews', *European Journal of Criminology* **3**(1), 69–84.

Goldkamp, J. S. & Gottfredson, M. R. (1988), 'Development of bail/pretrial release guidelines in Maricopa County Superior Court, Dade County Circuit Court and Boston Municipal Court. Washington, DC: National Institute of Justice', *National Institute of Justice* .

Grossman, G., Gazal-Ayal, O., Pimentel, S. D. & Weinstein, J. M. (2016), 'Descriptive representation and judicial outcomes in multiethnic societies', *American Journal of Political Science* **60**(1), 44–69.

Gupta, A., Hansman, C. & Frenchman, E. (2016), 'The heavy costs of high bail: Evidence from judge randomization', *The Journal of Legal Studies* **45**(2), 471–505.

Harlow, C. W. (2001), *Defense counsel in criminal cases*, US Department of Justice, Office of Justice Programs, Bureau of Justice.

Harris, A. P. & Sen, M. (2019), 'Bias and judging', *Annual Review of Political Science* **22**, 241–259.

Heaton, P., Mayson, S. & Stevenson, M. (2017), 'The downstream consequences of misdemeanor pretrial detention', *Stan. L. Rev.* **69**, 711.

Hoag, A. (2021), 'Black on Black representation', *NYU. L. Rev.* **96**, 1493.

Hodgson, S. & Pryor, B. (1984), 'Sex discrimination in the courtroom: Attorney's gender and credibility', *Psychological Reports* **55**(2), 483–486.

Kastellec, J. P. (2021), 'Race, context, and judging on the courts of appeals: Race-based panel effects in death penalty cases', *Justice System Journal* **42**(3-4), 394–415.

King, D. K. (1988), 'Multiple jeopardy, multiple consciousness: The context of a black feminist ideology', *Signs: Journal of women in culture and society* **14**(1), 42–72.

Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J. & Mullainathan, S. (2018), 'Human decisions and machine predictions', *The Quarterly Journal of Economics* **133**(1), 237–293.

Leslie, E. & Pope, N. G. (2017), 'The unintended impact of pretrial detention on case outcomes: Evidence from new york city arraignments', *The Journal of Law and Economics* **60**(3), 529–557.

Mikdash, M. & Oh, S. (2024), The role of race in the legal representation of low-income defendants. Working Paper.

Ojito, M. (2000), ‘Best of friends, worlds apart’.
URL: *https://archive.nytimes.com/www.nytimes.com/library/national/race/060500ojito-cuba.html*

Rachlinski, J. J., Johnson, S. L., Wistrich, A. J. & Guthrie, C. (2008), ‘Does unconscious racial bias affect trial judges’, *Notre Dame L. Rev.* **84**, 1195.

Sen, M. (2015), ‘Is justice really blind? Race and reversal in US courts’, *The Journal of Legal Studies* **44**(S1), S187–S229.

Shayo, M. & Zussman, A. (2011), ‘Judicial ingroup bias in the shadow of terrorism’, *The Quarterly Journal of Economics* **126**(3), 1447–1484.

Sloan, C. (2019), Racial bias by prosecutors: Evidence from random assignment, *in* ‘ICCJ 2019: International Conference on Criminal Justice June’, pp. 25–26.

Spencer, S. J., Logel, C. & Davies, P. G. (2016), ‘Stereotype threat’, *Annual Review of Psychology* **67**, 415–437.

Stevenson, M. T. (2018), ‘Distortion of justice: How the inability to pay bail affects case outcomes’, *The Journal of Law, Economics, and Organization* **34**(4), 511–542.

Szmer, J. J., Sarver, T. A. & Kaheny, E. B. (2010), ‘Have we come a long way, baby? The influence of attorney gender on Supreme Court decision making’, *Politics & Gender* **6**(1), 1–36.

Szmer, J., Kaheny, E. B., Sarver, T. & DeCamillis, M. (2013), ‘The impact of attorney gender on decision making in the United States Courts of Appeals’, *Journal of Women, Politics & Policy* **34**(1), 72–100.

A Finite Sample Test

To test for bias, it is necessary to determine whether the variation in the attorney racial and gender gaps across judges exceeds what would be expected due to sampling variability alone.

As Equation 1 is a probit model, a standard test for the joint significance of each set of relative bias coefficients would be a likelihood ratio test. Each set of bias coefficients $(\beta^B, \beta^H, \beta^F)$, would be tested separately with the full model run and then a restricted version estimated without the relevant set of coefficients. When testing for variation in the treatment of Black attorneys, the null hypothesis is $H_0 : \beta^B_j = \beta^B_k \ \forall \ j, k \in J$.[24] The alternate hypothesis is that at least one is not equal to the others.

The standard approach for this test would be to rely upon asymptotics. The ratio of the log-likelihoods from the full and restricted models would be compared to the χ^2 distribution with the number of degrees of freedom given by the number of coefficients being tested. However, in my setting, the asymptotic χ^2 is not appropriate. This is because while the overall sample is large, there are a relatively low number of observations at the judge $\times$ attorney demographic group level. As a result, there is finite-sample bias in estimating the judge-specific relative bias coefficients. To illustrate this, I simulate data from the theoretical model presented in Section C. I simulate this data with the bias of each judge set to 0. Thus, by construction, the null hypothesis should not be rejected. I run $2,000$ simulations. In each simulation, the likelihood ratio is computed and compared to the asymptotic χ^2 distribution to calculate the p-value. The results of these simulations are shown below in Figure 5. The p-values should be approximately uniformly distributed. For example, 5% of simulations should have a p-value below 5% *i.e.* using a 5% significance level, the type I error rate should be 5%. However, in Figure 5, over 25% of simulations have a p-value below 5%, which would lead to substantial overrejection of the null hypothesis.

[24]Note that the identification is in terms of the variation between the coefficients, not the level. Hence, they do not need to be equal to 0.

Figure 5: Distribution of Simulated p-values

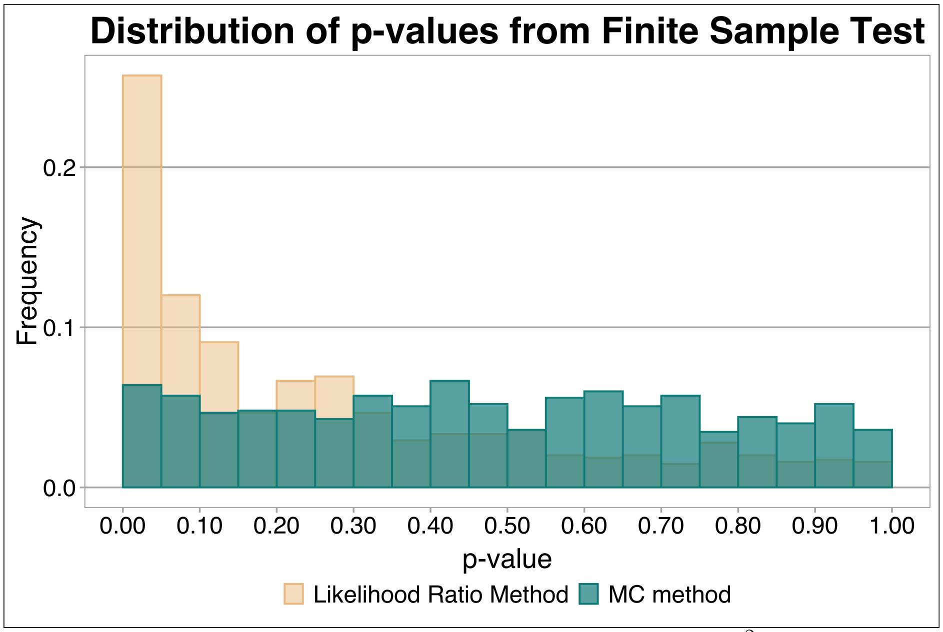

Distribution of p-values from the likelihood ratio test using the asymptotic χ^2 and the proposed Monte Carlo Method. Data simulated under the null hypothesis of no bias in the treatment of attorneys using the model outlined in Section C, 1000 simulations were run.

For this reason, I instead use a Monte Carlo simulation method to construct the finite-sample distribution of my test statistic as discussed in Section 3. For the purpose of this test, I use the variance of the set of coefficients as my test statistic to measure variation.

The relevant null hypotheses that I wish to test with respect to Equation 1 are:

- $H_0 : \beta_j^B = \beta_k^B \ \forall \ j, k \ \in \ J$
- $H_0 : \beta_j^H = \beta_k^H \ \forall \ j, k \ \in \ J$
- $H_0 : \beta_j^F = \beta_k^F \ \forall \ j, k \ \in \ J$

To construct the finite sample distribution of the variance under the null of no bias, I use the following process. Using the true dataset, equation 1 is estimated without the judge $\times$ public defender demographics terms.[25] The coefficients of all control variables in Equation 1 are held fixed at these estimated values for all simulations. Within each simulation, taking the original data, I randomize the race and gender of each public defender at the judge $\times$

[25] i.e., The terms for $J_{ij} \times B_i^D$, $J_{ij} \times H_i^{PD}$ and $J_{ij} \times F_i^{PD}$ are not included in this estimation

attorney level to create the simulated data. The randomization process is as follows. For a given judge, let PD_j denote the set of public defenders that argued before judge j while R_j^{PD} and G_j^{PD} denote the set of races and genders of these public defenders. In each simulation, each $pd \in PD_j$ is assigned a race r and gender g sampled without replacement from R_J^{PD} and G_j^{PD}.[26] Hence, in a given simulation, a public defender's race is constant for all cases with a given judge but may vary across judges. The outcome of each case is held constant.

Using this simulated data, I use maximum likelihood estimation to re-estimate $\beta_j^B, \beta_j^H, \beta_j^F$, and δ_j. As attorney fixed effects are not re-estimated, the connected sets do not change. Hence, in each simulation, I compare the relative bias coefficients of the same set of judges. I calculate the variance of each set of relative bias coefficients as a measure of dispersion of the coefficients. As attorney race and gender are randomly assigned in each simulation, this dispersion across judges is caused by sampling variation. The distribution of these variances thus provides the finite-sample distribution of the variance of relative bias coefficients under the null hypothesis of no bias. The variance from the true data is compared to this distribution to obtain the p-value for the probability of this variance occurring without judges varying in their treatment of minority and female attorneys.

Figure 5 shows the distribution of p-values when applying this Monte Carlo method to data simulated from the theoretical model under the null hypothesis of no bias. As can be seen in this figure, the resulting p-values are approximately uniformly distributed.

B Theoretical Model

C Conceptual Framework

This section presents a theoretical model of the bail hearing. For simplicity in the model's exposition, I will consider a setting where the only two demographic groups are Black and White. This is without loss of generality, and the implications extend to my empirical

[26]To ensure that the connected sets are unchanged, this randomization process is applied only to attorneys that have seen more than 15 cases with a given judge.

setting, where I consider bias against two distinct racial minorities as well as gender bias.

Consider a judge, j, who must decide whether to release a defendant, i, pending trial. The defendant has a latent type θ_i, representing his willingness to engage in pretrial misconduct,[27] normalized such that he will engage in misconduct if $\theta_i < 0$. The defendant's type consists of a deterministic component, $\bar{\theta}_{X_i}$, which depends on the characteristics of i observable by the judge (X_i)[28] and a stochastic component, $\delta_i \sim \mathcal{N}(0, \sigma_\delta^2)$. I assume that the probability of pretrial misconduct upon release is not affected by any conditions that may be imposed alongside the release. The defendant's type θ_i is given by:

$$\theta_i | X_i = \bar{\theta}_{X_i} + \delta_i.$$

Therefore, the judge's prior belief over the defendant's type given the observable characteristics is:

$$\theta | X_i \sim \mathcal{N}(\bar{\theta}_{X_i}, \sigma_\delta^2).$$

A hearing is held where the defense attorney d and prosecutor p respectively argue for the release/remand of the defendant. Each defense attorney (prosecutor) has a race denoted r_d (r_p) and is of varying quality denoted by q_d (q_p). From the hearing, the judge receives a noisy signal of the defendant i's pretrial misconduct potential given by:

$$v_{ij} = \theta_i + (q_d - b_j^{r_d}) - (q_p - b_j^{r_p}) + \zeta_i, \tag{4}$$

where $b_j^{r_d}$ $(b_j^{r_p})$ represents the racial bias the judge has against the defense attorney (prosecutor) in i's case. There is also a noise component $\zeta \sim \mathcal{N}(0, \sigma_\zeta^2)$. A judge is unaware of their

[27] In Miami-Dade, the relevant pretrial misconduct is either the commission of a new crime while released on bail or a failure to appear in court for future hearings related to the current charge. For simplicity, I will consider "misconduct" as a binary event for the purpose of this model.

[28] This includes both demographic and case characteristics.

own b_j and thus cannot adjust for this when forming their posterior. In addition, as judges deal with each attorney so infrequently, it is assumed that they do not know the true quality of each attorney (q_d, q_p). Instead, they know only the distribution of attorney quality for each type: $q_d \sim \mathcal{N}(\bar{q}^d, \sigma^2_{q^d})$ and $q_p \sim \mathcal{N}(\bar{q}^p, \sigma^2_{q^p})$. [29] This allows a higher-quality attorney to increase the probability of a successful outcome.[30] Note that attorney quality may be correlated with attorney race. Therefore, the unconditional quality distributions are mixtures of the race-specific quality distributions. Judges thus consider that the signal comes from the following process, which is misspecified due to failing to account for potential bias:

$$\hat{v}_i = \theta_i + \bar{q}^d - \bar{q}^p + \eta_i,$$

where $\eta_i \sim \mathcal{N}(0, \sigma^2_\zeta + \sigma^2_{q^d} + \sigma^2_{q^p})$. Given the signal, their model of the signal generating process, and their prior, the judge then forms a posterior $\hat{p}(\theta_i|v_i, X_i)$. The judge believes that

$$\theta|v_i, X_i \sim \mathcal{N}\left(\frac{\sigma^2_\eta \bar{\theta}_{X_i} + \sigma^2_\delta(v_i - \bar{q}^d + \bar{q}^p)}{\sigma^2_\eta + \sigma^2_\delta}, \frac{\sigma^2_\eta \sigma^2_\delta}{\sigma^2_\eta + \sigma^2_\delta}\right).$$

A risk-neutral judge will then release the defendant i if the expected cost of releasing the defendant is less than the benefit of release.

For simplicity, misconduct is modeled as a binary outcome with a constant societal cost of C if misconduct occurs and 0 if there is no misconduct. Hence, the judge's expectation of the cost of release is $\mathbb{E}[C|X_i, v_i] = \hat{P}(\theta_i < 0|v_i, X_i) \times C$. The perceived benefit of release (Π_j) is allowed to vary across judges to allow for the differences in leniency that have been

[29]Within the model it is assumed that q_p, q_d and ζ_i are independent. In Section **??**, I explain why this assumption is satisfied in my setting.

[30]If the judge knew the exact quality of an attorney, a rational judge would perfectly offset this when forming their posterior. This would imply that attorney quality had no impact on case outcomes. Such a result seems implausible. Relaxing this assumption would imply that any significant difference in release rates between Black and White attorneys was due to judicial bias.

documented in the literature.[31] The perceived benefit is assumed not to vary by defendant characteristics.[32]

The judge will release individual i if

$$\hat{P}(\theta_i < 0|v_i, X_i) \times C \leq \Pi_j.$$

The probability that defendant i is released by judge j is thus given by

$$P(Y_i = 1|X_i, q, b_j, \Pi_j) =$$
$$\Phi\left(\frac{1}{\sigma_\xi^2 + \sigma_\delta^2}\left[\frac{\sigma_\xi}{\sigma_\delta}\Phi^{-1}\left(\frac{\Pi_j}{C}\right)\sqrt{\sigma_\xi^2 + \sigma_\delta^2} + \left(\frac{\sigma_\xi^2}{\sigma_\delta^2} + 1\right)\bar{\theta}_{X_i} + \tilde{q}_d - b_j^{r_d} - \tilde{q}_p + b_j^{r_p}\right]\right),$$

where Y is a dummy for whether individual i is released pretrial, $\tilde{q}_d = q_d - \bar{q}^d$ and $\tilde{q}_p = q_p - \bar{q}^p$. Note it is assumed that $0 < \frac{\Pi_j}{C} < 1 \; \forall \; j \in J$, i.e., there is no judge who would never release a defendant nor one who would release every defendant.

This can be rearranged to give

$$P(Y_i = 1|X_i, q, b_j, \Pi_j) =$$
$$\Phi\Bigg(\underbrace{\frac{\tilde{q}_d - \bar{b}^{r_d}}{\sigma_\xi^2 + \sigma_\delta^2}}_{\text{net defender quality}} + \underbrace{\frac{\bar{b}^{r_p} - \tilde{q}_p}{\sigma_\xi^2 + \sigma_\delta^2}}_{\text{net prosecutor quality}} + \underbrace{\frac{\sigma_\xi}{\sigma_\delta}\frac{\sqrt{\sigma_\xi^2 + \sigma_\delta^2}}{\sigma_\xi^2 + \sigma_\delta^2}\Phi^{-1}\left(\frac{\Pi_j}{C}\right)}_{\text{effect of judge leniency}} \tag{5}$$
$$+ \underbrace{\left(\frac{\sigma_\xi^2}{\sigma_\delta^2} + 1\right)\frac{\bar{\theta}_{X_i}}{\sigma_\xi^2 + \sigma_\delta^2}}_{\text{effect of defendant demographics}} - \underbrace{\frac{\tilde{b}_j^{r_d}}{\sigma_\xi^2 + \sigma_\delta^2}}_{\substack{\text{idiosyncratic judicial bias} \\ \text{against defender}}} + \underbrace{\frac{\tilde{b}_j^{r_p}}{\sigma_\xi^2 + \sigma_\delta^2}}_{\substack{\text{idiosyncratic judicial bias} \\ \text{against prosecutor}}} \Bigg),$$

where $\bar{b}^{r_d}$ is the mean level of bias held by all judges against attorneys of race r_d while

[31] Arnold et al. (2018) for example, document substantial variation in pretrial release rates across judges in Miami-Dade.

[32] This perceived benefit also does not vary by attorney race. Hence, the effect of judicial bias against marginalized attorneys occurs purely through the information channel rather than taste-based bias.

$\tilde{b}_j^{r_d}$ is the deviation of judge j's bias from this mean level of bias.[33]

To illustrate the application of my identification method in the context of the theoretical model in Equation 5, consider the estimation of the following probit model. Once again, for simplicity of exposition, I consider White and Black attorneys as the only demographic groups for this illustration.

$$P(Y_i = 1|j, d, X_i) = \Phi\left(\alpha + \delta_j + \gamma_d + \sum_{j=1}^{J} \beta_j (J_{ij} \times B_i^D) + \kappa Z_i\right), \tag{6}$$

where Y_i is the outcome of whether individual i was released pretrial; δ_j and γ_d are fixed effects for the judge and public defender; J_{ij} is an indicator for judge identity; B_i^D is an indicator for whether a Black public defender represents the defendant i; and Z_i is a vector of observable defendant characteristics.[34] δ_j and γ_d are identified by the fact that a given public defender will argue before several different judges and a given judge will hear cases argued by several different public defenders. δ_j will thus capture variation in the leniency of judges relative to the omitted judge, i.e., variation in $\frac{\sigma_\eta}{\sigma_\delta}\frac{\sqrt{\sigma_\eta^2+\sigma_\delta^2}}{\sigma_\eta^2+\sigma_\delta^2}\Phi^{-1}\left(\frac{\Pi_j}{C}\right)$ in Equation 5. γ_d captures the terms which are constant for a given defense attorney, i.e., $\frac{\tilde{q}_d - \bar{b}^{r_d}}{\sigma_\eta^2+\sigma_\delta^2}$ in Equation 5. This consists of the individual attorney's ability net of the effect of the average level of judicial bias against attorneys of their race. From this, we can see that the mean level of bias across all judges is not separately identifiable from attorney quality in Equation 5. Therefore, I focus on identifying the relative bias term, $\tilde{b}_j^{r_d}$. β_j is identified through the idiosyncratic variation in the racial gap in public defender release rates for a given judge j relative to the omitted judge. These β_j coefficients estimate the term $\frac{\tilde{b}_j^{B_d} - \tilde{b}_j^{W_d}}{\sigma_\eta^2+\sigma_\delta^2}$ from Equation 5 for each judge as they are the only terms which vary across both public defender race and judge identity. Thus, provided that the level of bias is not uniform across judges, judicial bias against attorneys will cause inter-judge variation in β_j.

[33]Thus, we have that $b_j^{r_d} = \bar{b}^{r_d} + \tilde{b}_j^{r_d}$.

[34]This is distinct from X_i which included all characteristics observable to the judge some of which may not be observable to the econometrician.

The random assignment of cases to judges and cases to attorneys means that neither judge nor public defender identity is correlated with Z_i. Thus κZ_i separately estimates the effect of defendant characteristics, i.e., $\left(\frac{\sigma_\eta^2}{\sigma_\delta^2}+1\right)\frac{\bar{\theta}_{X_i}}{\sigma_\eta^2+\sigma_\delta^2}$ in Equation 5.[35]

[35]Note that allowing Π_j to depend on X_i would mean that variation in $\bar{\theta}_{X_i}$ would not be separately identified from variation in the perceived benefit. However, while this would mean that κ no longer identifies changes in the perceived probability of misconduct, it would not impact the identification of judicial bias.

D Appendix Figures

Figure A1: Distribution of Judge Specific Coefficients

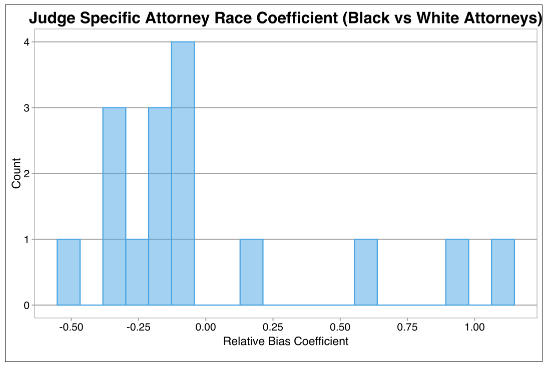

(a) Black vs. White Public Defenders

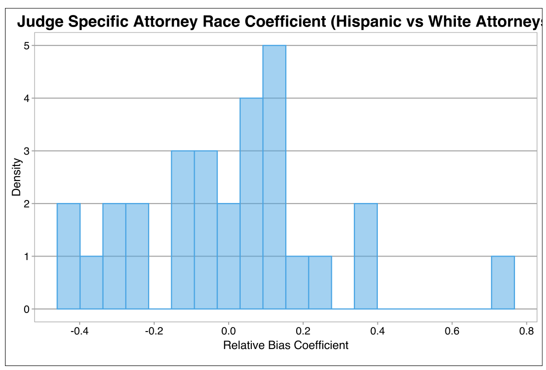

(b) Hispanic vs. White Public Defenders

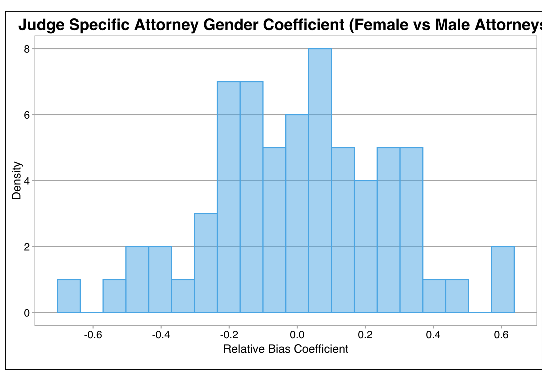

(c) Female vs Male Public Defenders

Histograms of the distribution across judges of the estimates of the judge dummy × attorney demographic group dummy from Equation 1. (a) shows the distribution of β_j^B, (b) the distribution of β_j^H and (c) the distribution of β_j^F. This coefficient measures the gap in the rate at which the relevant demographic group secures the release of their defendants relative to the rate at which White male attorneys do when arguing in front of this judge. Only judges who preside over at least 30 pretrial hearings involving public defenders of the relevant marginalized demographic group and 30 involving White male public defenders are included. Coefficients rescaled to be mean zero

Figure A2: Distribution of Marginal Effects

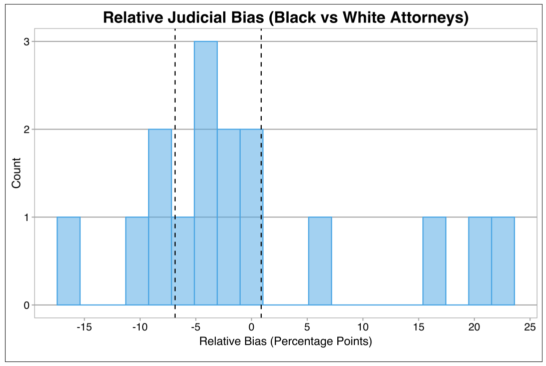

(a) Black vs. White Public Defenders

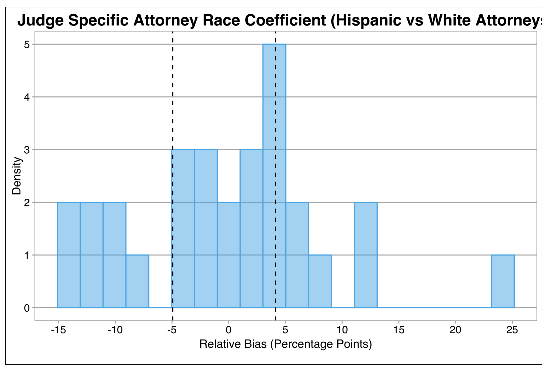

(b) Hispanic vs. White Public Defenders

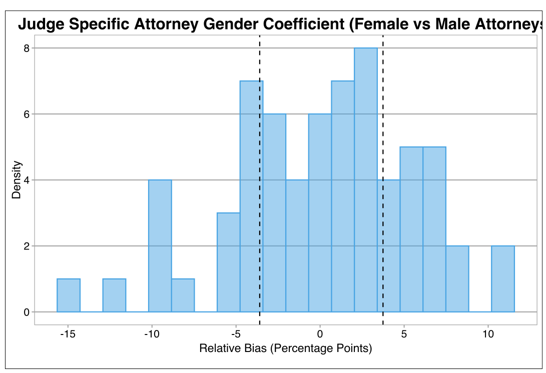

(c) Female vs Male Public Defenders

Histograms of the distribution across judges of the average marginal effects of the judge dummy $\times$ attorney demographic group dummy from Equation 1. The average marginal effect has been calculated across all cases involving a judge in the relevant connected set. (a) shows the distribution of β_j^B (b) shows the distribution of β_j^H and (c) shows the distribution of β_j^F. This coefficient measures the gap in the rate at which the relevant demographic group secures the release of their defendants relative to the rate at which White male attorneys do when arguing in front of this judge. Only judges who preside over at least 30 pretrial hearings involving public defenders of the relevant marginalized demographic group and 30 involving White male public defenders are included. Coefficients have been rescaled to be mean zero

E Appendix Tables