Judicial Bias Against Minority and Female Attorneys*

Priyanka Goonetilleke

University of Pennsylvania

January 2024

Click here for the latest version

Abstract

This paper tests for judicial bias in the treatment of racial minority and female attorneys by examining its impact on bail decisions. I account for potential omitted variable bias by exploiting the random assignment of first-appearance hearings in Miami-Dade County to public defenders and judges. I analyze the differences across judges in the rates at which attorneys of different demographic groups secure a pretrial release. I develop a finite sample test of significance that accounts for the small sample of cases at the individual judge level. I find significant variation in the release rate between Black and White attorneys measured across different judges. However, I do not find evidence of judicial bias in the treatment of female or Hispanic attorneys in my setting. Using my estimates, I sort judges by favorability towards Black attorneys. A defendant with a Black attorney is 2.8 percentage points less likely to be released when assigned a judge in the bottom quartile of this ranking compared with being assigned a judge in the top quartile. In settings where they can choose their representation, Black defendants hire Black attorneys at higher rates. Leveraging data from an alternative setting where defendants can choose their attorneys, I simulate client-attorney matches that correspond to these preferences. Under these simulated matches, I find that increasing judges' favorability toward Black attorneys by one standard deviation of the estimated relative favorability would decrease the defendant racial gap in pretrial release rates by 33%. As case outcomes affect not only defendants but also attorneys' productivity and wages, judicial bias may also help explain minorities' continued underrepresentation in the legal profession, particularly in its higher ranks.

*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, 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. 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 key underlying principle of many modern societies is equal justice under the law. Given this principle's importance, a significant body of research has analyzed the presence and effect of judicial bias,[1] particularly the unwarranted impact of an individual's demographic characteristics on judicial decisions. These papers have shown that in criminal cases in the U.S. and other countries, the conditions of bail, whether a defendant is found guilty, and the length of their sentence may be affected by judicial bias. However, the prior research has focused on the interaction between judges and defendants.[2] In an adversarial system, attorneys present their party's case, and the judge or jury makes findings of fact in light of the conflicting evidence. Thus, bias may influence the weight given to an attorney's arguments. For example, a recent survey found that most female respondents agreed that male judges give more credibility to the arguments of male attorneys than those of female attorneys (New York State Judicial Committee on Women in the Courts 2020).

The potential effects of bias against minority and female attorneys are twofold. First, prior research has shown that in the U.S., criminal defendants are more likely to be represented by attorneys from their demographic group.[3] Therefore, bias against minority attorneys may account for some of the racial disparities seen in the U.S. criminal justice system. Prior studies may have incorrectly classified this impact as bias against defendants. Second, bias against attorneys from marginalized groups would also impact their career success relative to White male colleagues. Racial minorities continue to be significantly underrepresented in the U.S. legal profession, with only 19% of attorneys identifying as people of color compared to 34% of law students (American Bar Association 2022). Although women account for almost 50% of associates at major U.S. law firms, they are similarly underrepresented in senior roles, comprising just 25% of partners at these firms (National Association for Law Placement 2022).[4] By reducing the productivity of attorneys from marginalized groups, judicial bias may partially explain this lack of diversity, as well as persistent wage gaps.

This paper uses random assignment of attorneys and judges to cases to test for judicial bias in the treatment of Black, Hispanic, and female attorneys.[5] I define a defense attorney's release rate as the proportion of their clients released while awaiting trial. Differences in release rates across defense attorneys are insufficient to prove judicial bias because of potential omitted variable bias (OVB). The two key potential sources of OVB are (1) the selection of cases by attorneys and (2) differences in the distribution of ability among demographic groups of attorneys. If minority and female attorneys represent defendants with a higher probability of pretrial misconduct, they

[1]See *e.g.*, Abrams et al. (2012), Albright (2019), Alesina & La Ferrara (2014), Arnold et al. (2018, 2022), Didwania (2020*a*), Fischman & Schanzenbach (2012), Grossman et al. (2016), Shayo & Zussman (2011).

[2]Papers such as Didwania (2022), Rehavi & Starr (2014), Sloan (2019), and Tuttle (2019) do consider the potential for bias in the application of prosecutorial discretion. However, they still focus exclusively on defendants as the subjects of this bias.

[3]For example, Agan et al. (2021) find that Black criminal defendants are more than twice as likely to be represented by an attorney of the same race when they are allowed to choose their attorney relative to when they are randomly assigned. Legal literature such as Troccoli (2002) and Hoag (2021) qualitatively support these findings.

[4]Women have comprised over 40% of law students since the late 1980s (Katz et al. 2023). Thus, this underrepresentation at the senior levels appears to be due to either a lower rate of promotion or greater exit.

[5]Prior literature on discrimination in the justice system has focused on racial discrimination due to the small proportion of female defendants. However, research has shown evidence of gender bias by judges in their dealings with colleagues (*e.g.*, Ash et al. 2021, Jacobi & Schweers 2017), and this bias could similarly impact outcomes for female attorneys.

should, on average, have a lower release rate than White male attorneys, even in the absence of judicial bias. Similarly, unobserved differences in the skill level of an attorney may explain differences in their release rates. If skill is correlated with race or gender, then skill differences alone could create an overall attorney racial or gender gap. Hence, such a gap is insufficient to prove the existence of bias.

Marginal outcome based tests are commonly used in the literature as a solution to OVB when testing for bias.[6] However, potential differences in attorney skill rule out the use of such a test to measure bias against attorneys. A skillful attorney would increase the probability of pretrial release but would not affect a client's likelihood of pretrial misconduct. Therefore, even in the absence of bias, the pretrial misconduct rate of the marginal defendant does not have to be equal across attorneys. Instead, the marginal releasee of higher skilled attorneys should have a higher probability of misconduct.

I account for potential omitted variable bias by exploiting the random assignment of both the judge and the public defender to cases in Miami-Dade first appearance hearings.[7] At these hearings, the judge decides whether to release the defendant on their own recognizance or alter the predetermined bail amount. In these hearings, all defendants are represented by a public defender, while an assistant state attorney represents the state. On weekends and public holidays, a single judge and public defender are assigned, on a rotating basis, to the felony bail shift, and handle all bail hearings on that day. This double random assignment, of both the judge and public defender to cases, is key to my identification.

The random assignment of cases to attorneys removes the concern of attorney case selection, ensuring that the unobserved characteristics of cases and defendants are uncorrelated with attorney race and gender. The random assignment of cases to judges means that the distribution of unobserved defendant characteristics is the same across judges. This allows me to account for potential differences in ability by analyzing inter-judge variation in the racial and gender gaps in attorney release rates. If judges do not exhibit bias in their treatment of attorneys based on demographics, these gaps would only reflect differences in mean ability across public defenders of different groups. Therefore, they should not vary significantly across judges. Thus, in this setting, inter-judge variation in the racial and gender gaps in attorney release rates indicates that some judges treat attorneys differently depending on their demographic group.

I initially test for homophily, *i.e.*, whether judges systematically favor attorneys of their own racial/gender group. I group judges by race and gender to test whether, on average, different groups of judges vary in their treatment of marginalized attorneys. After accounting for multiple hypothesis testing,[8] I do not find statistically significant evidence of bias in the treatment of Black, Hispanic, or female attorneys. This suggests that homophily does not play a substantial role in the treatment of attorneys in my setting. However, this test cannot rule out judicial bias against minor-

[6]Marginal outcome tests detect bias by comparing the expected effects of a decision across the marginal individual of each group. The marginal individuals are those for whom the decision-maker is indifferent in their decision.

[7]Many prior studies have exploited the random assignment of cases to judges to test for judicial bias against defendants (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. This is the first paper to exploit double random assignment in the criminal justice setting.

[8]This is done using the Benjamini-Hochberg procedure.

ity and female attorneys since offsetting bias by judges of the same race and gender may obscure potential bias by individual judges. I then develop a model of the judicial decision process incorporating the attorneys' role in the hearing. The model serves two main purposes. First, the model formally shows how judicial bias results in inter-judge variation in attorney racial and gender gaps. Second, I use it to simulate data to verify the validity of the finite sample distribution I construct to test for statistical significance.

In this model, a judge must form a posterior belief of a defendant's potential for pretrial misconduct based on the information they receive in the hearing. This information is represented by a signal that is influenced by the true misconduct potential of the defendant and the ability of the two attorneys involved in the hearing. Unconscious judicial bias against marginalized attorneys may reduce the persuasiveness of attorneys from these groups. This is incorporated through a modifier to the signal that a biased judge receives from the hearing, conditional on whether it involves an attorney of a demographic group against which the judge is biased. The judge then forms a posterior belief over the defendant's misconduct potential and releases the defendant if the expected cost of doing so is lower than the expected benefit. This model highlights that the mean level of bias across all judges is not separately identified from differences in mean ability across attorney demographic groups. Therefore, inter-judge variation is consistent with some judges being biased against marginalized attorneys or some (other) judges being biased in favor of these marginalized attorney groups.[9]

Informed by my theoretical model, I develop an identification strategy to test for inter-judge variation in the treatment of Black, Hispanic, and female attorneys. For each judge, I estimate the difference between the rate at which Black, Hispanic, and female defense attorneys secure the release of their defendants and the release rate of White male defense attorneys arguing in front of them. The inter-judge variation in these gaps represents my main empirical object of interest. I control for judge leniency and attorney ability through judge and public defender fixed effects.[10]

To test for statistical significance in the variation in attorney racial and gender gaps across judges, I build upon the method in Abrams, Bertrand & Mullainathan (2012).[11] The asymptotic distribution is not reliable for the estimated parameters because of the limited number of cases each judge hears which involve a given attorney demographic group. I address this by developing a finite sample test using Monte Carlo simulation to construct the distribution of these gaps under the null hypothesis of no bias.

I find statistically significant variation in the Black vs. White racial gap in public defender release rates across judges. This excess variation exists even when considering just the behavior of White male judges. Thus, it does not appear to be due to differences in the conduct of Black attorneys when arguing in front of judges of certain races. Therefore, I conclude that this variation is evidence of bias whereby some judges treat Black attorneys differently than White attorneys.

[9]This limitation is common to studies that utilize inter-judge variation in decisions to detect bias. See *e.g.*, Grossman et al. (2016) and Kastellec (2021). Sloan (2019) also acknowledges this limitation in identifying prosecutorial bias.

[10]Due to the substantial variation in the number of observations per judge, I employ empirical Bayes estimators to account for the fact that judges who heard fewer cases will have greater volatility in the estimates for their racial and gender gaps.

[11]In that setting, the random assignment of cases to judges allows them to detect judicial bias in the treatment of minority defendants from the variation in sentencing outcomes across judges.

However, I do not find evidence of statistically significant bias in the treatment of either Hispanic or female attorneys in Miami-Dade.

The magnitude of the bias I observe is such that when ranking judges from most negative toward Black attorneys to most positive, a defendant represented by a Black public defender is 2.8 percentage points less likely to be released if their case is heard by a judge at the 25th percentile of bias compared to a judge at the 75th percentile.

While the randomness that judicial bias introduces to decisions is of itself undesirable, I also quantify the potential effect of the observed bias against Black attorneys on racial inequality in defendant outcomes. I first calculate the proportion of defendants of each racial group represented by attorneys of each race and gender in a setting where defendants can choose their attorneys.[12] This is used as a proxy for defendant preferences to simulate the attorney-defendant matches that would occur in my dataset if defendants could choose their defense attorney. Under these simulated matches, I find that increasing judges' favorability toward Black attorneys by one standard deviation of the estimated relative favorability would reduce the defendant racial gap in pretrial release rates by 33% if defendants picked their attorney.

In analyzing the potential for judicial bias against attorneys from marginalized demographic groups, this paper adds to the literature on judicial bias. A substantial body of literature — in law, as well as economics and other social sciences — has disproven the societal ideal of judges as impartial, objective decision-makers (Harris & Sen 2019). Rachlinski et al. (2008) show, using the Implicit Association Test,[13] that judges appear to hold the same biases as the general population. 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 studies have attempted to answer whether bias against attorneys impacts the outcome of cases. Due to the empirical difficulties imposed by omitted-variable bias, the existing research generally relies on evidence from lab experiments (see *e.g.*, Cohen & Peterson 1981, Hahn & Clayton 1996, Hodgson & Pryor 1984, Sigal et al. 1985, Wood et al. 2019). For example, Hodgson & Pryor (1984) conduct a lab experiment in which participants, college students, listened to a mock trial where the role of the defense attorney was played by either a male or female actor. They find that the female "attorney" was considered less credible, with substantially lower scores on metrics related to competence. In addition, the female attorney's client was more likely to be found guilty. This underscores the potential for attorney characteristics to impact the assessment of attorney credibility and thus, case outcomes.

However, lab experiments testing for bias in trial settings have several shortcomings. In particular, research has shown that the extent of reliance on heuristic cues such as credibility varies depending on the importance of the task and the effort that the decision-maker is applying to it (Chaiken 1980). Therefore, the impact of bias on "verdicts" reached in experiments where the participants have little incentive to arrive at the correct decision may not have external validity. In a real courtroom setting, because decisions have a major impact on an individual's life, decision-

[12] This is calculated using data from Bexar County. This is the same data used by Agan et al. (2021).

[13] The Implicit Association Test is designed to detect subconscious associations by testing the strength of an individual's association between certain concepts and evaluations.

makers may be less reliant on heuristics.[14] By using data from actual hearings, this study circumvents this concern.

This paper is most similar to Szmer et al. (2010) and Szmer et al. (2013) which investigate gender bias by appellate court judges against female attorneys and Chen et al. (2017) which examines the effect of attorney vocal characteristics on Supreme Court outcomes. However, the ability of these studies to identify bias is limited. Szmer et al. (2010) and Szmer et al. (2013) account for neither potential unobserved differences between different attorney types nor potential unobserved differences in the types of cases that they choose to represent. While Chen et al. (2017) show an effect even after controlling for unobserved differences across attorneys, they acknowledge that they are limited in their ability to claim causation given that an attorney's tone is endogenous. This paper will causally identify the effect of bias by utilizing random assignment to account for potential confounding factors.

In addition, this paper also contributes to the labor discrimination literature, adding to our understanding of the factors that may give rise to employer beliefs about differences in ability between minority and majority workers. Several papers seek to explain how such beliefs could persist in the absence of inherent differences in ability.[15] Much of the research in this area has focused on the actions of (potential) employees in response to employer beliefs. Mechanisms that have been studied include lower skill investment (Coate & Loury 1993) and lower effort due to reduced supervision (Glover et al. 2017). In the current setting, judicial bias can directly reduce the productivity of attorneys, as measured by their ability to secure successful outcomes for their clients. This is comparable to the sports setting of Price & Wolfers (2010), which finds that the own-race preference of NBA referees directly affects the "output" of NBA players as measured through their points scored and the number of fouls committed. This paper is able to test for this effect in a setting where the implications of biased decisions are of greater societal concern.

The remainder of the paper is organized as follows: Section 2 provides an overview of the Miami-Dade pretrial system, describes the data, and presents the empirical tests of random assignment and the initial test of homophily by judges towards attorneys of the same demographic group. Section 3 introduces the theoretical model that I use to illustrate the potential mechanism by which bias against attorneys may affect case outcomes and to motivate my empirical test of bias. Drawing on the theoretical model, Section 4 outlines the identification method. Section 5 provides the results of my test for judicial bias and Section 6 concludes the paper.

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.

[14] Spamann & Klöhn (2023) highlight a secondary concern, showing that even within a lab setting, the decisions made by judges differ from those made by law students.

[15] See Bertrand & Duflo (2017) for a discussion of the literature relating to the persistence of these differential beliefs.

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, approximately 70% fail to pay this and will have a first appearance hearing (Goldkamp & Gottfredson 1988). The hearings are quite short, lasting just 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 are required to 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 attorney "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 the set of potential confounding factors. Second, the short nature of the hearing means there is more scope for heuristics to play a role in the outcome.[16] In addition, while the determination of guilt is the paramount decision in a criminal case, the prevalence of plea bargains in the U.S. means that this decision is not made in most cases. Meanwhile, first appearance hearings occur in the majority 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.[17]

Random assignment of cases to judges and public defenders is a crucial component of my identification strategy. On weekends, trial judges cover the bail shift on a rotating basis to ensure roughly balanced caseloads throughout the year. Each Saturday, Sunday, and federal holiday beginning at 9:00 AM, a single judge will work the felony first appearance shift hearing all felony first appearances.[18] Similarly, the public defender's office uses a rotation system to assign a public defender to handle the first appearance hearings on these days. The allocations are made separately for judges and public defenders. Importantly, the public defender is assigned without knowledge of 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. The exploitation of random assignment in my identification strategy is explained in more detail in Section 4.

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

[17]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 although Didwania (2020*b*) does not find an effect on conviction probability for federal felony defendants.

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

2.2 Data

The main data used in this analysis is administrative data on all felony cases from the Miami-Dade County Clerk of Courts between 2006 and 2020. This consists of $943,020$ distinct cases. I restrict the data in several ways. First, I drop all cases without a first appearance hearing. Next, I drop all cases from 2017 onwards as, after this date, the Miami-Dade Public Defenders Office changed its process for handling pretrial release.[19] I further restrict the period of my estimation sample to cases from January 2009 onwards. This allows me to calculate pretrial misconduct in the three years before the defendant's first appearance hearing for all defendants. In addition, as my identification will rely on the quasi-random assignment of both bail judges and public defenders, only cases where the first appearance hearing occurred on a weekend or federal holiday are considered. Finally, I restrict attention to attorneys who are either "White", "Black", or "Hispanic" as the remaining racial groups are too small to analyze separately. Following these exclusions, $37,171$ cases remain. These cases involve 153 unique public defenders and 164 unique judges. The data contains detailed demographic information about the defendant, including race/ethnicity, gender, date of birth, and residential zip code. Case details include the charge, offense type, and case disposition. For each hearing, the data includes the date and the identity of the judge and attorneys involved. Finally, the data includes whether the defendant was released pretrial and, if so, the amount that the defendant posted and any additional non-monetary conditions.[20]

As the courts do not store any data on the demographics of the judges or attorneys involved in cases, I use a multi-step process to determine these characteristics. First, I matched attorney names against the Florida Electoral rolls for Miami-Dade County. If an exact match existed, I took the self-reported race, ethnicity, and gender from the rolls.[21] Next, for attorneys and judges who are not matched to the electoral rolls, I used machine learning tools to predict race, ethnicity, and gender from their names. I assign the predicted race/gender if the tools estimate at least 65% accuracy to the prediction. Finally, a research assistant manually searched for each individual on public platforms such as LinkedIn and other websites such as those of private law firms.[22] The research assistant recorded their classification of race, gender, and ethnicity from pictures and other biographical data made available on these sites.[23] Table 1 reports the number of public defenders, judges, and prosecutors whose race/sex is assigned through each step. In this paper, "White" is defined as non-Hispanic White, and "Hispanic" is defined as Hispanic White. "Black" is defined as all Black individuals regardless of ethnicity. Throughout my analysis, I consider Black and White Hispanic individuals as separate groups rather than as a single "racial minority" group to allow for potential differences in their treatment. Eckstein (2014) discusses the unique position of the

[19] In 2017, the Public Defenders Office established an Early Representation Unit, which would continue to 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.

[20] Because ability to pay is a key consideration in the determination of bail, without information on the defendant's financial resources I cannot directly use the amount at which bail was set as an outcome.

[21] In the case that multiple matches are found for a name, then race/sex is assigned from the voter rolls if more than 90% of the potential matches are of the same race/sex.

[22] At the time this was completed, the research assistant was not aware of the research question in this study, and they did not have access to any of the data for this study other than the attorney and judge names.

[23] This is predominantly used in the classification of "White"/"Black" attorneys as the machine learning tool used was less able to predict race than ethnicity.

Cuban-American population in Florida. She 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.[24] Given this status as an "in-group" and the difficulty of separating them from other Hispanic individuals who may not share this status, my results for the Hispanic ethnic group may not represent the experience of the broader Hispanic community.

Finally, I use the 2017 American Community Survey for data on Food Stamps/SNAP eligibility, local unemployment rates, and education at the 5-digit Zip Code Tabulation Area level. These data are matched with the defendant's address for use as additional controls in the regression analysis.

Table 1: Assignment Method for Judge/Attorney Demographics

	Defender		Judge		Prosecutor	
	Race	Sex	Race	Sex	Race	Sex
Voter Rolls	103	101	83	82	204	189
Machine Learning	43	53	58	80	77	102
Manual	7	0	22	1	13	3

Note: This table reports the number of attorneys and judges whose race/sex were determined through each method.

Table 2 reports summary statistics for the public defenders in my estimation sample separated by race, while Table A4 reports statistics broken out by gender. [25] The final column reports the p-value from an F-test of equality of means across the three racial groups for each relevant statistic. Black attorneys make up less than 15% of public defenders, while Hispanic attorneys constitute just under 25%. In addition, the majority of public defenders are female. Defendant and case characteristics are very similar across all demographic groups, which is consistent with the inability of attorneys to select their cases in this setting. Male defendants make up roughly 85% of defendants, and just over half of defendants are Black. Roughly 11% of cases involve at least one first-degree felony.

Around 17% of cases involve a defendant who has previously engaged in pretrial misconduct (defined as being arrested for another felony while on pretrial release) within the past 3 years.[26] While the variation in judge characteristics across public defenders is greater than that of defendant characteristics, it must be noted that because the same judge and attorney pair will handle almost all first appearance hearings on a given day, the standard errors are substantially larger for the attorney characteristics than the defendant characteristics. Therefore, these differences are not statistically significant. As discussed in Section 4, the public defender who will cover a given weekend shift is assigned before the assigned bail judge is publicly known. Hence, it does not appear that there is any scope for selection between judges and attorneys.

[24]Ojito (2000) expresses a similar view of White Cuban males as representing a "majority" group within Miami-Dade.

[25]As mentioned above, public defenders who are neither "Black" nor "White" nor "Hispanic" are removed from the estimation sample as they do not form a homogeneous group that could be analyzed together, and each of these individual groups is too small to be analyzed separately.

[26]Due to data limitations, I cannot include arrest for a misdemeanor or failure to appear to a scheduled hearing.

Black attorneys are successful in 3% less of their cases than Hispanic attorneys, and the difference in release rates across attorney races is statistically significant at the 5% level. Although this may suggest judicial bias against Black attorneys, it could instead be due to differences in skill level across attorney demographic groups. Meanwhile, release rates are roughly equal between male and female attorneys.

Tables 3 and A3 report the analogous summary statistics for the judges in my estimation sample. White judges are overrepresented relative to their population share in Miami-Dade, while Black judges are substantially underrepresented. Overall, the judges are split evenly between males and females, although Black females make up a particularly small proportion of judges. The characteristics of the cases heard are once again very similar across the different racial groups, which is to be expected given the quasi-random assignment of judges to bail shifts. [27] In terms of leniency, there is little variation between judges at the race or gender level, with each racial group of judges releasing approximately 58% of defendants. However, this hides substantial variation in leniency between individual judges. Figure 1 shows the distribution in the pretrial release rate of judges who presided over at least 30 weekend hearings in my dataset.

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 marginalized group of attorneys and White male attorneys, which I refer to as the release rate "gap." Figures 2a, 2b and 2c show the distribution in these gaps across individual judges while Table 4 reports some summary statistics. These demonstrate substantial variation across judges in each of these gaps for each of the three demographic groups. For example, ranking judges by their Black vs. White attorney release rate gap, Black public defenders secured the release of 7.5% less of their defendants than White attorneys in front of the 25$^{\text{th}}$ percentile judge. Meanwhile, they secured 8.5% more in front of the 75$^{\text{th}}$ percentile judge. While this variation may be indicative of judges treating public defenders differently based upon their demographic characteristics as discussed in Section 4, it may merely reflect sampling variability. Hence, in Section 5, I formally test whether this variation is statistically significant.

[27] Some of the small differences across judge race appear to be statistically significant at the 5% level when not accounting for multiple hypothesis testing. However, when formally testing for random assignment in Section 2.3, I fail to reject random assignment. Thus, these differences likely are due to time trends, which are controlled for in my main analysis.

Table 2: Case Characteristics By Public Defender Race

	Black	Hispanic	White	p-value
N	20	35	98	
Proportion Male	0.45	0.34	0.47	
Cases	3662	8486	25044	
Male Defendant	0.85	0.84	0.85	0.218
Black Defendant	0.54	0.54	0.54	0.652
First Degree Felony	0.11	0.11	0.11	0.625
Prior Misconduct	0.18	0.16	0.17	0.051
Male Judge	0.49	0.5	0.49	0.962
Black Judge	0.1	0.08	0.08	0.686
Hispanic Judge	0.2	0.31	0.29	0.131
White Judge	0.68	0.62	0.62	0.281
Pretrial Release	0.57	0.6	0.58	0.020
Misconduct	0.1	0.09	0.09	0.795

Note: This table provides summary statistics for the Public Defenders who conducted first appearances in Miami-Dade between January 2009 and December 2016, broken out by the Public Defender's race. p-value denotes the p-value from an F-test of whether the row variable has explanatory power in predicting the race of the Public Defender.

Table 3: Case Characteristics By Judge Race

	Black	Hispanic	White	p-value
N	14	48	101	
Proportion Male	0.79	0.35	0.53	
Cases	3157	10404	23213	
Male Defendant	0.84	0.84	0.85	0.049
Black Defendant	0.52	0.55	0.53	0.018
First Degree Felony	0.11	0.11	0.11	0.898
Prior Misconduct	0.15	0.17	0.17	0.171
Male Attorney	0.54	0.51	0.51	0.794
Black Attorney	0.12	0.07	0.11	0.146
Hispanic Attorney	0.21	0.24	0.23	0.876
White Attorney	0.67	0.68	0.67	0.800
Pretrial Release	0.6	0.59	0.58	0.008
Misconduct	0.09	0.1	0.09	0.132

Note: This table provides summary statistics for the judges who conduct first appearances in Miami-Dade between January 2009 and December 2016, broken out by the judge's race. p-value denotes the p-value from an F-test of whether the row variable has explanatory power in predicting the race of the Judge.

Figure 1: Release Rates by Judge

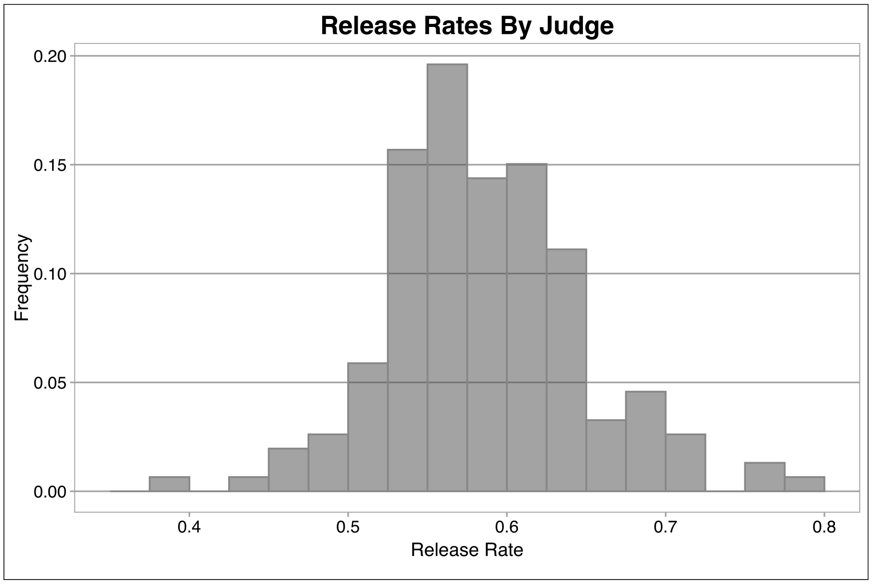

Note: Histograms of the distribution of release rates by judge. Only judges who preside over at least 30 pretrial hearings are included.

Figure 2: Attorney Demographic Group Gaps by Judge

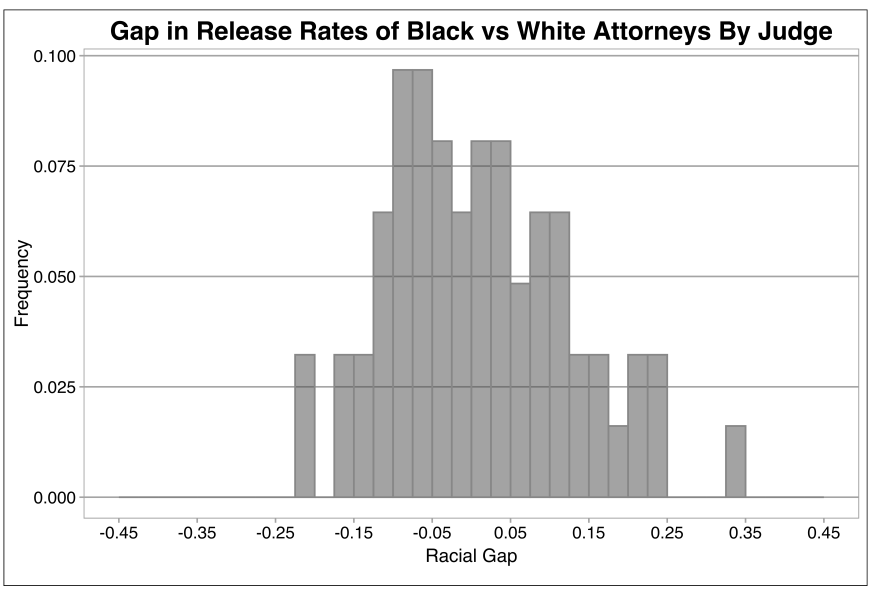

(a) Black vs. White Public Defenders

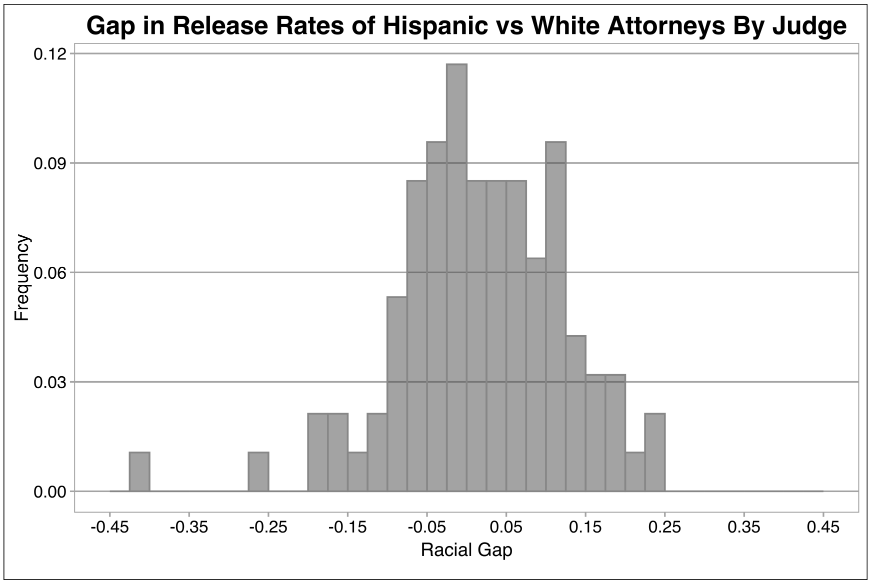

(b) Hispanic vs. White Public Defenders

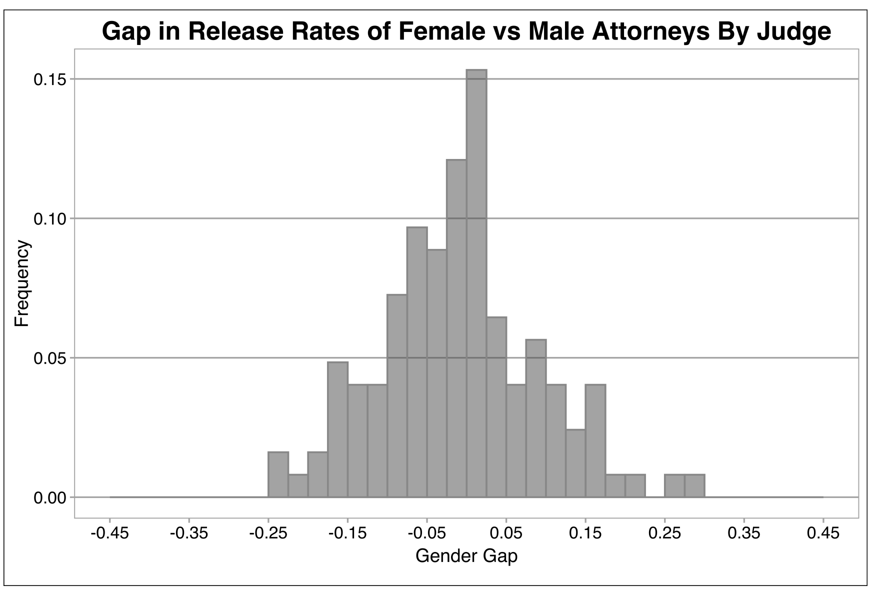

(c) Female vs Male Public Defenders

Note: 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 20 pretrial hearings involving public defenders of the relevant marginalized demographic group and 20 involving White male public defenders are included.

Table 4

	Black	Hispanic	Female
IQR	0.161	0.123	0.115
σ	0.118	0.106	0.100
Range	0.548	0.648	0.538

Note: 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 20 pretrial hearings involving public defenders of the relevant marginalized demographic group and 20 involving White male public defenders are included.

2.3 Empirical Test of Random Assignment

The observable case characteristics shown in Tables 2 and A4 were very similar across attorneys of different races and genders. Further, it does not seem feasible for any of the actors to influence the hearing date as it must occur within 24 hours of the arrest. Hence, it is extremely unlikely that there is a correlation between public defenders and the characteristics of the defendants or their cases. Given that this is central to my identification of bias, I formally verify this for observable characteristics. As discussed above, I compare the release rates of different attorney demographic groups. Therefore, it is necessary that, on average, attorneys of different groups represent similar sets of defendants. To test this, I regress case characteristics on dummies for public defender race and gender as well as time dummies to test whether the race and gender of the attorney involved in a case correlate with these characteristics.

$$\begin{aligned} Characteristic_i = &\alpha + \beta_1 B_i^{PD} + \beta_2 H_i^{PD} + \beta_3 F_i^{PD} \\ &+ month_t + year_t + holiday_t + \varepsilon_i, \end{aligned} \tag{1}$$

where $Characteristic_i$ denotes the relevant observable case/defendant characteristic in case i. B_i^{PD} (H_i^{PD}, F_i^{PD}) is a dummy for whether defendant i is represented by a Black (Hispanic, female) public defender. I include time dummies for month, year, and whether the first appearance hearing was on a public holiday.

If attorneys of different races represent similar defendants, the coefficients on β_1, β_2, and β_3 should equal zero. This would mean that the race and gender of the public defender have no explanatory power for each observable case characteristic tested. I test this using a partial F-test, with the results shown in Table 5. I test two demographic characteristics of the defendants (race and gender) and three case characteristics (whether the defendant has previously engaged in pretrial misconduct, the number of charges that are first-degree felonies, and the number of charges that are life felonies). I fail to reject the null hypothesis that the public defender demographic group does not explain the variation in defendant characteristics at any reasonable level of significance for all of these characteristics.

Table 5

	F	Pr(>F)
Defendant Male	0.909	0.435
Defendant Black	0.961	0.410
Prior Misconduct	0.530	0.662
# of First Degree Felonies	0.646	0.585
# of Life Felonies	0.036	0.991

This table reports the results of the regression in Equation 1.

In testing for random assignment of judges to bail shifts, I test whether a judge's identity has explanatory power in predicting either defendant or attorney characteristics. I estimate Equation 2 to do this. This is similar to Equation 1. However, while above I was able to aggregate attorneys by race/gender, here, I estimate a dummy for each judge. This is because I will compare the gaps in release rates of attorney demographic groups across individual judges. Therefore, I must ensure that each judge sees a similar pool of defendants. This is in contrast to the former test, where I only needed to ensure that, on average, each group of public defenders represented a similar pool of defendants.

$$Characteristic_i = \alpha + \delta_j + month_t + year_t + holiday_t + \varepsilon_i, \tag{2}$$

I estimate the set of $\{\delta_j\}_{j=1}^{J}$ for each defendant characteristic. δ_j estimates, conditional on time, the difference in the mean value of the defendant characteristic across the cases of judge j relative to the mean for the cases of the omitted judge. Under random assignment, all of these judge dummies should be approximately equal to zero. However, like in Abrams et al. (2012), the asymptotic F-test may not be appropriate to test whether these dummies are jointly equal to 0. This is because, while the overall number of cases is large, the number that any single judge presides over is substantially lower. As a result, an F-test would suffer from finite sample bias.[28] Hence, to test the null hypothesis that defendant characteristics are the same across all judges, I use a Monte Carlo simulation to generate the finite sample distribution of my test statistic, the Interquartile Range (IQR) of the coefficients, under the null hypothesis that defendants are randomly assigned to judges. I generate 5,000 simulations. In each simulation, starting from the true dataset, I randomize the allocation of cases to judges.Equation 2 is then estimated on each simulated dataset and the IQR of $\{\delta_j\}_{j=1}^{J}$ is separately computed for each defendant characteristic. For each characteristic, I then compare the IQR from the true dataset to the distribution of the IQRs from the simulated data to determine the probability of getting an IQR at least as extreme under the null of random assignment.[29] The simulated p-value for each characteristic is reported in Table 6. Once again, I fail to reject the null of random assignment at any standard level of significance.

[28]See Abrams et al. (2012) for a discussion of the potential for overrejection of the null hypothesis when using an F-test in a similar setting. In Appendix A, I discuss a related concern of finite sample bias in the application of my main empirical test of bias

[29]This is a modification of the Monte Carlo method employed by Abrams et al. (2012).

Table 6

	IQR	p-value
$IQR_{DefendantMale}$	0.045	0.180
$IQR_{DefendantBlack}$	0.063	0.317
$IQR_{LifeFelony}$	0.026	0.467
$IQR_{FirstDegreeFelony}$	0.068	0.589
$IQR_{PriorMisconduct}$	0.065	0.361

This table reports the IQR of the estimates of the judge dummies from the regression in Equation 2. This empirical IQR is compared to the distribution of IQRs estimated from 5,000 simulated datasets which are 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 IQRs that are larger than the one estimated in the true data.

2.4 Test of Homophily

To determine whether there is heterogeneous treatment of minority public defenders, I first examine whether judges exhibit homophily, a preference for attorneys of their demographic group. Tests of this form are widely used in papers that test for bias by looking at variation across decision-makers (*e.g.*, Grossman et al. 2016, Kastellec 2021, Sloan 2019). I test for this using the following linear probability model:

$$\begin{aligned} R_i = &\alpha + \beta_1(B_i^J \cdot B_i^{PD}) + \beta_2(B_i^J \cdot H_i^{PD}) + \beta_3(H_i^J \cdot B_i^{PD}) + \beta_4(H_i^J \cdot H_i^{PD}) + \beta_5(F_i^J \cdot F_i^{PD}) \\ &+ \delta_j + \gamma_d + \psi_p + \kappa X_i + year_t + month_t + holiday_t + B_i^D \cdot year_t + H_i^D \cdot year_t + \varepsilon_i, \end{aligned} \tag{3}$$

where R_i is a dummy for whether individual i was released pretrial. Judges and attorneys are grouped by race and gender with B_i^J (H_i^J, F_i^J) a dummy for whether i's bail was decided by a Black (Hispanic, female) judge and B_i^{PD} (H_i^{PD}, F_i^{PD}) is a dummy for whether i was represented by a Black (Hispanic, female) public defender. The coefficients of interest are β_1, β_2, β_3, β_4 and β_5. These measure the racial gap in release rates between Black vs. White and Hispanic vs. White public defenders across racial groups of judges and the gender gap between Female vs. Male public defenders by the judge's gender. Several sets of controls are included. Individual dummies for each judge (δ_j), public defender (γ_d), and prosecutor (ψ_p) are estimated. These control for variation in leniency across individual judges and skill across individual attorneys. In addition, I control for case/defendant characteristics, including gender, race, home zip code characteristics, and criminal history, which are represented by the vector X_i. I include time dummies for month, year, and whether the first appearance hearing was on a public holiday. Finally, I allow for time-varying treatment of minority defendants by interacting the year dummies with dummies for whether the defendant is either Black (B^D) or Hispanic (H^D).

Table 7

	Dependent variable:
	Pretrial Release
Black Judge * Black PD	0.004
	(0.058)
Hispanic Judge * Black PD	-0.040
	(0.028)
Black Judge * Hispanic PD	0.040
	(0.040)
Hispanic Judge * Hispanic PD	0.044
	(0.024)
Female Judge * Female PD	-0.009
	(0.014)
Mean	0.58
Observations	37,192
Adjusted R^2	0.124
Note:	*p<0.1; **p<0.05; ***p<0.01

This table reports the results of the regression in equation 3. Robust standard errors clustered at the Judge ID and Public Defender ID level. Significance thresholds are adjusted for multiple hypothesis testing using the Benjamini-Hochberg procedure.

The estimates for the coefficients of interest are shown in Table 7. As this is a linear probability model, the first estimate states that, on average, a Black judge is 0.4 percentage points more likely to release a defendant if they are represented by a Black public defender rather than a White public defender, conditional on the individual attorney fixed effects. However, this coefficient is not statistically significant. I account for the testing of multiple hypotheses using the Benjamini-Hochberg procedure.[30] Under this method, none of the interaction terms are statistically significant at the 5% level. Hence, I do not find evidence of homophily in this setting. Appendix Table A1 shows the results from an equivalent probit model. The results from this model are similar to those shown in the linear probability model. The average marginal effect of the interaction term between Black Judge and Black public defender is 0.7 percentage points and is once again statistically insignificant.[31]

However, while research has shown the impact of a judge's racial/ethnic identity on judicial biases (see *e.g.*,Gazal-Ayal & Sulitzeanu-Kenan (2010) and Grossman et al. (2016)), there is no reason to believe that bias within racial and gender groups would be uniform. As a result, 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. Therefore, while judge race and gender are predefined groupings that facilitate comparisons across judges, a null finding here is insufficient to conclude that there is no bias.

To test whether there is any bias by judges in the treatment of attorneys from marginalized groups, it is necessary to test for statistically significant variation across individual judges. To do so, I first develop a theoretical model that will allow me to formalize the bias that I will be testing for in Section 5.

3 Conceptual Framework

This section presents a theoretical model of the bail hearing. Section 4 then shows how it informs my empirical test of racial bias against attorneys. 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 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,[32] normalized such that he will engage in misconduct if $\theta_i < 0$. The defendant's type, θ_i consists of a deterministic component, $\bar{\theta}_{X_i}$, which depends on the characteristics of i observable by the judge (X_i)[33] 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:

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

[30]To calculate the Benjamini-Hochberg critical values, the p-values of each coefficient to be tested are ranked from smallest to largest. All null hypotheses such that $p_i < \frac{rank_i}{\#\,of\,tests} \times \alpha$ are rejected, where α is the chosen level of significant (Type-I error rate).

[31]Note that as discussed in Ai & Norton (2003), this is not equivalent to the "interaction effect." However, under the theoretical model, which I will outline in Section 3, it is the interaction term itself that represents bias, and thus, it is the marginal effect of this term in isolation that measures the effect of relative bias.

[32]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.

[33]This includes both demographic and case characteristics.

Therefore, the judge's prior 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 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: $q_d \sim \mathcal{N}(\bar{q}^d, \sigma_{q^d}^2)$ and $q_p \sim \mathcal{N}(\bar{q}^p, \sigma_{q^p}^2)$. [34] This allows a higher-quality attorney to increase the probability of a successful outcome. [35] 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 $\eta_i \sim \mathcal{N}(0, \sigma_\zeta^2 + \sigma_{q^d}^2 + \sigma_{q^p}^2)$. 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

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

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. As mentioned above, in determining pretrial release, judges must consider both the risk of a failure to appear and the commission of a new crime. 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 documented in the literature. [36] The

[34] Within the model it is assumed that q_p, q_d and ζ_i are independent. In Section 4, I will explain why this assumption is satisfied in my setting.

[35] 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.

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

perceived benefit is assumed not to vary by defendant characteristics. [37] The implications of relaxing this assumption are discussed in Section 4. 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

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

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 $0 \leq \frac{\Pi_j}{C} \leq 1 \, \forall \, j \in J$, *i.e.*, there is no judge which would never release a defendant nor one which would release every defendant.

This can be rearranged to give

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

where $\bar{b}^{r_d}$ is the mean level of bias held by all judges against attorneys of race r_d while $\tilde{b}_j^{r_d}$ is the deviation of judge j's bias from this mean level of bias.[38]

I describe my identification strategy in Section 4 below. This allows me to identify the idiosyncratic judicial bias component of Equation 5.

[37] 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.

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

4 Identification

I exploit the quasi-random assignment of judges and public defenders to first appearance shifts to test for bias against attorneys. In general, when comparing outcomes across attorney types, there are two key issues. First, there are potentially observed and unobserved differences in skills across attorneys (represented by $\tilde{q}_d$ in my theoretical model). If these differences are correlated with attorney race or sex, then this could justify a gap in attorney release rates across groups, even in the absence of bias. In addition, there could be unobserved differences in the cases argued by different attorney groups, which could drive differences in outcomes. For example, if Black attorneys were more likely to represent defendants with a history of pretrial misconduct or charges for serious violent offenses, they would likely, on average, have a lower release rate than White attorneys, even in the absence of judicial bias.

To identify judicial bias against minority and female attorneys, I compare the differences in release rates of public defenders from marginalized groups to those of White male public defenders across judges. I refer to these differences as "relative bias". Higher-skilled attorneys may be able to achieve more favorable outcomes for their clients, in this setting reflected through a higher proportion of clients obtaining pretrial release. Correlation between demographics and the skill of attorneys in my sample may explain an overall racial/gender gap in attorney release rates. However, if all judges treat attorneys equally regardless of race and gender, there should not be systematic variation in these gaps across judges. Therefore, excessive variation in the racial (gender) gaps in attorney release rates across judges would be evidence of bias by some judges in their treatment of attorneys based on their race (gender). This variation identifies the effect of $\tilde{b}_j^{r_d}$ in Equation 5 of the theoretical model. This method cannot identify the effect of bias that is common across all judges, in the model represented by $\bar{b}^{r_d}$, as it is not separately identified from attorney ability. However, for there to be no variation across judges in a setting where bias against attorneys affects case outcomes, the effect of bias would have to be equal across all judges, which seems implausible.[39] Thus, such a finding would provide strong evidence that bias against minority or female attorneys does not significantly impact case outcomes.

To identify each judge's relative bias, I exploit the random assignment of first appearance hearings in my setting.

The random assignment of cases to judges is necessary for the comparison of attorney racial and gender gaps across judges to be valid. This comparison requires that a given attorney's quality is constant across the different judges they argue before, *i.e.*, in Equation 5, $\tilde{q}_d$ for a given attorney d must not vary depending on the judge hearing the case. As attorneys also expend more effort or have particular expertise in representing certain types of clients, a correlation between judge identity and defendant characteristics would invalidate this method. However, given the random assignment of cases to judges, such a correlation cannot occur.[40] Note that the potential for public defenders to change their behavior depending on the judge's identity is discussed in detail in

[39]Papers such as Arnold et al. (2022), which have measured bias against defendants in similar settings, have found substantial variation in the level of this bias.

[40]Thus, while I do not explicitly model the potential correlation between X_i and q_d in Equation 5, relaxing this assumption would not impact the identification of $\tilde{b}_j^{r_d}$ but would instead become part of the error term.

Section 5.4.

Under my method, case selection by public defenders could affect the estimates of judicial bias if judges varied in their treatment of certain defendant/case characteristics. Assume, for example, that Black attorneys were more likely to represent defendants accused of drug crimes, and that some judges were stricter on drug crimes than others. Black attorneys would do worse in front of judges who were stricter on drug crimes than those who were more lenient. In the theoretical model, variation in a judge's decision based on certain defendant characteristics could be incorporated by relaxing the assumption that Π_j does not depend on defendant characteristics (*i.e.*, $\Pi_j(X_i)$).[41] Case selection could introduce correlation between X_i and attorney race thus creating correlation between an attorney's race and a judge's perceived benefit of releasing a defendant. However, because of the random assignment of cases to attorneys, there can be no correlation between attorney race and the defendant or case characteristics.[42]

In sum, the double random assignment of attorneys and judges to bail shifts ensures (1) no correlation between judge identity and defendant/case characteristics and (2) no correlation between attorney characteristics and defendant/case characteristics. The former allows me to make comparisons across judges. The latter ensures that observed differences in the racial gap across judges are due to their attitudes towards the attorneys themselves and not differential treatment of certain types of defendants.

It is worth noting that a marginal outcome based test of bias could not be used to answer the current question.[43] In a bail hearing, judges are tasked with setting bail based on the potential pretrial misconduct of the defendant. However, despite Supreme Court Justice Samuel Alito's claim that "it's the case, not the lawyer,"(**?**) it seems uncontroversial to consider that a more highly skilled attorney can secure better outcomes for their clients than a less skilled attorney.[44] Hence, a more skilled public defender should be able to secure the release of higher-risk defendants. However, the skill of the attorney at the first appearance hearing does not impact the probability that a given defendant will commit pretrial misconduct, *i.e.*, skill affects the release decision but not the misconduct probability. Thus, whether a public defender benefits from preferential treatment due to racial biases or has a higher skill level, the higher their release rate, the higher the misconduct probability of their marginal defendant. Therefore, using an outcome-based test to answer this question would be subject to the same confounder of unobservable differences in attorney skill as a naive comparison of attorney release rates. This is in contrast to studies of bias against defendants, such as Arnold et al. (2018), where the defendant's misconduct potential is the relevant unobserved "quality" that needs to be controlled for. In those studies, the unobserved quality affects both the judge's decision to release the defendant and the outcome of whether the defendant engaged in pretrial misconduct.

[41] This would allow judges to vary in their treatment of certain case characteristics and exhibit bias against defendants based on race or other characteristics.

[42] While judges may still vary in their treatment of certain defendants, as defendant characteristics are not correlated with public defender race, it will not bias the estimation of $\tilde{b}_j^{r_d}$. To the extent that there is variation across judges in the perceived benefit of the release of certain types of defendants, this will be absorbed into the error term.

[43] See Canay et al. (2020), Hull (2021), and Abrams et al. (2023) for a discussion of the use of these tests to detect bias.

[44] Abrams & Yoon (2007) finds substantial heterogeneity in client outcomes across attorneys. For example, they find that, on average, a public defender with 11 years of experience obtains sentences 17% shorter than a defender with 1 year of experience.

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.[45] 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.

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.[46] Because of estimation constraints, which I discuss in Section 5, I do not estimate prosecutor fixed effects. I instead control for prosecutor race, sex, law school tier, and experience.[47]

It should be noted that a non-linear estimator, in this case, a probit model, is used because a linear probability model may incorrectly attribute variation in judges' racial and gender gaps in attorney release rates caused by the interaction of variation in judge leniency and attorney ability to bias. To illustrate, consider two unbiased judges, j_1 and j_2. Assume that j_1 is relatively strict and, due to a quality difference between White and Black attorneys in the sample, releases 50% of defendants represented by White attorneys and 60% of those represented by Black attorneys. j_1 had

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

[46] 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.

[47] Given that prosecutors are also quasi-randomly assigned to bail shifts the effect of prosecutor quality and any potential racial bias against prosecutors should be absorbed within the error term.

an attorney racial gap of -10%. Meanwhile, j_2 is infinitely lenient, releasing 100% of all defendants, and thus has an attorney racial gap of 0. In a linear model, this could falsely appear to be evidence of bias. In contrast, a non-linear model accounts for the fact that the impact of bias on outcomes depends upon a defendant's base probability of release in the absence of bias.

5 Results

5.1 Inter-Judge Heterogeneity

I now proceed to estimate Equation 7.

$$\begin{aligned} R_i =&\alpha + \sum_{j=1}^{J} \beta_j^B (J_{ij} \times B_i^D) + \sum_{j=1}^{J} \beta_j^H (J_{ij} \times H_i^{PD}) + \sum_{j=1}^{J} \beta_j^F (J_{ij} \times F_i^{PD}) \\ &+ \delta_j + \gamma_d + P_i' \upsilon + \kappa X_i + year_t + month_t + holiday_t + \\ &+ B_i^D \cdot year_t + H_i^D \cdot year_t + \varepsilon_i \end{aligned} \tag{7}$$

This is similar to the model in Equation 3. However, instead of grouping judges by race, an indicator variable for each judge (J_{ij}) is interacted with a dummy for whether the defense attorney is Black, Hispanic, or female (B_i^{PD}, H_i^{PD}, and F_i^{PD}). The coefficients β_j^B, β_j^H, and β_j^F compare the racial and gender gap in attorney release rates for judge j against the gap of the omitted judge for each demographic group. I refer to these coefficients as the "relative bias" coefficients.

To control for variation in attorney ability, I again include public defender dummies alongside the attorney race and gender specific judge fixed effects. Akin to the estimation of employer-employee fixed effects, the individual judge and attorney dummies are only separately identified within a "connected set."[48] 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.

My test for bias involves the comparison across judges of the judge identity $\times$ attorney demographic fixed effects. Because the fixed effects are only identified within connected sets, comparisons between fixed effects are only valid within a connected set. Hence, the entire sample is used to estimate Equation 7, I restrict my comparison of individual judge dummies to the largest connected set for each demographic group of attorneys.[49] In constructing the connected sets, I restrict "con-

[48]The requirement of connected sets to separately identify firm and worker fixed effects is shown in Abowd et al. (2002).

[49]Note that judges must be connected both by attorneys of the relevant demographic group as well as by White male attorneys because attorney demographic specific judge fixed effects are relative to the judge fixed effect. Their leniency in cases represented by White male public defenders identifies the judge's fixed effect.

nections" between judges and White male attorneys to those that involve at least 20 cases between the relevant judge-attorney pair. For Black attorneys, this results in 18 connected judges, 34 for Hispanic attorneys, and 83 for female attorneys. As the number of cases each individual judge sees with a given attorney demographic group is relatively low and varies substantially across judges, I employ empirical Bayes estimators for the judge-specific attorney racial and gender gaps. This accounts for the fact that judges with fewer observations will have greater volatility in the estimated gaps.

To maximize the potential size of the connected groups, I no longer include individual prosecutor dummies.[50] Instead, I add dummies for prosecutor race, gender, experience tier, and law school rank. These dummies are denoted by the vector P_i.

To determine whether there is statistically significant variation in the gaps in defender release rates, I need to test whether these coefficients are statistically different from each other. As discussed in Appendix A, when considering the random assignment of defendants to judges, a likelihood ratio test (LRT) using the asymptotic χ^2 would suffer from finite sample bias. This is because of the relatively small number of cases at the individual judge $\times$ attorney demographic group level. Hence, I again use a Monte Carlo simulation to construct the finite sample distribution of my test statistic under the null of no judicial bias in the treatment of attorneys. My chosen test statistic for the variation across judges is once again the interquartile range. I compare the IQR of my estimated coefficients to this finite sample distribution to calculate the probability of the observed heterogeneity in the racial/gender gap in attorney release rates across judges under the null hypothesis that judges do not treat attorneys differently based on their demographic characteristics.

Figure 3a shows the histogram of the distribution of the IQR of β_j^B calculated from each simulation. The dashed line shows the IQR from the true dataset. Only 1.4% of simulated datasets yielded coefficients with more extreme variation than we see in the true data. Therefore, the variation in the treatment of Black attorneys relative to White attorneys exceeds what would be expected from sampling variation alone — *i.e.*, this variation is statistically significant. Under my theoretical model, I interpret this variation as evidence of relative bias in the treatment of Black attorneys, with some judges treating them more harshly than others. Turning to the racial gap between Hispanic and White attorneys, Figure 3b shows the histogram of the distribution of IQRs of $\{\beta_j^H\}_{j=1}^J$. Once again, the dashed line shows the IQR from the true dataset. In this case, over 25.3% of simulations had greater variation in the racial gap between Hispanic and White attorney release rates across judges. I cannot reject the null hypothesis that there is no variation in the treatment of Hispanic attorneys relative to White attorneys across judges in Miami-Dade. Figure 3c shows the histogram of the distribution of IQRs of $\{\beta_j^F\}_{j=1}^J$. Here 11% of simulations have a greater variation in the relative bias term than I observe in the true data. Therefore, at the 5% level, I cannot reject the null hypothesis that there is no variation across judges in the treatment of female attorneys relative to male attorneys. Given that I only have sufficient evidence to reject the null of equal treatment across judges in the case of Black public defenders, for the remainder of the paper I focus on them.

While the above analysis shows that there is statistically significant variation in the treatment

[50] The inclusion of prosecutor dummies would require that within a connected set judges are also connected by prosecutors.

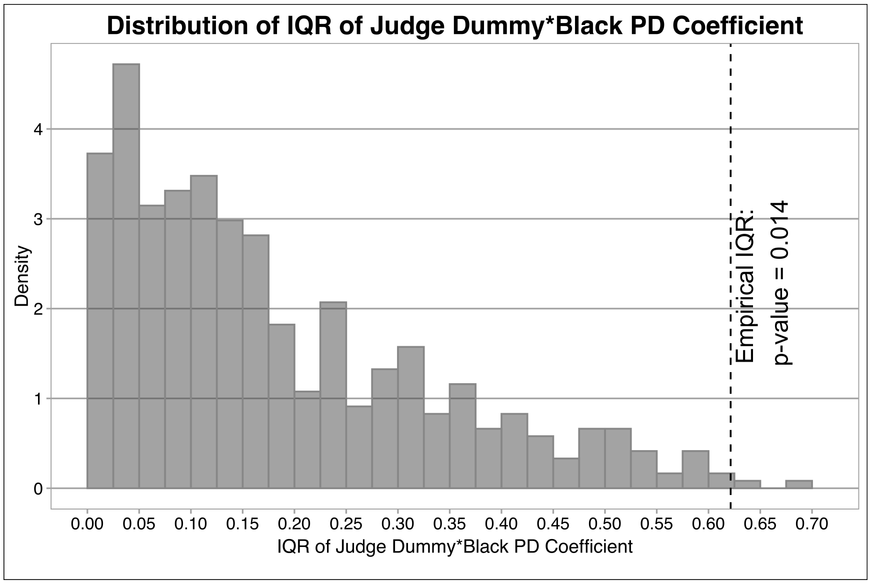

(a) Black vs. White Public Defenders

of Black attorneys, it is still necessary to determine whether the magnitude of this relative bias is "economically significant." Figure 4 shows the distribution of β_j^B, which under my model is each judge's relative bias towards Black public defenders.[51] I have converted the probit estimates to the average marginal effect for the subset of cases involving judges in this connected set to simplify interpretation. As these estimates reflect a judge's bias relative to any potential mean level of bias across judges, I have also demeaned the marginal effects. See Appendix Figure A1 for the distribution of the raw probit coefficients within this connected set. To provide a sense of the magnitude of this inter-judge variation, the difference in the probability of pretrial release for a defendant represented by a Black attorney is 2.8 percentage points lower when in front of a judge at the 25$^{\text{th}}$ percentile of relative bias compared to when in front of a judge at the 75$^{\text{th}}$ percentile, all else equal.[52] As a comparison, Arnold et al. (2022) estimate an IQR of roughly 5 percentage points for the unwarranted release rate disparity between Black and White defendants in the context of the NYC bail process.[53]

[51]Figures A2a and A2b in the appendix show the equivalent estimates for Hispanic and female attorneys.
[52]Note that this includes holding constant judge leniency and any bias they may have against the defendant.
[53]This is calculated from their estimate of a 3.1 standard deviation across judges in the unwarranted release rate disparity between Black and White defendants.

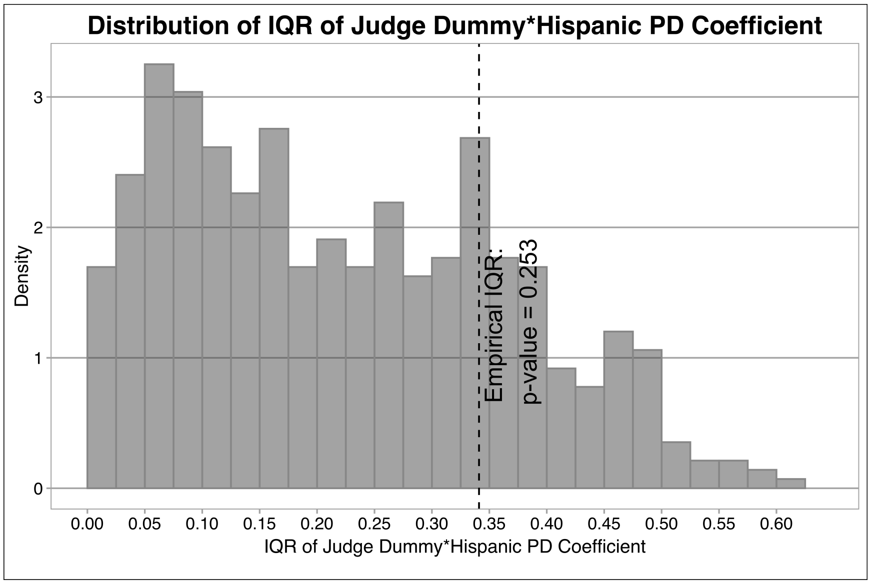

(b) Hispanic vs. White Public Defenders

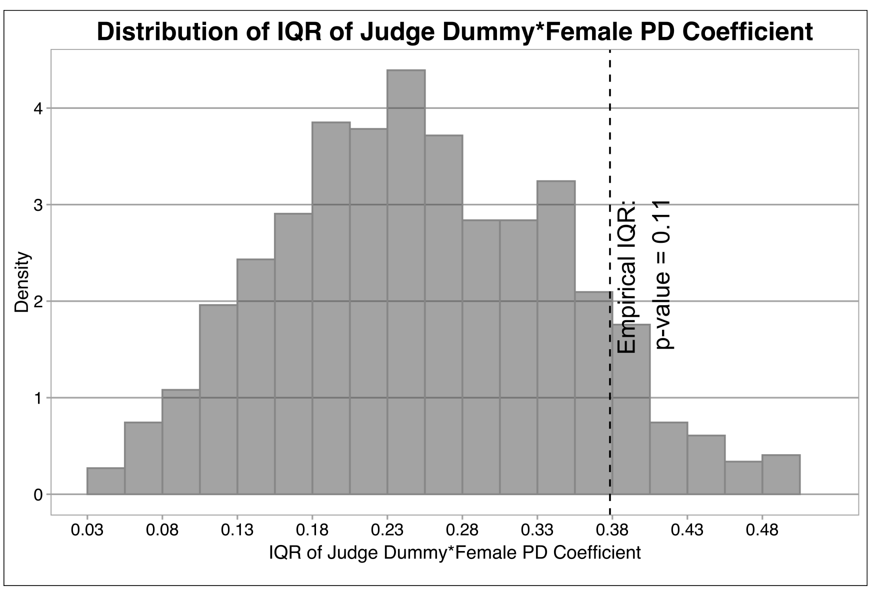

(c) Female vs Male Public Defenders

Figure 3: Density histograms of the distribution of the IQR of the estimates of the coefficient on the judge dummy × attorney demographic group dummy from Equation 7. IQRs are calculated on simulated data constructed under the null hypothesis that there is no bias against attorneys of the relevant demographic group. The dashed line marks the IQR of the estimated coefficients from the true dataset. (a) shows the distribution of IQRs of β_j^B, (b) the distribution of IQRs of β_j^H and (c) the distribution of IQRs of β_j^F.

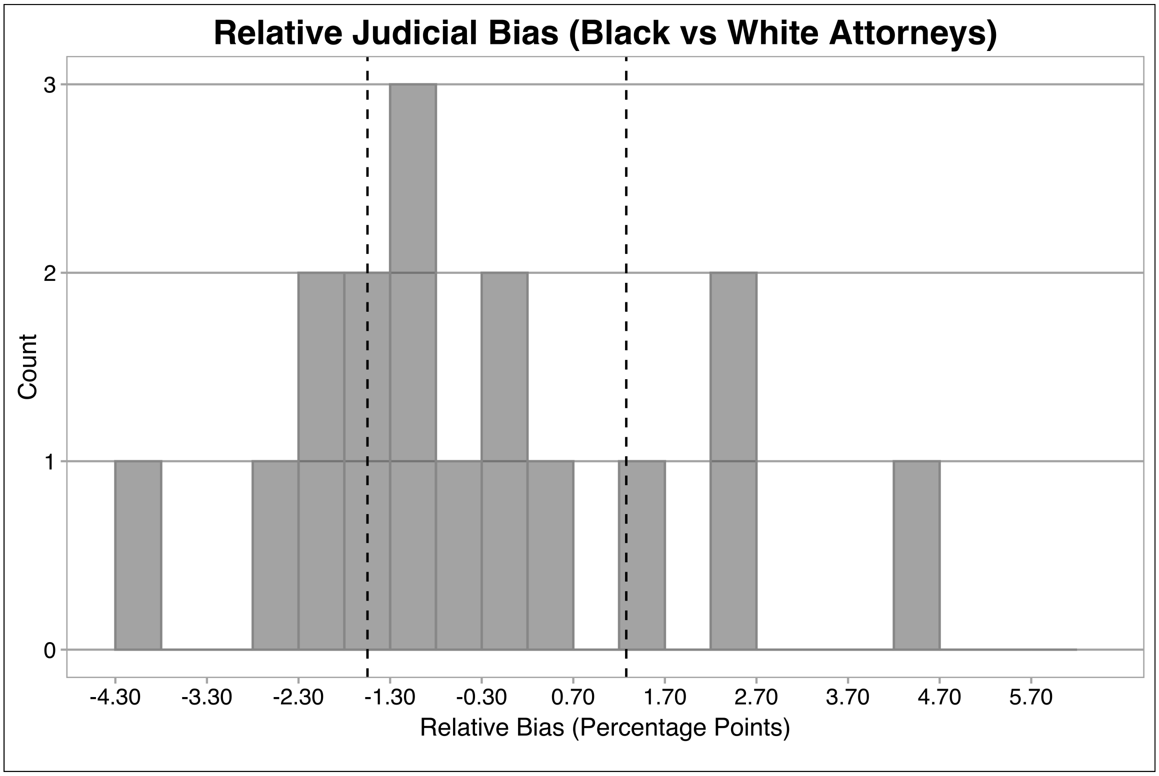

Figure 4: Histogram of the distribution across judges of the average marginal effects of the judge dummy $\times$ Black PD dummy (β_j^B) from Equation 7. The average marginal effect has been calculated across all cases involving a judge in the relevant connected set. This coefficient measures the gap in the rate at which Black male attorneys secure 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 20 pretrial hearings involving public defenders of the relevant marginalized demographic group and 20 involving White male public defenders are included.

I apply these estimates of the relative bias of judges against Black attorneys to determine whether idiosyncratic bias by some judges is a potential explanation for the gap in the rate at which Black attorneys secure the release of their defendants relative to attorneys of other demographic groups (*i.e.*, if the gap could be driven by a few "bad apples"). As shown in my theoretical model in Section 3, Equation 7 only allows me to identify relative bias and not the overall level of bias held by judges generally in Miami-Dade. This is because the level of bias common across all judges cannot be separately identified from attorney ability. However, while it cannot be concluded whether the gap is due to bias or differences in ability, I can test whether the idiosyncratic bias of a few judges is a potential cause. To do this, I simulate how many judges would need to change their behavior to close the racial gap in the release rates for Black and White public defenders and then for Black and Hispanic public defenders. I do this by replacing the judge-specific coefficient on the dummy for Black attorney for the judge which is harshest towards Black attorneys, with the value for the median judge. I continue doing this until the release rates for Black and White (Hispanic) attorneys are equal. As the gap between Black and White attorneys' release rates was initially only 0.5%, removing just the judge who has the most negative relative bias is sufficient to close this gap. The gap between Black and Hispanic attorneys is substantially larger at 2.1%. Even after removing all 11 judges in this connected set whose relative bias term is below the mean, a gap of 0.2% remains. Therefore, it seems unlikely that the bias of a few judges drives this racial gap. Instead, it appears to result from a widespread bias against Black attorneys common to all judges or a difference in the distribution of ability across the two racial groups in this sample.

5.2 Impact on Defendant Outcomes

Agan et al. (2021) find that minority defendants are significantly more likely to choose an attorney of their own race when given a choice, with Black defendants being 2.15 times more likely to retain a Black attorney than under random assignment.[54] 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).[55] Thus, Black defendants would primarily bear the impact of bias against Black attorneys.

I estimate the potential impact of judge 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.[56] Using these preferences, I simulate matches between defendants and attorneys in my dataset of Miami-Dade cases.[57]

Under these simulated matches, the pretrial release rate gap between Black and White defendants is 3.6%.[58] 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, if any, judge is unbiased in their treatment of Black attorneys,[59] I estimate the defendant racial gap at several quantiles of the relative bias coefficient.

Table 8 shows the results of the simulation exercise. If all judges had the same relative bias coefficient as the 25$^{\text{th}}$ percentile judge, which has an average marginal effect of -1.5 percentage points, Black defendants would have a 4.2 percentage point lower release rate than White defendants. This gap would decrease to 3 percentage points if all judges had the same level of bias as the 75$^{\text{th}}$ percentile judge, a decline of 28.5%. An increase in the favorability of all judges toward Black attorneys by one standard deviation of the relative bias coefficient would similarly decrease the racial gap in release rates by 33%.

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

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

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

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

[58] 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.

[59] 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.

Table 8

Bias Quantile	Bias Coefficient	B-W Release Gap
Baseline	–	-0.036
10th	−0.024	-0.046
25th	−0.015	-0.042
50th	−0.008	-0.039
75th	0.013	-0.030
90th	0.031	-0.020

Note: This table provides the forecast gap in the release rate of Black defendants and White defendants under the simulation in section 5.2. A negative value means that a higher proportion of White defendants are released than Black defendants. "Bias Coefficient" denotes the value to which each judge's bias coefficient is set and is expressed as an average marginal effect.

5.3 Intersectionality

In analyzing the potential role of bias in Section 5, 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. [60] Under the concept of "multiple jeopardy," King (1988) asserts that the disadvantages of race and gender are often compounded. This theory suggests, for example, that Black women face greater disadvantages than the addition of the two statuses would imply.

I cannot incorporate intersectionality into my analysis of individual judge heterogeneity due to data limitations. The requirement of a connected set, as discussed in Section 5.1, means that in this case, the number of judges across whom I could compare judge-specific coefficients would be too small to draw any inference. Given this, I again group judges by demographic characteristics and test for variation across these groups. To do so, I estimate the below linear probability model, which is a modified version of Equation 3

$$\begin{aligned} R_i =& \alpha + \beta_1(B_i^J \cdot B_i^{PD} \cdot F_i^{PD}) + \beta_2(F_i^J \cdot B_i^{PD} \cdot F_i^{PD}) + \beta_3(B_i^J \cdot B_i^{PD}) \\ &+ \beta_4(B_i^J \cdot F_i^{PD}) + \beta_5(F_i^J \cdot B_i^{PD}) + \beta_6(F_i^J \cdot F_i^{PD}) \\ &+ \delta_j + \gamma_d + \psi_p + \kappa X_i + year_t + month_t + holiday_t + B_i^D \cdot year_t + H_i^D \cdot year_t + \varepsilon_i, \end{aligned} \tag{8}$$

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 interaction terms between the judge and attorney demographic groups are also included, as well as the controls from Equation 3.

The results from estimating Equation 8 are shown in Table 9. Neither intersectionality term is statistically significant at the 5% level. However, the coefficient on Black Judge * Black Female PD is very large in magnitude and close to my chosen level of significance, with a p-value of 0.08. This estimate is noisy and insignificant at the 5% level as the small number of both Black female public defenders and Black judges means that the standard errors are large. Nevertheless, this does appear to highlight the need for further analysis in a setting where a larger sample size makes it possible to estimate the effect more precisely.

[60] See Browne & Misra (2003) for an overview of various theories on the operation of intersectionality in the labor market context.

Table 9

	Dependent variable:
	Pretrial Release
Black Judge * Black Female PD	0.170
	(0.089)
Female Judge * Black Female PD	−0.023
	(0.051)
Black Judge * Black PD	−0.033
	(0.042)
Female Judge * Black PD	−0.072
	(0.041)
Black Judge * Female PD	0.060
	(0.028)
Female Judge * Female PD	0.013
	(0.017)
Mean	0.58
Observations	36,753
Adjusted R^2	0.141
Note:	*p<0.1; **p<0.05; ***p<0.01

This table reports the results of the regression in Equation 8. Robust standard errors clustered at the Judge ID and Public Defender ID level.

5.4 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 a patient's illness.[61] In my context, there are

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

many potential reasons why a Black attorney may behave differently depending upon 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 only 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, it seems less likely that behavior should vary greatly before judges of the same race and gender. This seems particularly true in the current context, as attorneys have limited experience with the judges in these bail hearings. 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 6. Looking just at the variation in the racial gap between Black and White attorneys, I continue to find statistically significant variation, with only 2% of simulated datasets having a larger IQR for the coefficients of White male judges.

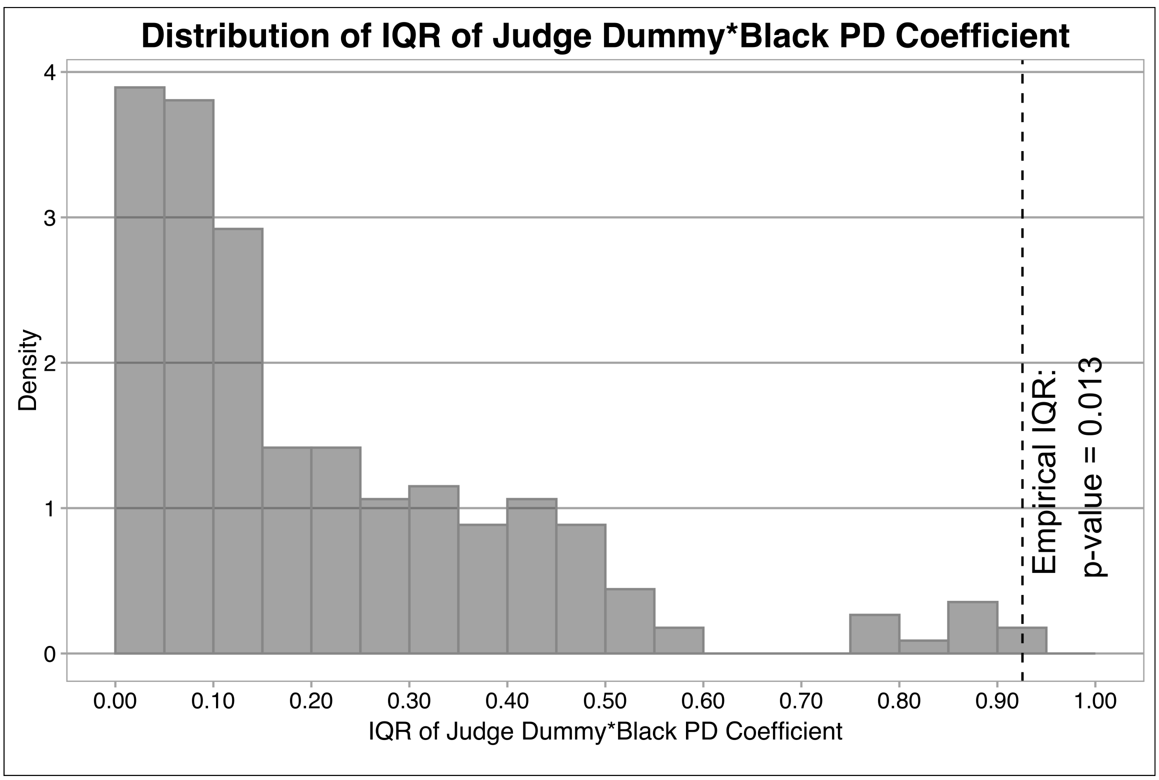

Figure 5: Black vs White Public Defenders - White, male judges

Figure 6: Density histograms of the distribution of the IQR of the estimates of the coefficient on the judge dummy $\times$ Black attorney dummy from Equation 7. IQRs of the coefficients of White male judges are calculated on simulated data constructed under the null hypothesis that there is no bias in the treatment of attorneys of the relevant demographic group. The dashed line marks the IQR of the estimated coefficients for White male judges from the true dataset.

6 Conclusion

This paper provides the first analysis of 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 am able to 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 across judges. In the setting of Miami-Dade County's first appearance hearings, I find statistically significant variation in the treatment of Black attorneys relative to White attorneys. The magnitude of this bias is such that a defendant with a Black attorney assigned to a 25^{th} percentile judge would be 2.8 percentage points less likely to be released pretrial than one assigned to a 75^{th} percentile judge. I do not find evidence of significant variation in the treatment of either Hispanic or female attorneys. To the extent that this does reflect bias against Black attorneys, this may explain the continued under-representation of Black attorneys in the legal profession. Further, given the preference of minority defendants to be represented by attorneys of the same racial background, it may contribute to the racial disparities seen in defendant outcomes. I find that increasing the favorability of all judges toward Black attorneys by one standard deviation of the estimated relative bias coefficients would decrease the racial gap in release rates amongst defendants who hire their attorneys by 33%. Regardless of its overall direction, the variation in the treatment of attorneys across judges would exacerbate the arbitrariness of a system that has substantial and long-lasting impacts on the defendants who go through it.

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., Fang, H. & Goonetilleke, P. (2023), Do cops know who to stop? assessing optimizing models of police behavior with a natural experiment, Technical report, National Bureau of Economic Research.

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.

Abrams, D. S. & Yoon, A. H. (2007), 'The luck of the draw: Using random case assignment to investigate attorney ability', *U. Chi. L. Rev.* **74**, 1145.

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.

American Bar Association (2022), 'ABA profile of the legal profession 2022'.
URL: *https://www.americanbar.org/content/dam/aba/administrative/news/2022/07/profile-report-2022.pdf*

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.

Ash, E., Chen, D. L. & Ornaghi, A. (2021), 'Gender attitudes in the judiciary: Evidence from US circuit courts', *Center for Law & Economics Working Paper Series* **2019**(02).

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

Bertrand, M. & Duflo, E. (2017), 'Field experiments on discrimination', *Handbook of Economic Field Experiments* **1**, 309–393.

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

Canay, I. A., Mogstad, M. & Mountjoy, J. (2020), On the use of outcome tests for detecting bias in decision making, Technical report, National Bureau of Economic Research.

Chaiken, S. (1980), 'Heuristic versus systematic information processing and the use of source versus message cues in persuasion.', *Journal of Personality and Social Psychology* **39**(5), 752.

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

Coate, S. & Loury, G. C. (1993), 'Will affirmative-action policies eliminate negative stereotypes?', *American Economic Review* pp. 1220–1240.

Cohen, D. L. & Peterson, J. L. (1981), 'Bias in the courtroom: Race and sex effects of attorneys on juror verdicts', *Social Behavior and Personality: an international journal* **9**(1), 81–87.

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.

Didwania, S. H. (2022), 'Gender favoritism among criminal prosecutors', *The Journal of Law and Economics* **65**(1), 77–104.

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.

Gazal-Ayal, O. & Sulitzeanu-Kenan, R. (2010), 'Let my people go: Ethnic in-group bias in judicial decisions—evidence from a randomized natural experiment', *Journal of Empirical Legal Studies* **7**(3), 403–428.

Glover, D., Pallais, A. & Pariente, W. (2017), 'Discrimination as a self-fulfilling prophecy: Evidence from French grocery stores', *The Quarterly Journal of Economics* **132**(3), 1219–1260.

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.

Hahn, P. W. & Clayton, S. D. (1996), 'The effects of attorney presentation style, attorney gender, and juror gender on juror decisions', *Law and Human Behavior* **20**(5), 533–554.

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.

Hull, P. (2021), What marginal outcome tests can tell us about racially biased decision-making, Technical report, National Bureau of Economic Research.

Jacobi, T. & Schweers, D. (2017), 'Justice, interrupted: The effect of gender, ideology, and seniority at supreme court oral arguments', *Va. L. Rev.* **103**, 1379.

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.

Katz, E. D., Rozema, K. & Sanga, S. (2023), 'Women in us law schools, 1948–2021', *Journal of Legal Analysis* **15**(1), 48–78.

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.

National Association for Law Placement (2022), '2021 report on diversity in U.S. law firms'.
URL: *https://www.nalp.org/uploads/2021NALPReportonDiversity.pdf*

New York State Judicial Committee on Women in the Courts (2020), 'Gender survey'.
URL: *https://www.nycourts.gov/LegacyPDFS/ip/womeninthecourts/Gender-Survey-2020.pdf*

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

Price, J. & Wolfers, J. (2010), 'Racial discrimination among NBA referees', *The Quarterly Journal of Economics* **125**(4), 1859–1887.

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.

Rehavi, M. M. & Starr, S. B. (2014), 'Racial disparity in federal criminal sentences', *Journal of Political Economy* **122**(6), 1320–1354.

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.

Sigal, J. et al. (1985), 'The effect of presentation style and sex of lawyer on jury decision-making behavior.', *Psychology: A Quarterly Journal of Human Behavior* **22**, 13–19.

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

Spamann, H. & Klöhn, L. (2023), 'Can law students replace judges in experiments of judicial decision-making?', *Available at SSRN 4362199* .

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.

Troccoli, K. P. (2002), 'I want a Black lawyer to represent me: Addressing a Black defendant's concerns with being assigned a white court-appointed lawyer', *Law & Ineq.* **20**, 1.

Tuttle, C. (2019), 'Racial disparities in federal sentencing: Evidence from drug mandatory minimums', *Available at SSRN* **3080463**.

Wood, S. M., DeVault, A., Miller, M. K., Kemmelmeier, M. & Summers, A. D. (2019), 'Decision-making in civil litigation: Effects of attorney credibility, evidence strength, and juror cognitive processing', *Journal of Applied Social Psychology* **49**(8), 498–518.

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 7 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 $H_0 : \beta_j^B = \beta_k^B \; \forall \; j, k \in J$.[62] 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 3. 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 $1,000$ 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 7. 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 7, over 30% of simulations have a p-value below 5%, which would lead to substantial overrejection of the null hypothesis.

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 5.1. For the purpose of this test, I use the interquartile range of the set of coefficients as my test statistic to measure variation. Figure 7 shows the distribution of p-values when applying this Monte Carlo method to data simulated from the theoretical model under the null of no bias. As can be seen in this figure, the resulting p-values are approximately uniformly distributed. Moreover, using a Kolmogorov-Smirnov. I fail to reject the null hypothesis that the p-values from my finite sample test are uniformly distributed.

[62]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 7: Distribution of Simulated p-values

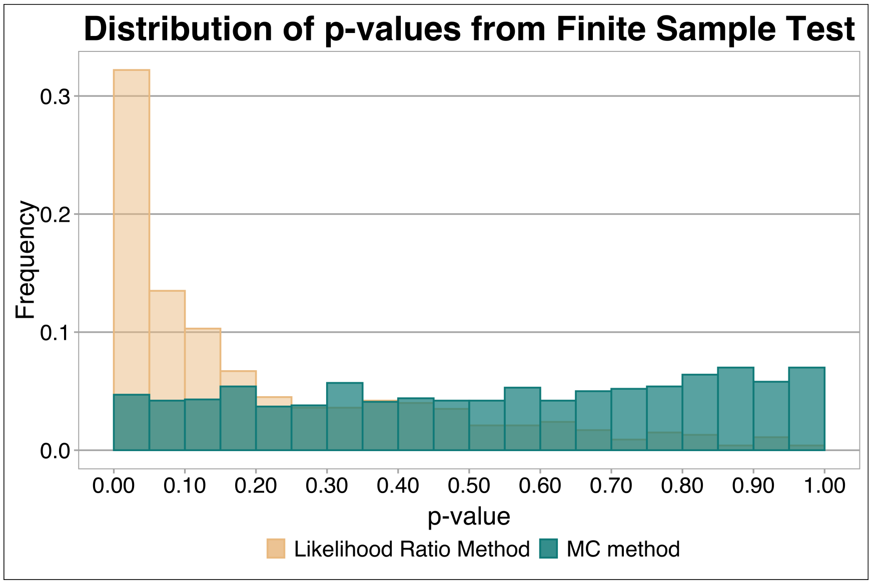

Note: Distribution of p-values from the likelihood ratio test using the asymptotic χ^2 and the proposed Monte Carlo Method. Data simulated under the null hypothesis of no bias in the treatment of attorneys using the model outlined in Section 3, 1000 simulations were run.

B Appendix Tables

Table A1

	Dependent variable:
	Pretrial Release
Black Judge * Black PD	0.007
	(0.165)
Hispanic Judge * Black PD	−0.125
	(0.092)
Black Judge * Hispanic PD	0.122
	(0.120)
Hispanic Judge * Hispanic PD	0.131
	(0.052)
Female Judge * Female PD	−0.027
	(0.050)
Observations	37,152

This table reports the results of the probit estimation of Equation 3. Robust standard errors clustered at the Judge ID and Public Defender ID level. Significance thresholds are adjusted for multiple hypothesis testing using the Benjamini-Hochberg procedure.

Table A2

	Dependent variable:
	Pretrial Release
Black Judge * Black Female PD	0.513
	(0.265)
Female Judge * Black Female PD	−0.063
	(0.156)
Black Judge * Black PD	−0.112
	(0.123)
Female Judge * Black PD	−0.218
	(0.127)
Black Judge * Female PD	0.179*
	(0.088)
Female Judge * Female PD	0.035
	(0.051)
Observations	36,732
Note:	*p<0.1; **p<0.05; ***p<0.01

This table reports the results of the probit estimation of equation 8. Robust standard errors clustered at the Judge ID and Public Defender ID level.

Table A3: Case Characteristics By Judge Gender

	Male	Female	p-value
N	82	82	
Cases	18043	19149	
Male Defendant	0.85	0.85	0.030
Black Defendant	0.54	0.54	0.825
First Degree Felony	0.1	0.11	0.081
Prior Misconduct	0.16	0.17	0.063
Male Attorney	0.59	0.55	0.216
Black Attorney	0.09	0.09	1.000
Hispanic Attorney	0.2	0.2	0.843
White Attorney	0.71	0.71	0.857
Pretrial Release	0.58	0.58	0.860
Misconduct	0.09	0.09	0.725

Note: This table provides summary statistics for the judges who conduct first appearances in Miami-Dade between January 2009 and December 2016, broken out by the gender of the judge. p-value denotes the p-value from an F-test of whether the row variable has explanatory power in predicting the race of the Judge.

Table A4: Case Characteristics By Public Defender Gender

	Male	Female	p-value
N	67	86	
Cases	19097	18074	
Male Defendant	0.85	0.85	0.073
Black Defendant	0.54	0.54	0.599
First Degree Felony	0.11	0.11	0.364
Prior Misconduct	0.22	0.22	0.862
Male Judge	0.51	0.47	0.216
Black Judge	0.08	0.08	0.878
Hispanic Judge	0.29	0.27	0.495
White Judge	0.62	0.63	0.376
Pretrial Release	0.58	0.58	0.527
Misconduct	0.09	0.1	0.021

Note: This table provides summary statistics for the Public Defenders who conduct first appearances in Miami-Dade between January 2009 and December 2016, broken out by the gender of the Public Defender. p-value denotes the p-value from an F-test of whether the row variable has explanatory power in predicting the race of the Public Defender.

C Appendix Figures

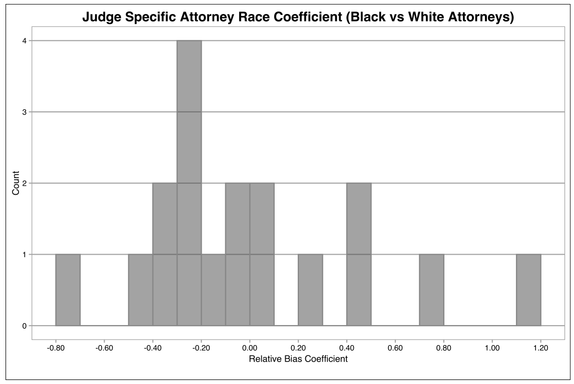

(a) Black vs. White Public Defenders

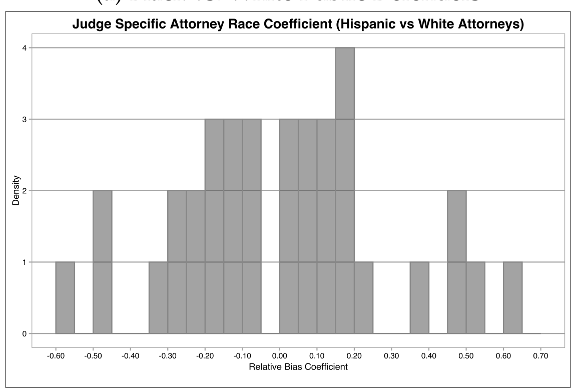

(b) Hispanic vs. White Public Defenders

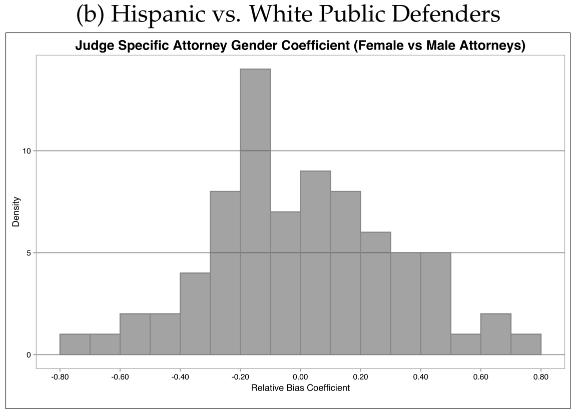

(c) Female vs Male Public Defenders

Figure A1: Histograms of the distribution across judges of the estimates of the judge dummy $\times$ attorney demographic group dummy from Equation 7. (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 20 involving White male public defenders are included. Coefficients have been rescaled to be mean zero

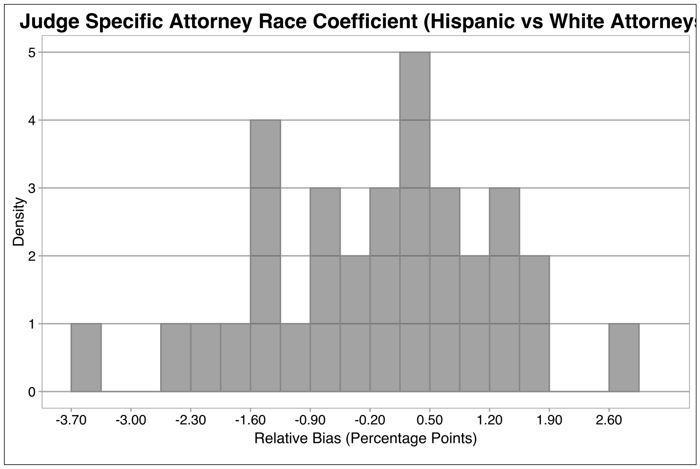

(a) Hispanic vs. White Public Defenders

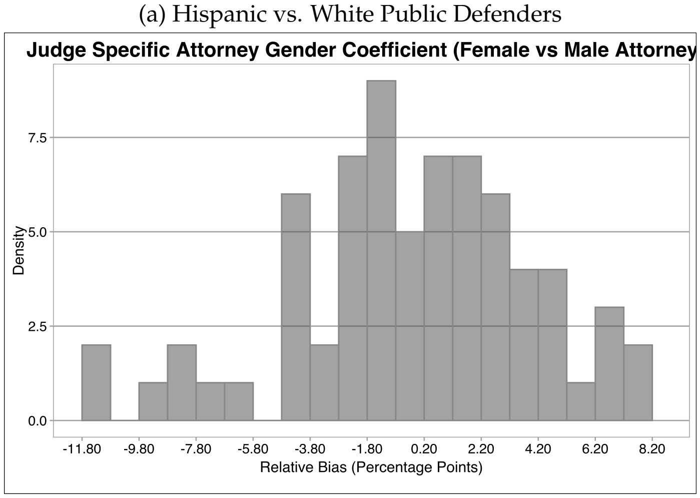

(b) Female vs Male Public Defenders

Figure A2: Histograms of the distribution across judges of the average marginal effects of the judge dummy $\times$ attorney demographic group dummy from Equation 7. 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^H and (b) 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 20 pretrial hearings involving public defenders of the relevant marginalized demographic group and 20 involving White male public defenders are included. Coefficients have been rescaled to be mean zero