Racial Bias in Bail Decisions*

David Arnold† Will Dobbie‡ Crystal S. Yang§

February 2017

PRELIMINARY AND INCOMPLETE
PLEASE DO NOT CIRCULATE OR CITE WITHOUT PERMISSION

Abstract

This paper develops a new test for identifying racial bias in the context of bail decisions – a high-stakes setting with large disparities between white and black defendants. We motivate our analysis using a simple model that predicts that, unless bail judges are racially biased, rates of pre-trial misconduct will be identical for marginal white and marginal black defendants. In contrast, marginal white defendants will have a higher probability of misconduct than marginal black defendants if bail judges are racially biased. To test the model, we develop a new estimator that uses the release tendencies of quasi-randomly assigned bail judges to identify the relevant race-specific misconduct rates. Estimates from Miami and Philadelphia show that bail judges are racially biased against black defendants, with substantially more racial bias among both inexperienced and part-time judges. We also find descriptive evidence that experienced and inexperienced judges place different weight on other, non-race characteristics that predict misconduct. We argue that these results are consistent with a model of on-the-job learning where racial bias decreases as judges adjust the predictive weight they place on the most salient observable defendant characteristics such as race.

*We thank Hank Farber, Roland Fryer, Jonah Gelbach, Paul Goldsmith-Pinkham, Lawrence Katz, Michal Kolesár, Alex Mas, Steven Shavell, and numerous seminar participants for helpful comments and suggestions. Kevin DeLuca and Amy Wickett provided excellent research assistance.

†Princeton University. Email: dharnold@princeton.edu

‡Princeton University and NBER. Email: wdobbie@princeton.edu

§Harvard Law School. Email: cyang@law.harvard.edu

Racial disparities exist at every stage of the criminal justice process. Compared to whites, blacks are more likely to be searched for contraband (Antonovics and Knight 2009), more likely to experience police use of force (Fryer 2016), more likely to be charged with a severe offense (Rehavi and Starr 2014), more likely to be convicted of an offense (Anwar, Bayer, and Hjalmarrson 2012), and more likely to be incarcerated (Abrams, Bertrand, and Mullainathan 2012). Racial disparities are particularly prominent in the setting of bail: in our data, black defendants are 11.2 percentage points more likely to be assigned monetary bail than white defendants and, conditional on receiving monetary bail, are assigned bail amounts that are $14,376 greater than whites.[1] However, determining whether these racial disparities are due to racial bias or statistical discrimination remains an empirical challenge.

To distinguish between racial bias and statistical discrimination, Becker (1957) proposed a method for measuring racial bias broadly known as the "outcome" test. In our setting, Becker's outcome test is based on the idea that racial bias against blacks will result in lower rates of pre-trial misconduct for marginally released black defendants than for marginally released white defendants. If, on the other hand, racial gaps in bail setting are due to statistical discrimination, there will be no difference in the misconduct rates of marginal black and white defendants. Importantly, Becker's test relies on comparing outcomes for marginal black and white defendants, not average black and white defendants. In practice, however, researchers usually cannot observe which defendants are and are not on the margin of release (e.g., Ayres 2002).

In recent years, two seminal papers have developed tests of racial bias that partially circumvent this infra-marginality problem. In the first paper, Knowles, Persico, and Todd (2001) show that all motorists of a given race will carry contraband with equal probability if motorists respond to the race-specific probability of being searched. In this scenario, the marginal and average hit rates of police searches will be identical. Knowles et al. (2001) find no difference in the average success rate of police searches for white and black drivers in their data, leading them to conclude that there is no racial bias in their setting. In a second important paper, Anwar and Fang (2006) develop a test of *relative* racial bias that is based on the idea that the ranking of search rates and search success rates by white and black police officers should be unaffected by the race of the motorist when there is no racial bias. This empirical test has the advantage of allowing for infra-marginality problems, but at the cost of only being able to detect relative, and not absolute, racial bias. Anwar and Fang (2006) also find no evidence of racial bias in their data.[2]

In this paper, we propose a new test for identifying absolute racial bias in the context of bail setting. Bail is an ideal setting to test for racial bias for a number of reasons. First, the legal

[1] Authors' calculation using the data described in Section II. Racial disparities in bail setting are also observed in other bail data. For example, black felony defendants in state courts are nine percentage points more likely to be detained pre-trial compared to otherwise similar white defendants (McIntyre and Baradaran 2013).

[2] We replicate the Knowles et al. (2001) and Anwar and Fang (2006) tests in our data, finding no evidence of racial bias. The difference between our preferred approach and the Knowles et al. (2001) and Anwar and Fang (2006) tests is that (1) we identify treatment effects for marginal defendants rather than the average defendant and (2) we identify absolute rather than relative bias. See Section III.D for additional details on why the Knowles et al. (2001) and Anwar and Fang (2006) tests yield different results.

objective of bail judges is narrow, straightforward, and measurable – to set bail conditions to minimize the risk of pre-trial misconduct. In contrast, the objectives of judges at other stages of the criminal justice process, such as sentencing, are complicated by multiple hard-to-measure objectives, such as retribution and mercy. Second, bail judges must make on-the-spot judgments with limited information and little to no interaction with defendants. In addition, many bail judges are relatively inexperienced and untrained. Both of these institutional features may make bail decisions particularly prone to stereotypes or categorical heuristics (e.g., Fryer and Jackson 2008, Bordalo et al. 2016). Finally, bail decisions are extremely consequential for both black and white defendants, with prior work suggesting that detained defendants suffer about $40,000 in lost earnings and government benefits alone (Dobbie, Goldin, and Yang 2016).

To motivate our empirical analysis, we begin with a simple model of bail setting that builds on Becker (1957). In the model, bail judges form an expectation of the probability of pre-trial misconduct and release all defendants with a probability of misconduct below a judge-specific threshold. We allow each judge-specific threshold to vary by defendant race, meaning that a judge is racially biased if his threshold is different for black and white defendants. The key empirical prediction from the model is that we can infer the average judge-by-race threshold from the pre-trial misconduct rates of marginal white and marginal black defendants. It therefore follows that racial bias against blacks should result in lower rates of pre-trial misconduct for marginally released black defendants than for marginally released white defendants.

To test the model, we develop an estimator that identifies the probability of pre-trial misconduct for marginal white and black defendants using the release tendencies of quasi-randomly assigned judges. Intuitively, our empirical strategy uses the local nature of instrumental variable (IV) estimators to identify treatment effects for individuals on the margin of release – the parameter required by the Becker outcome test. We show that the local average treatment effect (LATE) identified by our judge IV strategy is exactly equal to the average treatment effect of marginal defendants if the treatment effect is constant among the complier population. We then loosen this assumption and provide a formula for the upper bound of the estimation bias between our IV-based estimate for racial bias and the true level of racial bias. We show that this estimation bias shrinks as the number of judges increases and as the heterogeneity in treatment effects among compliers shrinks. In practice, we find that any potential estimation bias in our setting is less than 0.5 percentage points.[3,4]

We divide the empirical analysis into two sections. First, we test for racial bias in bail set-

[3]There is a large literature on the use of instrumental variables to estimate local average treatment effects (e.g., Imbens and Angrist 1994, Angrist, Imbens, and Rubin 1996) and marginal treatment effects (e.g., Heckman and Vyltacil 2005, and Heckman, Urzua, and Vyltacil 2006) in the presence of heterogenous treatment effects. Though IV estimates are often criticized for the local nature of the estimate, our empirical test relies on treatment effects for exactly this kind of local population at the margin of release.

[4]Our empirical test does not allow us to examine whether bail decisions for any defendant race are optimal. For instance, we cannot conclusively determine whether differences in the treatment of black and white defendants are due to judges being too lenient on white defendants (relative to the optimum) or judges being too harsh on black defendants. Rather, our empirical test simply identifies whether there is differential treatment of black and white defendants due to racial bias.

ting by comparing the pre-trial misconduct rates of marginal black defendants and marginal white defendants. Using administrative court data from Miami and Philadelphia, we find evidence of significant racial bias in bail setting. Marginally released white defendants are approximately 18 percentage points more likely to be rearrested prior to disposition than marginally released black defendants. Among high-risk defendants, marginally released white defendants are 38.5 percentage points more likely to be rearrested compared to marginally released black defendants. Racial bias against black defendants is also consistently larger among drug offenders, prior offenders, and defendants charged with felonies. Our estimates are nearly identical if we account for other observable crime and defendant differences by race, implying that our results cannot be explained by black-white differences in certain types of crimes (e.g., the proportion of felonies versus misdemeanors) or black-white differences in defendant characteristics (e.g., the proportion with a prior offense versus no prior offense).

In the second part of the analysis, we investigate whether the relative inexperience of bail judges can explain racial bias in our setting. As discussed above, bail judges must make quick judgments on the basis of limited information and virtually no training. Moreover, in many jurisdictions, bail hearings are heard by "generalist" judges that only assist with bail hearings a few days a year. It is possible that these relatively inexperienced bail judges may be more likely to rely on incorrect inferences of risk based on defendant race, leading to the relative over-detention of low-risk black defendants (e.g., Fryer and Jackson 2008, Bordalo et al. 2016). Inexperienced bail judges may also have limited experience interacting with black defendants, potentially making it more difficult for these judges to discern the true level of risk for black defendants compared to observably similar white defendants. Finally, it is possible that inexperience leads judges to value the freedom of black defendants less than the freedom of observably similar white defendants.

We document three sets of facts which suggest that the relative lack of on-the-job experience explains racial bias in bail setting. First, we find that racial bias is higher in Miami, where the bail judges are generalists that handle bail hearings only a few weekends a year, compared to Philadelphia, where judges are specialists that handle bail hearings throughout the year. Second, we find that racial bias is significantly higher among inexperienced judges in Miami compared to experienced judges in Miami. Finally, we find more descriptive evidence that the relationship between pre-trial release and the predicted risk of pre-trial misconduct is substantially different between experienced and inexperienced judges, with the differences largely driven by inexperienced judges being relatively less lenient for defendants charged with violent offenses. Taken together, these results are consistent with inexperienced judges placing relatively more weight on particularly salient case and defendant characteristics such as race and the nature of the offense.[5]

We conclude that racial bias on the part of bail judges is a substantial driver of black-white disparities in pre-trial detention and that increasing on-the-job experience provides a potential

[5]Note that differences in the pre-trial release tendencies of inexperienced versus experienced judges do not necessarily lead to a violation of the monotonicity assumption required for our IV estimator. The monotonicity assumption is valid so long as individuals released by a strict judge would also be released by a more lenient judge and individuals detained by a lenient judge would also be detained by a stricter judge.

way to alleviate any unwarranted disparities. If our interpretation of the results is correct, this suggests that providing judges with increased opportunities for training or on-the-job feedback could play an important role in decreasing racial disparities in the criminal justice system. Our results also suggest that providing judges with data-based risk assessments, including those generated by machine learning algorithms (e.g., Kleinberg et al. 2015), could help decrease racial disparities, although our estimates do not directly speak to this type of policy.

Our results contribute to an important literature testing for racial bias in the criminal justice system, in particular analyzing the stop and search decisions of police. As discussed above, Knowles et al. (2001) and Anwar and Fang (2006) are seminal works in this area. Subsequent work by Antonovics and Knight (2009) finds that police officers in Boston are more likely to conduct a search if the race of the officer differs from the race of the driver, which they interpret as evidence of racial bias. Also related is work by Ayres and Waldfogel (1994) showing that bail bond dealers in New Haven charge lower rates to minorities, leading them to conclude that bail judges are biased against minorities. Most recently, Bushway and Gelbach (2011) use a parametric framework that accounts for unobserved heterogeneity across defendants to show that bail judges are biased against black defendants. To date, however, the literature has been unable to develop a test for absolute racial bias that does not depend on strong modeling assumptions.[6]

Our paper is also related to an emerging literature documenting the costs of pre-trial detention using the quasi-random assignment of cases to bail judges (e.g., Dobbie, Goldin, and Yang 2016, Gupta, Hansman, and Frenchman 2016, Leslie and Pope 2016, and Stevenson 2016). Pre-trial detention increases the probability of pleading guilty before trial and decreases formal sector employment and the receipt of employment- and tax-related government benefits (Dobbie et al. 2016). A partial cost-benefit analysis that accounts for administrative jail expenses, the costs of apprehending defendants, the costs of future crime, and the economic impacts on defendants suggests that the net cost of pre-detention is between $37,031 and $40,048 per defendant, largely due to the significant collateral consequences of having a criminal conviction on labor market outcomes and the relatively low costs of apprehending defendants who fail to appear in court (Dobbie et al. 2016).

The remainder of the paper is structured as follows. Section I develops the theoretical model underlying our analysis and proposes an empirical test based on the theoretical predictions of the model. Section II provides background on bail decisions in our setting and describes our data. Section III presents the results, Section IV explores the role of experience and learning, and Section V concludes. An online appendix provides additional results and detailed information on the outcomes and setting of our analysis.

I. An Empirical Test of Racial Bias

In this section, we develop an empirical test for racial bias in bail setting. Our theoretical framework closely follows the previous literature on the outcome test in the criminal justice system (e.g., Becker

[6]There is a large literature examining racial bias in other settings. See, for example, Fryer (2011) and Bertrand and Duflo (2016) for reviews of the literature.

1957, Knowles et al. 2001, Anwar and Fang 2006, Antonovics and Knight 2009). Consistent with the prior literature, we show that we can test for racial bias by comparing treatment effects for the marginal black and marginal white defendant. We then develop an estimator that identifies these race-specific treatment effects using an instrumental variables approach that exploits the quasi-random assignment of cases to judges.

A. Overview of the Bail System

In the United States, bail judges are granted considerable discretion to determine which defendants should be released before trial. Bail judges are meant to balance two competing objectives when deciding whether to detain or release a defendant before trial. First, bail judges are directed to release all but the most dangerous defendants before trial to reduce jail expenses and increase defendant well-being. Second, bail judges are instructed to minimize the risk of pre-trial misconduct by setting the appropriate conditions for release.

These bail conditions are set at a bail hearing held within 24 to 48 hours of a defendant's arrest. In most jurisdictions, bail hearings last only a few minutes and are held through a video-conference to the detention center. During the bail hearing, the assigned bail judge considers factors such as the nature of the alleged offense, the weight of the evidence against the defendant, the nature and probability of danger that the defendant's release poses to the community, and any record of prior flight or bail violations, among other factors (Foote 1954). Because bail judges are granted considerable discretion in setting the appropriate bail conditions, there are substantial differences across judges in the same jurisdiction (e.g., Dobbie et al. 2016, Gupta et al. 2016, Leslie and Pope 2016, Stevenson 2016).

The assigned bail judge has a number of potential options when setting a defendant's bail conditions. For example, the bail judge can release low-risk defendants on a promise to return for all court appearances, known as release on recognizance (ROR). For defendants who pose a higher risk of flight or new crime, the bail judge can allow release but impose non-monetary conditions such as electronic monitoring or periodic reporting to pre-trial services. The judge can also require defendants to post a monetary amount to secure release, typically 10 percent of the total bail amount. If the defendant fails to appear at the required court appearances or commits a new crime while out on bail, either he or the bail surety forfeits the 10 percent payment and is liable for the remaining 90 percent of the total bail amount. In practice, the median bail amount is $5,000 in our sample, and only 31 percent of defendants are able to meet the required monetary conditions to secure release. Bail may also be denied altogether for defendants who commit the most serious crimes such as first- or second-degree murder.

One important difference between jurisdictions is the degree to which bail judges specialize in conducting bail hearings. For example, in our setting, Philadelphia bail judges are full-time specialists who are tasked with setting bail 24 hours a day, seven days a week. In contrast, the bail judges we study in Miami are part-time nonspecialists who assist the bail court by serving weekend shifts. These weekend bail judges spend their weekdays as trial court judges. We will return to this

distinction as a possible source of racial bias in Section IV.

B. Model of Judge Behavior

This section develops a theoretical framework that allows us to define an outcome-based test of racial bias in bail setting. Let i denote defendants and $\mathbf{V}_i$ denote all case and defendant characteristics considered by the bail judge, excluding defendant race r_i. The expected cost of release for defendant i conditional on observable characteristics $\mathbf{V}_i$ and race r_i is equal to the expected probability of pre-trial misconduct $\mathbb{E}[\alpha_i|\mathbf{V}_i, r_i]$ times the cost of misconduct C.[7] For simplicity, we normalize $C = 1$, so that the expected cost of release conditional on observable characteristics is equal to $\mathbb{E}[\alpha_i|\mathbf{V}_i, r_i]$. We also simplify our notation by letting the expected cost of release conditional on observables be denoted by $\mathbb{E}[\alpha_i|r_i]$.

The benefit of releasing defendant i assigned to judge j is denoted by $t_r^j(\mathbf{V}_i)$, where we explicitly allow for the benefits to be a function of the observable case and defendant characteristics $\mathbf{V}_i$. The benefit of release $t_r^j(\mathbf{V}_i)$ includes cost savings from reduced jail time and private gains to defendants such as an improved bargaining position with the prosecutor and increased labor force participation. Importantly, we allow the benefit of release $t_r^j(\mathbf{V}_i)$ to vary by race $r \in W, B$ to allow for judge preferences to vary for white and black defendants.

Definition 1. Following Becker (1957), we define judge j as racially biased against black defendants if $t_W^j(\mathbf{V}_i) > t_B^j(\mathbf{V}_i)$. Thus, for racially biased judges, there is a higher benefit of releasing white defendants than releasing observably identical black defendants.

Finally, we assume that bail judges are risk neutral and maximize the net benefit of pre-trial release. Thus, bail judge j will release defendant i if and only if the benefit of pre-trial release is greater than the expected cost of release:

$$\mathbb{E}[\alpha_i|r_i = r] \leq t_r^j(\mathbf{V}_i) \tag{1}$$

Given this decision rule, the marginal defendant for judge j and race r is the defendant i for whom the expected cost of release is exactly equal to the benefit of release, i.e., $\mathbb{E}[\alpha_i^j|r_i = r] = t_r^j(\mathbf{V}_i)$. We simplify our notation moving forward by letting this expected cost of release for the marginal defendant for judge j and race r be denoted by α_r^j.

Based on the above framework and Definition 1, the model yields the familiar outcome-based test for racial bias from Becker (1957):

Proposition 1. If judge j is racially biased against black defendants, then $\alpha_W^j > \alpha_B^j$. Thus, for racially biased judges, the expected cost of release for the marginal white defendant is higher than the expected cost of release for the marginal black defendant.

Proposition 1 predicts that the marginal white and black defendant should have the same probability of pre-trial misconduct if judge j is racially unbiased, but that the marginal white defendant should

[7] In the model, we abstract away from the underlying determinants of the expected probability of pre-trial misconduct, $\mathbb{E}[\alpha_i|\mathbf{V}_i, r_i]$. Thus, we assume that bail judges are only concerned with the probability of pre-trial misconduct per se and not with the reasons a defendant commits pre-trial misconduct.

have a higher probability of misconduct than the marginal black defendant if judge j is racially biased against black defendants.

C. Empirical Test of Racial Bias in Bail Setting

Overview: The goal of our analysis to formally establish the conditions under which we can identify the rate of pre-trial misconduct for the marginal white and marginal black defendant across all judges. Defendant i's probability of pre-trial misconduct, Y_i, is given by the following relationship:

$$Y_i = \alpha_W Released_i \cdot White_i + \alpha_B Released_i \cdot Black_i + \beta \mathbf{X}_i + \mathbf{U}_i + \varepsilon_i \tag{2}$$

where $Released_i$ is an indicator for being released before trial, $White_i$ and $Black_i$ are race indicators, $\mathbf{X}_i$ denotes characteristics of the defendant observed by both the econometrician and bail judge, and $\mathbf{U}_i$ denotes characteristics observed by the bail judge but not the econometrician. In practice, $\mathbf{X}_i$ includes variables such age, gender, type of crime, and prior offenses, while $\mathbf{U}_i$ include characteristics such as the defendant's physical appearance and any information conveyed during the bail hearing. ε_i is the idiosyncratic defendant-level variation that is unobserved by both the econometrician and the judge.

OLS estimates of α_W and α_B from Equation (2) will typically not recover unbiased estimates of the causal rate of pre-trial misconduct for marginal black and marginal white defendants across all judges, $\mathbb{E}[\alpha_B^j]$ and $\mathbb{E}[\alpha_W^j]$, for two reasons. First, characteristics observable to the judge but not the econometrician, $\mathbf{U}_i$, may be correlated with $Released_i$, resulting in omitted variable bias. For example, bail judges may be more likely to release defendants who both appear to be less dangerous during the bail hearing and who are, in fact, less likely to have an incident of pre-trial misconduct. In this scenario, OLS estimates of Equation (2) will be downwards biased from the true average treatment effect.

The second, and more important, reason OLS estimates will not recover unbiased estimates of $\mathbb{E}[\alpha_B^j]$ and $\mathbb{E}[\alpha_W^j]$ is that the treatment effect of pre-trial release may be correlated with judges' decision rules, meaning that the average treatment effect identified by OLS will not be equal to the marginal treatment effect required by our test (e.g., Ayres 2002). Thus, even if the econometrician observes the full set of observables known to the bail judge, $\mathbf{X}_i$ and $\mathbf{U}_i$, OLS estimates are still not sufficient to test for racial bias unless one is willing to assume constant treatment effects across the entire distribution of defendants (i.e., assuming that the average treatment effect is equal to the marginal treatment effect). In our model, we explicitly rule out constant treatment effects by allowing judges' race-specific decision rules to be correlated with the expected treatment effect, $\mathbb{E}[\alpha_i | r_i = r]$ (see Equation 1). In this scenario, the average treatment effect will be an underestimate of the marginal treatment effect required by our outcome test.

In this paper, we identify racial bias in the presence of both omitted variables and infra-marginality issues using the local nature of instrumental variables estimators to estimate causal treatment effects for individuals exactly on the margin of release, the parameter required by the

outcome-based test described above. As we will show below, the local average treatment effect identified by our empirical strategy is exactly equal to the treatment effect of the marginal defendant if the treatment effect is constant among the complier population. We then loosen this assumption and provide a formula for the upper bound of the bias between our IV estimate for racial bias and the true level of racial bias. We show that this bias shrinks as the number of judges increases and as the heterogeneity in treatment effects among compliers shrinks.

Simple Case with Two Judges: To build intuition for why an IV estimator can be used to test for racial bias under reasonable assumptions, we first consider a simple setting where there are only two bail judges ($j = 2$). We begin by noting that we can express the expected cost of release for a defendant of race r, $\mathbb{E}[\alpha_i | r_i = r]$, in the potential outcomes framework as:

$$\mathbb{E}[\alpha_i | r_i = r] = \mathbb{E}[Y_i(1) - Y_i(0) | r_i = r]$$

where $Y_i(1)$ is an indicator for pre-trial misconduct for defendant i if released before trial, and $Y_i(0)$ is the probability of pre-trial misconduct for defendant i if detained before trial.

Let bail judge 1's and bail judge 2's preferences for release be given by $t_r^1(\mathbf{V}_i)$ and $t_r^2(\mathbf{V}_i)$ with $t_r^1(\mathbf{V}_i) < t_r^2(\mathbf{V}_i)$, such that judge 2 is more lenient to defendants of the same race compared to judge 1. Using judge-specific leniency measures that proxy for judge preferences, z_1 and z_2, as instrumental variables for pre-trial release in Equation (2), we can express the resulting estimates of α_W and α_B as follows:

$$\begin{aligned} \alpha_W &= \mathbb{E}\big[\alpha_i | R_i(z_2) - R_i(z_1) = 1, r_i = W] \\ \alpha_B &= \mathbb{E}\big[\alpha_i | R_i(z_2) - R_i(z_1) = 1, r_i = B] \end{aligned} \tag{3}$$

where $R_i(z_j)$ equals 1 if defendant i is released if assigned to judge j. Specifically, α_r is the expected treatment effect for compliers of race r, that is, defendants who would have been released if assigned to judge 2 but detained if assigned to judge 1. Turning to the potential outcomes framework, α_r can be expressed as $\mathbb{E}\big[Y_i(1) - Y_i(0) | R_i(z_2) - R_i(z_1) = 1, r_i = r]$.

Given that judge j releases defendant i of race r if $\mathbb{E}[\alpha_i | r_i = r] \leq t_r^j(\mathbf{V}_i)$, compliers in this scenario are defendants of race r for whom $t_r^1(\mathbf{V}_i) < \mathbb{E}[\alpha_i | r_i = r] \leq t_r^2(\mathbf{V}_i)$. In other words, the complier population is comprised of defendants whose expected cost of release falls between judge 1 and judge 2's threshold for release. In contrast, recall that the marginal defendant of race r for judge j is α_r^j where $\alpha_r^j = t_r^j(\mathbf{V}_i)$ (Proposition 1).

If we assume the individual-specific treatment effect $\mathbb{E}[\alpha_i | r_i = r]$ is constant among the complier population of a given race, then the LATE identified by our IV estimator, α_r, is equivalent to the treatment effect for the marginal defendant, $\mathbb{E}[\alpha_r^j]$. Specifically, if $\mathbb{E}[\alpha_i | r_i = r] = \bar{\alpha_r}$ for all defendants i for whom $t_r^1(\mathbf{V}_i) < \mathbb{E}[\alpha_i | r_i = r] \leq t_r^2(\mathbf{V}_i)$, the local average treatment effect for these compliers is precisely equal to the average treatment effect of the marginal defendants ($\bar{\alpha_r} = \alpha_r = \mathbb{E}[\alpha_r^j]$). If this condition holds, we are able to precisely identify the treatment effects for marginal white and

marginal black defendants as required by our outcome test.

If, however, there is not a constant treatment effect among compliers, IV estimates of Equation (2) will not be equivalent to the treatment effect for the marginal defendant. It is easy to see that as the differences in the preferences of the two judges decreases, $t_r^1(\mathbf{V}_i) \rightarrow t_r^2(\mathbf{V}_i)$, the IV estimate converges to the race-specific benefit of release because the local average treatment effect for compliers converges to the treatment effect for the marginal defendant, $\alpha_r \rightarrow \mathbb{E}[\alpha_r^j]$. Therefore, any potential bias in our estimator shrinks as a result of either (1) decreased distance in leniency between any two judges and (2) reduced heterogeneity in treatment effects among compliers.

Generalized Case with Many Judges: The intuition behind our estimator can be easily expanded to the more general case with many judges. Suppose for instance that there are $J+1$ judges in the bail system and that we can rank judge preferences $\{t_r^0(\mathbf{V}_i), ..., t_r^J(\mathbf{V}_i)\}$ in order of leniency toward pre-trial release where $t_r^0(\mathbf{V}_i)$ is the strictest judge and $t_r^J(\mathbf{V}_i)$ is the most lenient judge. Again, we use judge-specific leniency measures, $\{z_0, ..., z_J\}$, as instrumental variables for pre-trial release. Under this more general case, our estimators α_W and α_B are comprised of a weighted sum of pairwise local average treatment effects where each pairwise treatment effect is the treatment effect of compliers whose pre-trial release status is altered by assignment to judge j versus $j-1$. Specifically, each pairwise estimate can be characterized as:

$$\begin{aligned} \alpha_W^{j,j-1} &= \mathbb{E}[\alpha_i | R_i(z_j) - R_i(z_{j-1}) = 1, r_i = W] \\ \alpha_B^{j,j-1} &= \mathbb{E}[\alpha_i | R_i(z_j) - R_i(z_{j-1}) = 1, r_i = B] \end{aligned} \tag{4}$$

where $R_i(z_j)$ equals 1 if defendant i is released if assigned to judge j. Returning again to the potential outcomes framework, $\alpha_r^{j,j-1}$ can be expressed as $\mathbb{E}\big[Y_i(1) - Y_i(0) | R_i(z_j) - R_i(z_{j-1}) = 1, r_i = r\big]$.

Under the same logic as the simple two judge example above, our LATE precisely identifies the treatment effects for the marginal white and marginal black defendants if we assume that the treatment effect is constant among the complier population. In the absence of this constant treatment effects assumption, any potential bias in our estimator decreases as the heterogeneity in treatment effects among compliers shrinks, and as $t_r^{j-1}(\mathbf{V}_i) \rightarrow t_r^j(\mathbf{V}_i)$ for all $j, j-1$.

Identifying Assumptions: Building on the above example, we can now formally define the assumptions underlying our IV estimator. Let Z_i be a measure of the assigned judge's propensity for pre-trial release that will serve as our instrumental variable for pre-trial release. For illustrative purposes, assume for the moment that Z_i is a scalar that takes on values ordered $\{z_0, ..., z_J\}$, where $J+1$ is the number of total judges in the bail system. For example, a value of $z_j = 0.5$ indicates that judge j releases 50 percent of all defendants.

As will be described in further detail in Section II.B, we construct our judge leniency measure using a leave-out procedure that captures the pre-trial release tendency of judges across both white and black defendants, implicitly assuming that the judge ordering produced by the scalar Z_i is the same for both white and black defendants. We relax this assumption in Section III.C by separately

calculating our leave-out judge leniency measure by defendant race.

We begin with the three standard assumptions for the identification of a well-defined LATE using an instrumental variables approach (Angrist and Imbens 1994):

Assumption 1. [Existence]. Pre-trial release is a nontrivial function of Z_i such that a first stage exists:

$$Cov(Released_i, Z_i) \neq 0$$

Assumption 1 ensures that there is a first stage relationship between our instrument Z_i and the probability of pre-trial release.

Assumption 2. [Exclusion Restriction]. Z_i is uncorrelated with unobserved determinants of Y_i:

$$Cov(Z_i, \mathbf{v}_i) = 0$$

where $\mathbf{v}_i = \mathbf{U}_i + \varepsilon_i$. Assumption 2 ensures that our instrument Z_i is orthogonal to characteristics unobserved by the econometrician, $\mathbf{v}_i$. In other words, Assumption 2 assumes that the assigned judge only affects pre-trial misconduct through the channel of pre-trial release.

Assumption 3. [Monotonicity]. The impact of judge assignment on the probability of pre-trial release is monotonic if for each z_{j-1}, z_j pair:

$$R_i(z_j) - R_i(z_{j-1}) \geq 0$$

where $R_i(z_j)$ equals 1 if defendant i is released if assigned to judge j. Assumption 3 implies that any defendant released by a strict judge would also be released by a more lenient judge, and any defendant detained by a lenient judge would also be detained by a strict judge.

Following Angrist and Imbens (1994), our IV estimator is valid and well-defined under Assumptions 1-3. Specifically, our race-specific estimators utilizing a multi-valued instrument are:

$$\alpha_r = \sum_{j=1}^{J} \lambda_r^j \cdot \alpha_r^{j,j-1} \tag{5}$$

where λ_r^j are non-negative weights which sum to one and $\alpha_r^{j,j-1} = \mathbb{E}[Y_i(1) - Y_i(0) | R_i(z_j) - R_i(z_{j-1}) = 1, r_i = r]$. See Imbens and Angrist (1994) for additional details.

The weights λ_r^j depend on the size of the subpopulation whose treatment status is altered by changing the value of the instrument from z_j to z_{j-1}, as well as the probability of being assigned a particular judge, denoted π^j. For example, in a case with three judges, if judge 1 is much stricter than judge 2, while judge 3 is only slightly stricter than judge 2, then the IV estimator will attribute more weight to the pairwise LATE $\alpha_r^{2,1}$ than the pairwise LATE $\alpha_r^{3,2}$.

Our estimator for racial bias D is given by:

$$D = \alpha_W - \alpha_B \tag{6}$$

where D captures the weighted average difference in treatment effects for marginal white defendants and marginal black defendants across all judges within a court. $D > 0$ indicates that the bail system is, on average, racially biased against black defendants and $D < 0$ indicates that the bail system is, on average, racially biased against white defendants.

Potential Bias: As discussed previously, if we assume that the treatment effect is constant among the complier population, our IV estimate D exactly identifies the treatment effects for marginal white and marginal black defendants, the parameter required by our outcome test. We now loosen this assumption and characterize the maximum bias of our estimator. For tractability, we make the following linearity assumption:

Assumption 4. [Linear First Stage]. The first stage relationship between pre-trial release and Z_i is given by a linear probability model of the form:

$$Released_i = \gamma_0 + \gamma_1 Z_i \cdot White_i + \gamma_2 Z_i \cdot Black_i + \pi \mathbf{X}_i + \mathbf{U}_i + \varepsilon_i \tag{7}$$

Assumption 4 allows us to calculate a simple closed-form formula for the maximum bias between our IV estimate of racial bias and the true amount of racial bias. We show below that a linear first stage is also consistent with our data (see Figure 1).

What does the assumption of a linear first stage mean for the interpretation of our estimates? Specifically, by assuming the first stage is linear for each defendant race r, but allowing possibly different coefficients by race ($\gamma_1 \neq \gamma_2$), we enforce that the weights λ_r^j are the same for every pairwise LATE by defendant race r. To understand, consider again a three judge example with $t_r^1(\mathbf{V}_i) < t_r^2(\mathbf{V}_i) < t_r^3(\mathbf{V}_i)$. Linearity of the first stage implies that moving from judge 1 to judge 2 shifts the same proportion of compliers for each race. For example, if 10 percent of all white compliers are compliers between judge 1 and 2 and 90 percent are compliers between judge 2 and 3, then linearity assumes that 10 percent of all black compliers are compliers between judge 1 and 2 and 90 percent are compliers between judge 2 and 3. With $J + 1$ judges, this linearity assumption can be expressed as:

$$\frac{Pr(t^{j-1}(\mathbf{V}_i) < \mathbb{E}[\alpha_i] \leq t^j(\mathbf{V}_i) | r_i = W)}{Pr(t^{j-1}(\mathbf{V}_i) < \mathbb{E}[\alpha_i] \leq t^j(\mathbf{V}_i) | r_i = B)} = constant \tag{8}$$

for each $j, j-1$.[8]

Proposition 2. If Assumptions 1-4 are satisfied, the maximum bias of our IV estimator D is given by $\max_j(\lambda^j)(\alpha_{max} - \alpha_{min})$ under the assumption of no racial bias, where α_{max} is the largest treatment effect among compliers, α_{min} is the smallest treatment effect among compliers, and λ^j is

[8] An even stronger assumption of linearity is used in most prior studies utilizing quasi-random assignment of judges (Aizer and Doyle 2015, Dahl et al. 2014, Dobbie et al. 2016, Kling 2006), which implicitly assume that moving a defendant from a judge with leniency z_j to a judge with leniency z_{j+1} increases the probability of treatment uniformly across all defendants.

given by:

$$\lambda^j = \frac{(z_j - z_{j-1}) \cdot \sum_{l=j}^{J} \pi^l (z_l - \mathbb{E}[Z])}{\sum_{m=1}^{J} (z_j - z_{j-1}) \cdot \sum_{l=m}^{J} \pi^l (z_l - \mathbb{E}[Z])} \tag{9}$$

where π^j is the probability of being assigned to judge j.

Proof. See Appendix B.

Proposition 2 formalizes the intuition behind the potential bias of our estimator. The maximum bias of our estimator decreases with both (1) the distance in leniency between any two judges and (2) the heterogeneity in treatment effects among compliers decreases. Intuitively, the distance between any two leniency measures decreases as the number of judges increases, therefore decreasing the weight attributed to any particular pairwise LATE. Similarly, the bias in our estimator, holding fixed the distance between the judge leniency measures, decreases as the heterogeneity in treatment effects among compliers decreases.

In Appendix B, we calculate the bias in our preferred estimates by plugging in the empirical counterparts into the formula given by Equation (9). We estimate that under the null of no racial bias, the difference in treatment effects between marginal white and marginal black defendants could be no greater than 0.5 percentage points.[9]

D. Discussion and Extensions

In this section, we discuss some important assumptions underlying our model, possible extensions, and how they affect the interpretation of our results.

Racially Biased Prediction Errors: In our model, we have assumed that judges agree on the expected cost of release, $\mathbb{E}[\alpha_i | r_i]$, but not the benefit to release, $t_r^j(\mathbf{V}_i)$, which we allow to vary across judges. An alternative approach is to assume that judges instead vary in their predictions of the expected cost of release, as would be the case if there were race-specific prediction errors (e.g., judges systematically overestimate the cost of release for black defendants relative to white defendants).

In Appendix C, we develop this alternative model of judge behavior. A model motivated by racially-biased prediction errors yields similar predictions as a model of taste-based discrimination. An outcome-based test based on this alternative model yields the empirical prediction that larger treatment effects for marginal white defendants compared to marginal black defendants implies that judges systematically overestimate the cost of release for black defendants relative to white defendants.

Racial Differences in Arrest Probability: The outcome-based test developed above requires that any measurement error in the outcome is uncorrelated with race. Thus, we assume that the probability of pre-trial misconduct, as proxied for by the probability of rearrest prior to case disposition, is measured similarly for white and black defendants on the margin of release.

[9] Results are similar using less parametric specifications. For example, we find that the maximum estimation bias is 1.5 percentage points in a specification where we break the judge leniency measure into 100 separate bins.

This assumption would be violated if the police are more likely to rearrest black defendants who commit a new crime compared to white defendants who commit a new crime. In this scenario, our outcome-test would incorporate the potential bias of the police. If the police are racially biased against black defendants, we will overestimate the probability of pre-trial misconduct of black versus white defendants, biasing our outcome-based test towards a finding of no racial bias towards blacks. Our estimates of racial bias D are therefore likely to represent a lower bound of the true level of racial bias in bail setting.

Judge Preferences for Non-Race Characteristics: Judge behavior may also vary across non-race characteristics such as crime type, i.e., $t_r^j(\mathbf{V}_i)$ can vary with race or non-race defendant characteristics. For example, some judges may attribute a lower benefit to releasing defendants charged with violent offenses, leading to fewer violent defendants being released in equilibrium for those judges. As a result, released violent defendants for these judges will have lower expected probabilities of pre-trial misconduct. If black defendants are also more likely to be charged with violent offenses, then our outcome-based test will capture the combined effect of racial bias and this "offense type" bias.

This issue suggests two conceptually distinct outcome-based tests. Our preferred test for racial bias incorporates both direct racial bias and any indirect bias that results from harsher treatment of black defendants, on average, through different treatment of defendants based on characteristics such as crime type that are nonetheless correlated with race. We consider both direct and indirect forms of bias to be racial bias because differences in release thresholds by crime type might themselves be racially motivated, i.e., the treatment of drug possession versus DUIs.

A second form of the outcome-based test isolates only the direct racial bias holding fixed any non-race characteristics such as crime type. We control for differences in observable characteristics following the approach developed by Barsky et al. (2002) and Chandra and Staiger (2010). Specifically, we re-weight our data such that black and white defendants have the same average characteristics. To account for unobserved differences between black and white defendants, we assume that judge preferences vary only by observable characteristics:

Assumption 5. [Threshold Independence of Unobservables]. $t_r^j(\mathbf{V}_i) = t_r^j(\mathbf{X}_i)$

Assumption 5 asserts that any unobservable heterogeneity which enters judges' preferences is independent of race. If there is an unobserved characteristic which is correlated with race and this characteristic results in a specific release threshold, then the assumption is violated.

These weighted estimates allow us to test for direct racial bias among black and white defendants with the same distribution of observable characteristics:

Proposition 3. Let α_r^ω be the instrumental variables estimate of the effect of pre-trial release on pre-trial misconduct for defendants of race r on the re-weighted sample. If Assumption 5 is satisfied, then the average level of taste-based bias is given by:

$$D^\omega = \alpha_W^\omega - \alpha_B^\omega \tag{10}$$

Proof. See Appendix B.

Under this re-weighting procedure, we can interpret any differences in expected outcomes for the marginal black and marginal white defendant as due to taste-based racial bias on the part of judges. Alternatively, the re-weighted results may be interpreted as a partial correction for differences in other characteristics by defendant race. For instance, these re-weighted estimates will eliminate any treatment effect differences due to differences in crime types by race.

However, we emphasize that this procedure may lead to an underestimate of the true level of racial bias in the bail setting in some plausible scenarios. For example, if drug cases are held to stricter standards than property cases due to implicit racial bias, then re-weighting the data will eliminate this implicit source of racial bias.

II. Data and Instrument Construction

This section summarizes the most relevant information regarding our administrative court data from Philadelphia and Miami-Dade and the construction of our judge leniency measure. Further details on the cleaning and coding of variables are contained in Appendix D.

A. Data Sources and Descriptive Statistics

Philadelphia court records are available for all defendants arrested and charged between 2010-2014 and Miami-Dade court records are available for all defendants arrested and charged between 2006-2014. For both jurisdictions, the court data contain information on defendant's name, gender, race, date of birth, and zip code of residence. Because our ethnicity identifier does not distinguish between non-Hispanic white and Hispanic white, we match the surnames in our dataset to census genealogical records of surnames. If the probability a given surname is Hispanic is greater than 80 percent, we label this individual as Hispanic. In our main analysis, we include all defendants and compare outcomes for marginal black and marginal white (Hispanic and non-Hispanic) defendants. In robustness checks, we present results comparing marginal black and marginal non-Hispanic white defendants.[10]

The court data also include information on the original arrest charge, the filing charge, and the final disposition charge. We also have information on the severity of each charge based on state-specific offense grades, the outcome for each charge, and the punishment for each guilty charge. Finally, the case-level data include information on attorney type, arrest date, and the date of and judge presiding over each court appearance from arraignment to sentencing. Importantly, the case-level data also include information on bail type, bail amount when monetary bail is set, and whether bail was met. Because the data contain defendant identifiers, we can measure whether a defendant committed pre-trial misconduct by whether the defendant was subsequently arrested for a new crime before the case was resolved.

[10] Appendix Table A1 presents results for marginal Hispanic defendants compared to non-Hispanic white defendants. Perhaps in some part because of measurement error in our coding of Hispanic ethnicity, we find no evidence of racial bias against Hispanics.

We make three restrictions to the court data to isolate cases that are quasi-randomly assigned to judges. First, we drop a small set of cases with missing bail judge information. Second, we drop the 30 percent of defendants in Miami-Dade who never have a bail hearing because they post bail immediately following the arrest; below we show that the characteristics of defendants who have a bail hearing are uncorrelated with our judge leniency measure. Third, we drop all weekday cases in Miami-Dade because, as explained in Appendix E, bail judges in Miami-Dade are assigned on a quasi-random basis only on the weekends. The final sample contains 193,431 cases from 116,583 unique defendants in Philadelphia and 93,572 cases from 66,003 unique defendants in Miami-Dade.

Table 1 reports summary statistics for our estimation sample separately by race and pre-trial release status measured at three days of the bail hearing, as recent policy initiatives focus on this time period. On average, black defendants are more 11.2 percentage points more likely to be assigned monetary bail compared to white defendants and receive bail amounts that are $14,376 greater than white defendants. Compared to white defendants, released black defendants are also 6.4 percentage points more likely to be rearrested for a new crime before case disposition.

B. Construction of the Instrumental Variable

We estimate the causal impact of pre-trial release for the marginal defendant using a measure of the tendency of a quasi-randomly-assigned bail judge to release a defendant pre-trial as an instrument for release. In both Philadelphia and Miami-Dade, there are multiple bail judges serving at each point in time in both jurisdictions, allowing us to utilize variation in bail setting across judges. Both jurisdictions also assign cases to bail judges in a quasi-random fashion in order to balance caseloads: Philadelphia utilizes a rotation system where three judges work together in five day shifts, with one judge working an eight-hour morning shift (7:30AM-3:30PM), another judge working the eight-hour afternoon shift (3:30PM-11:30PM), and the final judge working the eight-hour evening shift (11:30PM-7:30AM). Similarly, bail judges in Miami-Dade rotate through the weekend felony and misdemeanor bail hearings. Additional details on the setting can be found in Appendix E.

We construct our instrument using a residualized, leave-out judge leniency measure that accounts for case selection following Dahl et al. (2014) and Dobbie et al. (2016). Because the judge assignment procedures in Philadelphia and Miami-Dade are not truly random as in other settings, selection may impact our estimates if we used a simple leave-out mean to measure judge leniency following the previous literature (e.g., Kling 2006, Aizer and Doyle 2015). For example, bail hearings following DUI arrests disproportionately occur in the evenings and on particular days of the week, leading to case selection. If certain bail judges are more likely to work evening or weekend shifts due to shift substitutions, the simple leave-out mean will be biased.

Given the rotation systems in both counties, we account for court-by-bail year-by-bail day of week fixed effects and court-by-bail month-by-bail day of week fixed effects. In Philadelphia, we add additional bail-day of week-by-bail shift fixed effects. Including these exhaustive court-by-time effects effectively limits the comparison to defendants at risk of being assigned to the same set of judges. With the inclusion of these controls, we can interpret the within-cell variation in

the instrument as variation in the propensity of a quasi-randomly assigned bail judge to release a defendant relative to the other cases seen in the same shift and/or same day of the week.

Let the residual pre-trial release decision after removing the effect of these court-by-time fixed effects be denoted by:

$$Released^*_{ict} = Released_{ic} - \gamma \mathbf{X}_{ict} = Z_{ctj} + v_{ict} \tag{11}$$

where $\mathbf{X}_{ict}$ includes the respective court-by-time fixed effects. The residual release decision, $Released^*_{ict}$, includes our measure of judge leniency Z_{ctj}, as well as unobserved defendant level variation v_{ict}.

For each case, we then use these residual bail release decisions to construct the leave-out mean decision of the assigned judge within a bail year:

$$Z_{ctj} = \left(\frac{1}{n_{tj} - n_{itj}}\right)\left(\sum_{k=0}^{n_{tj}}(Released^*_{ikt}) - \sum_{c=0}^{n_{itj}} Released^*_{ict}\right) \tag{12}$$

where n_{tj} is the number of cases seen by judge j in year t and n_{itj} is the number of cases of defendant i seen by judge j in year t. We calculate the instrument across all case types (i.e., both felonies and misdemeanors), but allow the instrument to vary across years. In robustness checks, we allow judge tendencies to vary by defendant race.

The leave-out judge measure given by Equation (12) is the release rate for the first assigned judge after accounting for the court-by-time fixed effects. This leave-out measure is important for our analysis because regressing outcomes for defendant i on our judge leniency measure without leaving out the data from defendant i would introduce the same estimation errors on both the left- and right-hand side of the regression and produce biased estimates of the causal impact of being released pre-trial. In our two-stage least squares results, we use our predicted judge leniency measure, Z_{ctj}, as an instrumental variable for whether the defendant is released pre-trial.

Figure 1 presents the distribution of our residualized judge leniency measure for pre-trial release at the judge-by-year level for all defendants, white defendants, and black defendants. Our sample includes seven total bail judges in Philadelphia and 170 total bail judges in Miami-Dade. In Philadelphia, the average number of cases per judge is 27,633 during the sample period of 2010-2014, with the typical judge-by-year cell including 6,239 cases. In Miami-Dade, the average number of cases per judge is 550 during the sample period of 2006-2014, with the typical judge-by-year cell including 187 cases. Controlling for our vector of court-by-time effects, the judge release measure ranges from -0.164 to 0.205 with a standard deviation of 0.036. In other words, moving from the least to most lenient judge increases the probability of pre-trial release by 37.1 percentage points, a 72.3 percent change from the mean three-day release rate of 50.6 percentage points.

In practice, variation in our judge leniency measure comes from a combination of different bail decisions. Dobbie et al. (2016) show that defendants assigned to judges who are more likely to use conditional non-monetary bail are more likely to be released before trial. Conversely, defendants assigned to judges who are more likely to use monetary bail and assign higher monetary bail amounts

are less likely to be released pre-trial. In contrast, there is no significant relationship between the residualized judge leniency measure for ROR and the probability of pre-trial release. In combination, these results suggest that defendants on the margin of pre-trial release are those for whom judges disagree about the appropriateness of conditional non-monetary bail versus monetary bail, not those for whom judges disagree about the appropriateness of ROR versus other bail decisions.

One question might be why judges differ in their bail decisions. While interesting for thinking about the design of the bail determination process, it is not critical to our analysis to know precisely why some judges are more lenient than others. What is critical is that some judges are systematically more lenient than others, that cases are randomly assigned to judges conditional on our court-by-time fixed effects, and that defendants released by a strict judge would also be released by a lenient one. We consider below whether each of these conditions holds in our data.

Another question is how many and what types of defendants are compliers in our setting. In Dobbie et al. (2016), we find that 13 percent of defendants in our sample are "compliers," meaning that they would have received a different bail outcome had their case been assigned to the most lenient judge instead of the most strict judge. In comparison, 53 percent of our sample are "never takers," meaning that they would be detained by all judges, and 34 percent are "always takers," meaning that they would be released pre-trial regardless of the judge assigned to the case. In Appendix Table A2, we describe the characteristics of compliers in our sample following the approach developed by Abadie (2003) and extended by Dahl et al. (2014). Compliers in our sample are 12 percentage points more likely to be charged with a misdemeanor and 17 percentage points more likely to be charged with non-violent offenses compared to the average defendant. Compliers are not systematically different from the average defendant by race or prior criminal history, however.

C. Instrument Validity

Existence and Linearity of First Stage: To examine the first stage relationship between bail judge leniency and whether a defendant is released pre-trial (*Released*), we estimate the following equation for individual i and case c, assigned to judge j at time t using a linear probability model:

$$Released_{ictj} = \gamma_0 + \gamma_1 Z_{ctj} \cdot White_i + \gamma_2 Z_{ctj} \cdot Black_i + \pi \mathbf{X}_{ict} + \mathbf{v}_{ict} \tag{13}$$

where the vector $\mathbf{X}_{ict}$ includes court-by-time fixed effects. The error term $\mathbf{v}_{ict} = \mathbf{U}_i + \varepsilon_{ict}$ is composed of characteristics unobserved by the econometrician but observed by the judge, $\mathbf{U}_i$, and idiosyncratic variation, ε_{ict}. As described previously, Z_{ctj} are leave-out (jackknife) measures of judge leniency that are allowed to vary across years. Robust standard errors are two-way clustered at the individual and judge-by-shift level.

Figure 1 provides graphical representations of the first stage relationship, pooled and separately by race, between our residualized measure of judge leniency and the probability of pre-trial release controlling for our exhaustive set of court-by-time fixed effects, overlaid over the distribution of judge leniency. The graphs are a flexible analog to Equation (13), where we plot a local linear

regression of actual individual pre-trial release against judge leniency. The individual rate of pre-trial release is monotonically increasing for both races, and approximately linearly, increasing in our leniency measure.

Table 2 presents formal first stage results from Equation (13) for all defendants, white defendants, and black defendants. Columns 1, 3, and 5 begin by reporting results only with court-by-time fixed effects. Columns 2, 4, and 6 add our baseline crime and defendant controls: race, gender, age, whether the defendant had a prior offense in the past year, the number of charged offenses, indicators for crime type (drug, DUI, property, violent, other) and crime severity (felony or misdemeanor), and indicators for missing characteristics.

We find that our residualized judge instrument is highly predictive of whether a defendant is released pre-trial, with an F-statistic for the instrument of 501.8. Our results show that a defendant assigned to a bail judge that is 10 percentage points more likely to release a defendant pre-trial is 5.9 percentage points more likely to be released pre-trial. Judge leniency is also highly predictive of pre-trial release for both white and black defendants. A white defendant assigned to a bail judge that is 10 percentage points more likely to release a defendant pre-trial is 5.4 percentage points more likely to be released pre-trial and a black defendant assigned to a bail judge that is 10 percentage points more likely to release a defendant pre-trial is 6.4 percentage points more likely to be released pre-trial.

Exclusion Restriction: Table 3 verifies that assignment of cases to bail judges is random after we condition on our court-by-time fixed effects. Columns 1, 3, and 5 of Table 3 uses a linear probability model to test whether case and defendant characteristics are predictive of pre-trial release. These estimates capture both differences in the bail conditions set by the bail judges and differences in these defendants' ability to meet the bail conditions. We control for court-by-time fixed effects and two-way cluster standard errors at the individual and judge-by-shift level. For example, we find that black male defendants are 12.6 percentage points less likely to be released pre-trial compared to similar female defendants, while white male defendants are 11.5 percentage points less likely to be released pre-trial compared to similar female defendants. White defendants with a prior offense in the past year are 20.1 percentage points less likely to be released compared to defendants with no prior offense, while black defendants with a prior offense in the past year are 14.5 percentage points less likely to be released compared to defendants with no prior offense. Columns 2, 4, and 6 assess whether these same case and defendant characteristics are predictive of our judge leniency measure using an identical specification. We find that judges with differing leniencies are assigned cases with very similar defendants.

Even with random assignment, the exclusion restriction could be violated if bail judge assignment impacts the probability of pre-trial misconduct through channels other than pre-trial release. The assumption that judges only systematically affect defendant outcomes through pre-trial release is fundamentally untestable, and our estimates should be interpreted with this potential caveat in mind. However, we argue that the exclusion restriction assumption is reasonable in our setting. Bail judges exclusively handle one decision, limiting the potential channels through which they

could affect defendants. In addition, we are specifically interested in short-term outcomes (pre-trial misconduct) which occur prior to disposition, further limiting the role of alternative channels that could affect longer-term outcomes.

Monotonicity: The final condition needed to interpret our estimates as the LATE of pre-trial release is that the impact of judge assignment on the probability of pre-trial release is monotonic across defendants. In our setting, the monotonicity assumption requires that individuals released by a strict judge would also be released by a more lenient judge and that individuals detained by a lenient judge would also be detained by a stricter judge. If the monotonicity assumption is violated, our two-stage least squares estimates would still be a weighted average of pairwise local average treatment effects, but the weights would not sum to one (Angrist et al. 1996, Heckman and Vytlacil 2005). The monotonicity assumption is therefore necessary to interpret our estimates as a well-defined LATE.

An implication of the monotonicity assumption is that the first stage estimates should be non-negative for all subsamples. Appendix Table A3 present these first stage results using the full sample of cases to calculate our measure of judge leniency. We find that our residualized measure of judge leniency is consistently non-negative and sizable in all subsamples, in line with the monotonicity assumption. Appendix Figure A1 further explores how judges treat cases of observably different defendants by plotting our residualized judge leniency measures calculated separately by offense type, offense severity, and prior criminal history. Each plot reports the coefficient and standard error from an OLS regression relating each measure of judge leniency. Consistent with our monotonicity assumption, we find that the slopes relating the relationship between judge leniency in one group and judge leniency in another group are non-negative, suggesting that judge tendencies are similar across observably different defendants and cases.

III. Results

In this section, we present our main results applying our empirical test for racial bias. We then compare the results from our empirical test with the alternative outcome-based tests developed by Knowles et al. (2001) and Anwar and Fang (2006).

A. Empirical Tests for Racial bias

We apply our proposed method to estimate the probability of pre-trial misconduct for white and black defendants on the margin of release. Specifically, we estimate the following two-stage least squares specification for individual i and case c, assigned to judge j at time t:

$$Y_{ict} = \beta_0 + \alpha_W Released_{ic} \cdot White_i + \alpha_B Released_{ic} \cdot Black_i + \beta_1 \mathbf{X}_{ict} + \mathbf{v}_{ict} \tag{14}$$

where the vector $\mathbf{X}_{ict}$ includes court-by-time fixed effects and defendant gender, age, whether the defendant had a prior offense in the past year, the number of charged offenses, indicators for crime type (drug, DUI, property, violent, or other), crime severity (felony or misdemeanor), and indicators

for any missing characteristics. As described previously, the error term $\mathbf{v}_{ict} = \mathbf{U}_i + \varepsilon_{ict}$ consists of characteristics unobserved by the econometrician but observed by the judge, $\mathbf{U}_i$, and idiosyncratic variation unobserved by both the econometrician and judge, ε_{ict}. We instrument for pre-trial release with the interaction of defendant race and our measure of judge leniency, Z_{ctj}. Robust standard errors are two-way clustered at the individual and judge-by-shift level.

Table 4 presents estimates of Equation (14). Columns 1-2 reports two-stage least squares estimates of the causal effect of pre-trial release on the probability of rearrest prior to case disposition for marginal white defendants, α_W, and marginal black defendants, α_B, respectively. Column 3 reports the difference between α_W and α_B. Panel A presents results for the probability of rearrest for any crime prior to case disposition, while Panel B presents results for rearrest rates for drug, property, and violent offenses separately. In total, 20.8 percent of defendants are rearrested for a new crime prior to disposition, with 9.1 percent of defendants being rearrested for drug offenses and 5.9 percent of defendants being rearrested for property offenses.

We find convincing evidence of racial bias against black defendants. In Panel A, we find that marginally released white defendants are 18.5 percentage points more likely to be rearrested for any crime compared to marginally detained white defendants (column 1). In contrast, the effect of pre-trial release on rearrest rates for the marginally released black defendants is a statistically insignificant 0.5 percentage points (column 2). Taken together, these estimates imply that marginally released white defendants are 18.0 percentage points more likely to be rearrested prior to disposition than marginally released black defendants (column 3), consistent with racial bias against blacks. Importantly, we can reject the null hypothesis of no racial bias even assuming the maximum potential bias in our IV estimator (see Appendix B).

In Panel B, we find suggestive evidence of racial bias against black defendants across all crime types, although the point estimates are too imprecise to make definitive conclusions. Most strikingly, we find that marginally released white defendants are 9.7 percentage points more likely to be rearrested for a drug crime prior than marginally released black defendants (p-value = 0.024). Marginally released white defendants are also 3.0 percentage points more likely to be rearrested for a property crime compared to marginally released black defendants (p-value = 0.579), and marginally released whites are about 8.2 percentage points more likely to be rearrested for a violent crime prior to disposition than marginally released blacks (p-value = 0.036).

In Appendix Table A4, we present results comparing outcomes for marginal non-Hispanic white defendants and marginal black defendants. We find very similar results consistent with racial bias against black defendants. Overall, these findings indicate significant racial bias against black defendants, driven largely by differences in the probability of committing a new drug crime for marginal white and marginal black defendants.

B. Subsample Results

To explore heterogeneous treatment effects, we combine all observable demographic and crime characteristics into a single risk index. In Table 5, we divide defendants into above and below median

predicted risk, with those in the below median group having a 13.0 percent probability of rearrest prior to case disposition compared to 31.0 percent among defendants in the above median group.[11] We find that racial bias against black defendants is almost exclusively driven by those with the highest predicted risk of rearrest. Among high-risk defendants, marginally released white defendants are 38.5 percentage points more likely to be rearrested prior to case disposition than marginally released black defendants (p-value = 0.008). In contrast, we find no evidence of racial bias against black defendants among low-risk defendants (p-value = 0.967).

In Appendix Tables A5-A8, we explore additional subsample results. In Appendix Table A5, we analyze whether racial bias against black defendants is larger among those charged with drug offenses versus non-drug offenses. This subsample split is of particular interest because black defendants in our sample are more likely to be charged with drug offenses compared to white defendants, and conditional on being charged with a drug offense, are less likely to be released before trial. We find that our main results are largely driven by the differential treatment of white and black defendants charged with drug offenses. Among drug offenders, marginally released white defendants are 36.0 percentage points more likely to be rearrested prior to case disposition than marginally released black defendants (p-value = 0.024). In contrast, we find limited evidence of racial bias among defendants arrested for all other non-drug crimes (p-value = 0.313).

Another important dimension on which black and white defendants differ, and which affects the likelihood of pre-trial release, is the likelihood of having a prior offense from the last year. In Appendix Table A6, we find evidence that racial bias against black defendants is also driven by defendants with a prior in the past year. Among prior offenders, marginally released white defendants are 31.1 percentage points more likely to be rearrested prior to case disposition than marginally released black defendants (p-value = 0.014), whereas we find limited evidence of racial bias among defendants with no recent priors (p-value = 0.434). In Appendix Tables A7-A8, we also find that racial bias against black defendants is larger among defendants charged with felonies (p-value = 0.011) and defendants from below median income zip codes (p-value = 0.058).

C. Robustness

Our main results are robust to a number of alternative specifications. In Appendix Table A9, we present analogous re-weighted two-stage least squares with the weights chosen to match the distribution of observable characteristics by race (see Proposition 3). After re-weighting on observables, we find that marginally released white defendants are 15.9 percentage points more likely to be rearrested prior to case disposition than marginally released black defendants (p-value = 0.061), driven largely by differences in rearrest rates for drug crimes among marginal white and marginal black defendants (p-value = 0.025). These results indicate that even after accounting for differences in other observable characteristics by defendant race, bail judges appear to be directly racially biased

[11] In small samples, endogenous stratification may lead to biased results. See Abadie, Chingos, and West (2014). In our setting, given the large sample size, we find identical results if we use a split-sample estimator to predict risk in a five percent random sample and estimate our two-stage least squares results in the remaining 95 percent.

against black defendants.

In Appendix Table A10, we present our main results clustering more conservatively at the individual and judge level. In Appendix Table A11, we reestimate the main results using a version of our instrument constructed separately for white and black defendants. By calculating the instrument separately by defendant race, we relax the monotonicity assumption and specifically allow for judge tendencies to vary across white and black defendants. Under these alternative specifications, we continue to find that marginally released white defendants are significantly more likely to be rearrested prior to disposition than marginally released black defendants, evidence of racial bias against black defendants.

D. Comparison to Other Outcome Tests

In this section, we replicate the alternative outcome tests from the seminal papers of Knowles et al. (2001) and Anwar and Fang (2006). In our context, the test of Knowles et al. (2001) relies on the prediction that under the null hypothesis of no racial bias, the average pre-trial misconduct rate does not depend on defendant race. The test of Anwar and Fang (2006) relies on the prediction that under the null hypothesis of no relative racial bias, the relative treatment of white defendants relative to black defendants does not depend on judge race.

Appendix Table A12 replicates the Knowles et al. (2001) test for absolute racial bias. To implement this test, we estimate an OLS regression of pre-trial release on the probability of rearrest for black and white defendants in our full sample. This OLS specification compares the average rearrest rates for black and white defendants conditional on observables. In contrast to our preferred IV test, the OLS results indicate that judges are not racially biased against black defendants (p-value = 0.424), implying that either there are omitted variables biasing the OLS estimates or that the marginal effect of pre-trial release is not equal to the average effect of pre-trial release. While it is not possible to distinguish between these explanations using our data, these results suggest that the Knowles et al. (2001) test may be invalid in our setting.

Appendix Tables A13-A14 replicate the Anwar and Fang (2006) test for relative racial bias. We are unable to use data from Philadelphia for this test, as all seven Philadelphia judges in our sample are white, making it impossible to detect relative racial bias. We therefore restrict the sample to defendants in Miami where we have both white (Hispanic and non-Hispanic) and black judges.[12]

Appendix Table A13 presents average release rates and average rearrest rates conditional on release by both judge and defendant race. Appendix Table A14 then presents bootstrapped p-values from a test of relative racial bias, that is, whether white judges are more lenient on white defendants than black defendants and black judges are more lenient on black defendants than white defendants, against the null hypothesis of no reversal in relative treatment by judge race.

Unlike Anwar and Fang (2006), we find very little evidence of non-monolithic judge behavior.

[12]To implement this test, we collected information on the race of each bail judge in our sample using official court directories and internet searches. In Miami, there are 91 white judges, 61 Hispanic judges, and 15 black judges. See Appendix D for additional details.

Instead, we find that judges do not differ substantially in their treatment of black versus white defendants. For example, Panel A of Appendix Table A13 indicates that 34.5 percent of white defendants are released by white judges and 33.9 percent of white defendants are released by black judges. Similarly, black defendants are generally less likely to be released by both white judges (31.1 percent) and black judges (31.8 percent). Perhaps unsurprisingly then, we find no evidence of relative bias using the Anwar and Fang (2006) test. In Appendix Table A14, we are unable to reject the null hypothesis of no relative bias in either pre-trial release rates (p-value = 0.364) or rearrest rates conditional on release (p-value = 0.412). As a result, in sharp contrast to our preferred IV results, we cannot reject the null hypothesis of no relative bias using the method proposed by Anwar and Fang (2006). Rather, these results suggest that both white and black judges are racially biased against black defendants. Consistent with this explanation, in results available upon request, we find that our IV estimate of racial bias is similar among white and black judges in Miami, although these confidence intervals for these estimates are extremely large, making definitive conclusions impossible. Finally, in other results available upon request, we find no evidence of relative bias if we disaggregate white judges and defendants into Hispanic whites and non-Hispanic whites.

There are three potential explanations for why the results from our IV approach differ from the outcome tests proposed by Knowles et al. (2001) and Anwar and Fang (2006). First, our IV estimator uses the quasi-random variation in the assignment of defendants to bail judges to identify causal treatment effects that are not biased by selection, while past work has relied on a strong selection on observables assumption. Second, our IV estimator specifically identifies treatment effects for marginal defendants rather than the average defendant, while past work has again had to rely on a strong homogenous treatment effects assumption. Finally, our estimator is able to detect absolute rather than relative bias and is therefore applicable in settings in which both minority and non-minority judges are similarly biased.

IV. Learning and Racial Bias

In this section, we explore whether the relative inexperience of bail judges can explain racial bias in our setting. As previously discussed, bail judges must make quick judgments on the basis of limited information and virtually no training, leading some to call bail judges the "amateur link in the criminal justice chain." For example, some jurisdictions do not require bail judges to have any legal education or certification other than an one-day training session.[13] In other jurisdictions, bail hearings are heard by "generalist" judges that have no specific training in bail setting and who only assist with bail hearings a few days a year.

[13] See, for example, https://bangordailynews.com/2011/03/22/business/maine's-bail-system-a-19th-century-holdoverpart-1-of-4people-who-set-bail-in-maine-have-almost-no-legal-training/?ref=series. Recent reforms include increased training for bail judges and mandatory review of all bail determinations by a second judge. See http://www.nytimes.com/2015/10/02/nyregion/jonathan-lippman-bail-incarceration-new-york-state-chief-judge.html?_r=0. Other jurisdictions encourage new bail judges to shadow experienced ones. See http://pinetreewatchdog.org/maines-bail-system-best-state-can-afford-or-a-threat-to-due-process/.

There are a number of reasons why inexperience may influence racial bias. First, it is possible that these relatively inexperienced bail judges may be more likely to use race as a proxy for risk, leading to the relative over-detention of black defendants on the margin. Consistent with this "prediction error" channel, Bordalo et al. (2016) show that the use of heuristics in probability judgments can lead to stereotypes that amplify systematic differences between groups even when these differences are very small, e.g., the risk of pre-trial misconduct for black and white defendants. Similarly, Fryer and Jackson (2008) show that categorical decision-making can lead to biased decisions even when there is no "taste" for bias. The use of these kinds of simple race-based heuristics or categorical decision-making by bail judges may decrease with on-the-job learning, however, as information on whether a defendant commits a new crime before case disposition is salient and generally known to the bail judge. In turn, it is possible that the decreased use of these heuristics or categories could decrease racial bias among those judges.[14]

A second reason experience may influence racial bias is that bail judges' "taste" for bias may change with increased exposure to black defendants. There is a large literature suggesting that intergroup contact can increase tolerance towards minority groups. For example, Van Laar et al. (2005) and Boisjoly et al. (2006) show that living with a minority group increases tolerance among white college students, Dobbie and Fryer (2013) show that teaching in a school with mostly minority children increases racial tolerance, and Clingingsmith et al. (2009) show that winning a lottery to participate in the Hajj pilgrimage to Mecca increases belief in equality and harmony of ethnic groups. These studies suggest that the increased exposure to black defendants that comes with on-the-job experience may decrease any taste-based bias in bail setting.

We explore the combined importance of these channels in three ways: (1) by examining racial bias among bail judges who do and do not specialize in bail setting, (2) by examining racial bias among more and less experienced judges who are all non-specialists, and (3) by descriptively examining how experienced and inexperienced judges use observable characteristics to make release decisions.

Specialist versus Generalist Judges: We begin by taking advantage of the institutional differences between the Philadelphia and Miami-Dade bail systems. In Philadelphia, bail judges are full-time judges who serve on a court that specializes in setting bail 24 hours a day, seven days a week, with Philadelphia judges in our sample hearing an average of 6,239 cases each year. In sharp contrast, the Miami bail judges that are in our sample are part-time generalists who work as trial court judges on weekdays and "help out" the bail court on weekends. As a result, the Miami judges in our sample hear an average of only 187 bail cases each year. If inexperience plays a role in explaining racial bias, the degree of racial bias as estimated under our test should be different across the two jurisdictions.

Columns 1-3 of Table 6 presents estimates separately by court. Column 1 reports the difference in pre-trial misconduct rates for marginal white and marginal black defendants in Miami, column

[14]Consistent with this kind of on-the-job learning, there is evidence that new bail judges rely heavily on external information while more experienced judges primarily rely on their own judgment. As one bail judge told ProPublica, "the [risk] scores were helpful when he was a new judge, but now that he has experience he prefers to rely on his own judgment." See https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

2 reports the difference in pre-trial misconduct rates for marginal white and marginal black defendants in Philadelphia, and column 3 reports the difference between the two jurisdictions. We find suggestive evidence that racial bias is higher in Miami than Philadelphia (p-value = 0.094). In Miami, marginally released white defendants are 29.1 percentage points more likely to be rearrested compared to marginally released black defendants (p-value = 0.071). In Philadelphia, we find no statistically significant evidence of racial bias.

Inexperienced versus Experienced Judges: Next, we explore the role of learning and experience within only the Miami system. There is substantial variation in the experience profiles of the Miami bail judges in our sample: splitting by the median number of years hearing bail cases, the experienced Miami judges have an average 9.5 years of experience in the bail system, compared to an average of only 2.5 years among the inexperienced Miami judges. Columns 4 through 6 of Table 6 present results separately for Miami judges with above- and below-median levels of experience. We find suggestive evidence that inexperienced judges are more racially biased than experienced judges (p-value = 0.290). Among inexperienced judges, marginally released white defendants are 48.0 percentage points more likely to be rearrested compared to marginally released black defendants (p-value = 0.071). Among experienced judges, marginally released white defendants are 19.3 percentage points more likely to be rearrested compared to marginally released black defendants (p-value = 0.228).

Use of Observables in Bail Decisions: We conclude with a more descriptive analysis that asks whether judges make different bail decisions on the basis of observables as they gain more on-the-job experience. Ideally, we would estimate a series of two-stage least squares estimates separately by both race and the relevant observable characteristics. In practice, however, these subsample estimates are too imprecise to be informative given the available data. We therefore proceed with a more descriptive analysis, with the caveat that it is difficult to draw definitive conclusions from these descriptive results alone given the obvious selection and inframarginality concerns.

Figure 2 plots the relationship between the probability of pre-trial release and the predicted risk of pre-trial misconduct based on the full set of observables. Predicted risk is calculated using an OLS regression of the probability of rearrest on the full set of observables for defendants who are released before trial. The smoothed line in Figure 2 is then estimated using a local linear regression of pre-trial release against predicted risk. Panel A presents the relationship for all defendants, Panel B presents the relationship for white defendants, and Panel C presents the relationship for black defendants. In general, there is an expected negative relationship between predicted risk and the probability of release, but in most of the risk distribution, black defendants are less likely to be released compared to white defendants. These results are further evidence that bail judges treat black defendants differently than observably similar white defendants.

In Figure 3, we examine the importance of experience by plotting the relationship between predicted risk and release separately for experienced and inexperienced judges (across both courts) by splitting based on median years of experience on the job. For both white and black defendants,

the relationship between predicted risk and release is different for experienced and inexperienced judges, suggesting that the role of observables changes with experience. In particular, the risk-release relationship is shifted up for more experienced judges, indicating that (conditional on observables) the same defendant is more likely to be released by an experienced judge than an inexperienced judge. In results available upon request, we find that the differences between experienced and inexperienced judges are primarily driven by inexperienced judges being relatively less lenient for defendants charged with violent offenses. These results are consistent with these more inexperienced judges overweighting particularly salient case and defendant characteristics (relative to experienced judges), as would be expected by simple heuristic- or category-based decision-making.

Taken together, we view these results as being consistent with a model of on-the-job learning where racial bias decreases as judges learn to place different predictive weight on the most salient case and defendant characteristics such as race and the nature of the charge. Recall that there are several potentially relevant channels for this on-the-job learning result that we cannot examine with our data, including more experienced bail judges putting less weight on inaccurate stereotypes and more experienced judges becoming less racially biased over time as they are exposed to black defendants. Unfortunately, we are unable to directly measure bail judges' racial bias or risk assessments using our data. The precise mechanisms for these results therefore remain unclear and may include a combination of these factors.

V. Conclusion

In this paper, we test for racial bias in bail setting using the quasi-random assignment of bail judges to identify pre-trial misconduct rates for marginal white and marginal black defendants. We find evidence that there is substantial bias against black defendants, with the largest bias against black defendants with the highest predicted risk of rearrest. Our estimates are nearly identical if we account for observable crime and defendant differences by race, indicating that our results cannot be explained by black-white differences in the probability of being arrested for certain types of crimes (e.g., the proportion of felonies versus misdemeanors) or black-white differences in defendant characteristics (e.g., the proportion of defendants with a prior offense versus no prior offense).

We document three facts that suggest the relative inexperience on bail judges can explain the racial bias we find in our setting. First, we find that racial bias is higher in Miami, where judges in our sample are part-time generalists that only hear a few hundred cases each year, than in Philadelphia, where judges are full-time specialists that hear thousands of cases each year. Second, we find that racial bias is higher among inexperienced Miami judges compared to experienced Miami judges. Finally, we find descriptive evidence that inexperienced bail judges overweighting particularly salient case and defendant characteristics, as would be expected by simple heuristic- or category-based decision-making. These results suggest that increasing the on-the-job experience or on-the-job feedback of bail judges provides a potential way to alleviate any unwarranted disparities.

The empirical test developed in this paper can be used to test for bias in a variety of other

settings. Our test for racial bias is appropriate when there is the quasi-random assignment of judges or examiners and the objective of these judges or examiners is both known and well-measured. Our test can therefore help to explore bias in settings as varied as parole board decisions, Disability Insurance applications, bankruptcy filings, and hospital care decisions.

References

[1] Abadie, Alberto. 2003. "Semiparametric Instrumental Variable Estimation of Treatment Response Models." Journal of Econometrics, 113(2): 231-263.

[2] Abadie, Alberto, Matthew M. Chingos, and Martin R. West. 2014. "Endogenous Stratification in Randomized Experiments." NBER Working Paper No. 19742.

[3] Abrams, David S., Marianne Bertrand, and Sendhil Mullainathan. 2012. "Do Judges Vary in Their Treatment of Race?" Journal of Legal Studies, 41(2): 347-383.

[4] Aizer, Anna, and Joseph Doyle, Jr. 2015. "Juvenile Incarceration, Human Capital, and Future Crime: Evidence from Randomly-Assigned Judges." Quarterly Journal of Economics, 130(2): 759-803.

[5] Alesina, Alberto, and Eliana La Ferrara. 2014. "A Test of Racial Bias in Capital Sentencing." American Economic Review, 104(11): 3397-3433.

[6] Angrist, Joshua, Guido Imbens, and Donald Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." Journal of the American Statistical Association, 91(434): 444-455.

[7] Antonovics, Kate, and Brian Knight. 2009. "A New Look at Racial Profiling: Evidence from the Boston Police Department." The Review of Economics and Statistics, 91(1): 163-177.

[8] Anwar, Shamena, and Hanming Fang. 2006. "An Alternative Test of Racial bias in Motor Vehicle Searches: Theory and Evidence." The American Economic Review, 96(1): 127-151.

[9] Anwar, Shamena, Patrick Bayer, and Randi Hjalmarsson. 2012. "The Impact of Jury Race in Criminal Trials." Quarterly Journal of Economics, 127(2): 1017-1055.

[10] Ayres, Ian, and Joel Waldfogel. 1994. "A Market Test for Race Discrimination in Bail Setting." Stanford Law Review, 46(5): 987-1047.

[11] Ayres, Ian. 2002. "Outcome Tests of Racial Disparities in Police Practices." Justice Research and Policy, 4: 131-142.

[12] Barsky, Robert, John Bound, Kerwin Kofi Charles, and Joseph P. Lupton. 2002. "Accounting for the Black-White Wealth Gap: A Nonparametric Approach." Journal of the American Statistical Association, 97(459): 663-673.

[13] Becker, Gary S. 1957. *The Economics of Discrimination.* Chicago: University of Chicago Press.

[14] Bertrand, Marianne, and Esther Duflo. 2016. "Field Experiments on Discrimination." NBER Working Paper No. 22014.

[15] Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Loken, and Magne Mogstad. 2016. "Incarceration, Recidivism and Employment." Unpublished Working Paper.

[16] Bushway, Shawn D., and Jonah B. Gelbach. 2011. "Testing for Racial Discrimination in Bail Setting Using Nonparametric Estimation of a Parametric Model." Unpublished Working Paper.

[17] Boisjoly, Johanne, Greg J. Duncan, Michael Kremer, Dan M. Levy, and Jacque Eccles. 2006. "Empathy or Antipathy? The Impact of Diversity." American Economic Review, 96(5): 1890-1905.

[18] Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2016. "Stereotypes." Quarterly Journal of Economics, 131(4): 1753-1794.

[19] Chandra, Amitabh, and Douglas O. Staiger. 2010. "Identifying Provider bias in Healthcare." NBER Working Paper No. 16382. National Bureau of Economic Research.

[20] Clingingsmith, David, Asim Ijaz Khwaja, and Michael Kremer. 2009. "Estimating the Impact of the Hajj: Religion and Tolerance in Islam's Global Gathering." Quarterly Journal of Economics, 124(3): 1133-1170.

[21] Demuth, Stephen. 2003. "Racial and Ethnic Differences in Pretrial Release Decisions and Outcomes: A Comparison of Hispanic, Black, and White Felony Arrestees." Criminology, 41(3): 873-908.

[22] Dahl, Gordon B., Andreas Ravndal Kostol, and Magne Mogstad. 2014. "Family Welfare Cultures." Quarterly Journal of Economics, 129(4): 1711-1752.

[23] Dobbie, Will, and Roland G. Fryer. 2013. "The Impact of Voluntary Youth Service on Future Outcomes: Evidence from Teach For America." The B.E. Journal of Economic Analysis and Policy, 15(3): 1031-1065.

[24] Dobbie, Will, Jacob Goldin, and Crystal Yang. 2016. "The Effects of Pre-Trial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges." NBER Working Paper No. 22511.

[25] Durose, Matthew R., and Patrick A. Langan. 2007. "Felony Sentences in State Courts, 2004." Bureau of Justice Statistics.

[26] Foote, Caleb. 1954. "Compelling Appearance in Court: Administration of Bail in Philadelphia." University of Pennsylvania Law Review, 102: 1031-1079.

[27] Fryer, Roland G., and Matthew Jackson. 2008. "A Categorical Model of Cognition and BiasedDecision-Making." B.E. Journal of Theoretical Economics, 8(1): 1-42.

[28] Fryer, Roland G. 2011. "Racial Inequality in the 21st Century: The Declining Significance of Discrimination." Handbook of Labor Economics, Volume 4, Part B, 855-971.

[29] Fryer, Roland G. 2016. "An Empirical Analysis of Racial Differences in Police Use of Force." NBER Working Paper No. 22399.

[30] Gelman, Andrew, Jeffrey Fagan, and Alex Kiss. 2012. "An Analysis of the New York City Police Department's "Stop-and-Frisk" Policy in the Context of Claims of Racial Bias." Journal of the American Statistical Association, 102(479): 813-823.

[31] Goldkamp, John S., and Michael R. Gottfredson. 1988. Development of Bail/Pretrial Release Guidelines in Maricopa County Superior Court, Dade County Circuit Court and Boston Municipal Court. Washington, D.C.: National Institute of Justice.

[32] Grogger, Jeffrey, and Greg Ridgeway. 2006. "Testing for Racial Profiling in Traffic Stops From Behind a Veil of Darkness." Journal of the American Statistical Association, 101(475): 878-887.

[33] Gupta, Arpit, Christopher Hansman, and Ethan Frenchman. 2016. "The Heavy Costs of High Bail: Evidence from Judge Randomization." Journal of Legal Studies, 45(2): 471-505.

[34] Heckman, James J., Sergio Urzua, and Edward Vytlacil. 2006. "Understanding Instrumental Variables in Models with Essential Heterogeneity." The Review of Economics and Statistics, 88(3): 389-432.

[35] Heckman, James J., and Edward Vytlacil. 2005. "Structural Equations, Treatment Effects, and Econometric Policy Evaluation." Econometrica, 73(3): 669-738.

[36] Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." Econometrica, 62(2): 467-475.

[37] Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. 2015. "Prediction Policy Problems." American Economic Review, 105(5): 491-95.

[38] Kling, Jeffrey. 2006. "Incarceration Length, Employment, and Earnings." American Economic Review, 96(3): 863-876.

[39] Knowles, John, Nicola Persico, and Petra Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." Journal of Political Economy, 109(1): 203-232.

[40] Leslie, Emily, and Nolan G. Pope. 2016. "The Unintended Impact of Pretrial Detention on Case Outcomes: Evidence from NYC Arraignments." Unpublished Working Paper.

[41] McIntyre, Frank, and Shima Baradaran. 2013. "Race, Prediction, and Pretrial Detention." Journal of Empirical Legal Studies, 10(4): 741-770.

[42] Rehavi, M. Marit, and Sonja B. Starr. 2014. "Racial Disparity in Federal Criminal Sentences." Journal of Political Economy, 122(6): 1320-1354.

[43] Shubik-Richards, Claire, and Don Stemen. 2010. "Philadelphia's Crowded, Costly Jails: The Search for Safe Solutions." Technical Report, Pew Charitable Trusts Philadelphia Research Initiative.

[44] Stevenson, Megan. 2016. "Distortion of Justice: How the Inability to Pay Bail Affects Case Outcomes." Unpublished Working Paper.

[45] Van Laar, Colette, Shana Levin, Stacey Sinclair, and Jim Sidanius. 2005. "The Effect of University Roommate Contact on Ethnic Attitudes and Behavior." Journal of Experimental Social Psychology, 41: 329-345.

Table 1: Descriptive Statistics

	All Defendants		White		Black	
	Detained	Released	Detained	Released	Detained	Released
Panel A: Bail Type	(1)	(2)	(3)	(4)	(5)	(6)
Release on Recognizance	0.024	0.369	0.028	0.384	0.021	0.353
Non-Monetary Bail	0.052	0.227	0.056	0.209	0.049	0.247
Monetary Bail	0.925	0.404	0.917	0.407	0.930	0.400
Bail Amount (in thousands)	53.262	15.162	43.980	17.948	60.066	12.150
Panel B: Subsequent Bail Outcomes						
Bail Modification Petition	0.463	0.056	0.458	0.050	0.466	0.064
Released in 14 days	0.079	1.000	0.087	1.000	0.074	1.000
Released before Trial	0.376	1.000	0.376	1.000	0.376	1.000
Panel C: Defendant Characteristics						
Male	0.875	0.775	0.869	0.752	0.880	0.801
Age at Bail Decision	34.357	33.987	34.990	33.959	33.893	34.020
Prior Offense in Past Year	0.385	0.220	0.376	0.193	0.392	0.251
Panel D: Charge Characteristics						
Number of Offenses	3.349	2.402	2.956	2.420	3.638	2.381
Felony Offense	0.627	0.344	0.586	0.315	0.656	0.376
Misdemeanor Only	0.373	0.656	0.414	0.685	0.344	0.624
Any Drug Offense	0.287	0.407	0.278	0.377	0.295	0.443
Any DUI Offense	0.023	0.112	0.026	0.123	0.021	0.100
Any Violent Offense	0.264	0.200	0.223	0.214	0.294	0.184
Any Property Offense	0.351	0.193	0.358	0.189	0.346	0.197
Panel E: Outcomes						
Rearrest Prior to Disposition	0.190	0.202	0.177	0.173	0.200	0.236
Drug	0.068	0.102	0.060	0.083	0.073	0.124
Property	0.068	0.042	0.069	0.037	0.068	0.047
Violent	0.047	0.022	0.037	0.018	0.054	0.026
Observations	141,689	145,314	59,917	77,678	81,772	67,636

Note: This table reports descriptive statistics for the sample of defendants from Philadelphia and Miami-Dade counties. The sample consists of bail hearings that were quasi-randomly assigned from Philadelphia between 2010-2014 and from Miami-Dade between 2006-2014. We define pre-trial release based on whether a defendant was released within the first three days after the bail hearing. Information on race, gender, age, and criminal outcomes is derived from court records. See Appendix D for additional details on the sample and variable construction.

Table 2: Judge Leniency and Pre-Trial Release

	All Defendants		White		Black	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-trial Release	0.572***	0.587***	0.533***	0.537***	0.612***	0.640***
	(0.035)	(0.033)	(0.046)	(0.043)	(0.044)	(0.041)
	[0.506]	[0.506]	[0.565]	[0.565]	[0.453]	[0.453]
Court x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Crime Controls	No	Yes	No	Yes	No	Yes
Observations	287,003	287,003	137,595	137,595	149,408	149,408

Note: This table reports first stage results. The regressions are estimated on the sample as described in the notes to Table 1. Judge leniency is estimated using data from other cases assigned to a bail judge in the same year following the procedure described in Section II.B. Columns 1, 3, and 5 begin by reporting results only with court-by-time fixed effects. Columns 2, 4, and 6 add the demographic and crime controls discussed in Section II.C. The sample mean of the dependent variable is reported in brackets. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 3: Test of Randomization

	All		White		Black	
	Pre-Trial Release	Judge Leniency	Pre-Trial Release	Judge Leniency	Pre-Trial Release	Judge Leniency
	(1)	(2)	(3)	(4)	(5)	(6)
Male	−0.12394***	0.00008	−0.11527***	0.00013	−0.12574***	0.00002
	(0.00255)	(0.00019)	(0.00322)	(0.00024)	(0.00378)	(0.00027)
Age at Bail Decision	−0.01361***	−0.00002	−0.01621***	−0.00007	−0.01234***	0.00004
	(0.00081)	(0.00007)	(0.00114)	(0.00010)	(0.00109)	(0.00009)
Prior Offense in Past Year	−0.17756***	0.00026*	−0.20655***	0.00001	−0.14502***	0.00042**
	(0.00200)	(0.00014)	(0.00292)	(0.00021)	(0.00261)	(0.00018)
Number of Offenses	−0.02685***	0.00004	−0.02320***	0.00005	−0.02712***	0.00003
	(0.00046)	(0.00003)	(0.00072)	(0.00004)	(0.00054)	(0.00003)
Felony Offense	−0.32579***	0.00002	−0.30763***	−0.00012	−0.33212***	0.00017
	(0.00277)	(0.00013)	(0.00380)	(0.00018)	(0.00365)	(0.00017)
Any Drug Offense	0.08375***	−0.00001	0.06379***	0.00004	0.10462***	−0.00004
	(0.00243)	(0.00022)	(0.00330)	(0.00027)	(0.00333)	(0.00029)
Any Violent Offense	−0.00052	−0.00013	0.04381***	−0.00048	−0.03541***	0.00023
	(0.00381)	(0.00022)	(0.00486)	(0.00031)	(0.00434)	(0.00026)
Any Property Offense	−0.02291***	−0.00026	−0.03273***	0.00002	−0.00801**	−0.00055**
	(0.00274)	(0.00022)	(0.00364)	(0.00030)	(0.00352)	(0.00028)
Joint F-Test	[0.00000]	[0.38774]	[0.00000]	[0.78955]	[0.00000]	[0.10715]
Observations	287,003	287,003	137,595	137,595	149,408	149,408

Note: This table reports reduced form results testing the random assignment of cases to bail judges. The regressions are estimated on the sample as described in the notes to Table 1. Judge leniency is estimated using data from other cases assigned to a bail judge in the same year following the procedure described in Section II.B. Columns 1, 3, and 5 report estimates from an OLS regression of pre-trial release on the variables listed and court-by-time fixed effects. Columns 2, 4, and 6 report estimates from an OLS regression of judge leniency on the variables listed and court-by-time fixed effects. The p-value reported at the bottom of the columns is for a F-test of the joint significance of the variables listed in the rows. Robust standard errors two-way clustered at the individual and the judge-by-shift level are reported in parentheses. ***=significant at 1 percent level, **=significant at 5 percent level, *=significant at 10 percent level.

Table 4: Pre-trial Release and Criminal Outcomes

	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)
Rearrest Prior to Disposition	0.185***	0.005	0.180**
	(0.067)	(0.057)	(0.087)
	[0.174]	[0.216]	–
Panel B: Rearrest by Crime Type			
Drug	0.077**	−0.020	0.097**
	(0.034)	(0.037)	(0.049)
	[0.073]	[0.096]	–
Property	0.029	−0.001	0.030
	(0.045)	(0.033)	(0.054)
	[0.051]	[0.059]	–
Violent	0.044	−0.038	0.082**
	(0.028)	(0.027)	(0.039)
	[0.026]	[0.042]	–
Observations	137,595	149,408	–

This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by race. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 5: Results for High Risk and Low Risk Offenders

	High Risk			Low Risk		
	White	Black	Difference	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)	(4)	(5)	(6)
Rearrest Prior to Disposition	0.384***	−0.002	0.385***	0.006	0.010	−0.004
	(0.120)	(0.083)	(0.145)	(0.065)	(0.068)	(0.097)
	[0.268]	[0.282]	–	[0.108]	[0.127]	–
Panel B: Rearrest by Crime Type						
Drug Crime	0.147**	−0.015	0.162**	0.005	−0.018	0.023
	(0.061)	(0.057)	(0.081)	(0.028)	(0.028)	(0.041)
	[0.127]	[0.147]	–	[0.034]	[0.027]	–
Property Crime	0.108	−0.001	0.109	−0.037	−0.005	−0.032
	(0.085)	(0.051)	(0.098)	(0.029)	(0.030)	(0.043)
	[0.093]	[0.082]	–	[0.021]	[0.028]	–
Violent Crime	0.063	−0.035	0.098*	0.024	−0.046	0.070
	(0.040)	(0.035)	(0.053)	(0.039)	(0.040)	(0.057)
	[0.024]	[0.037]	–	[0.028]	[0.048]	–
Observations	57,418	86,083	–	80,177	63,325	–

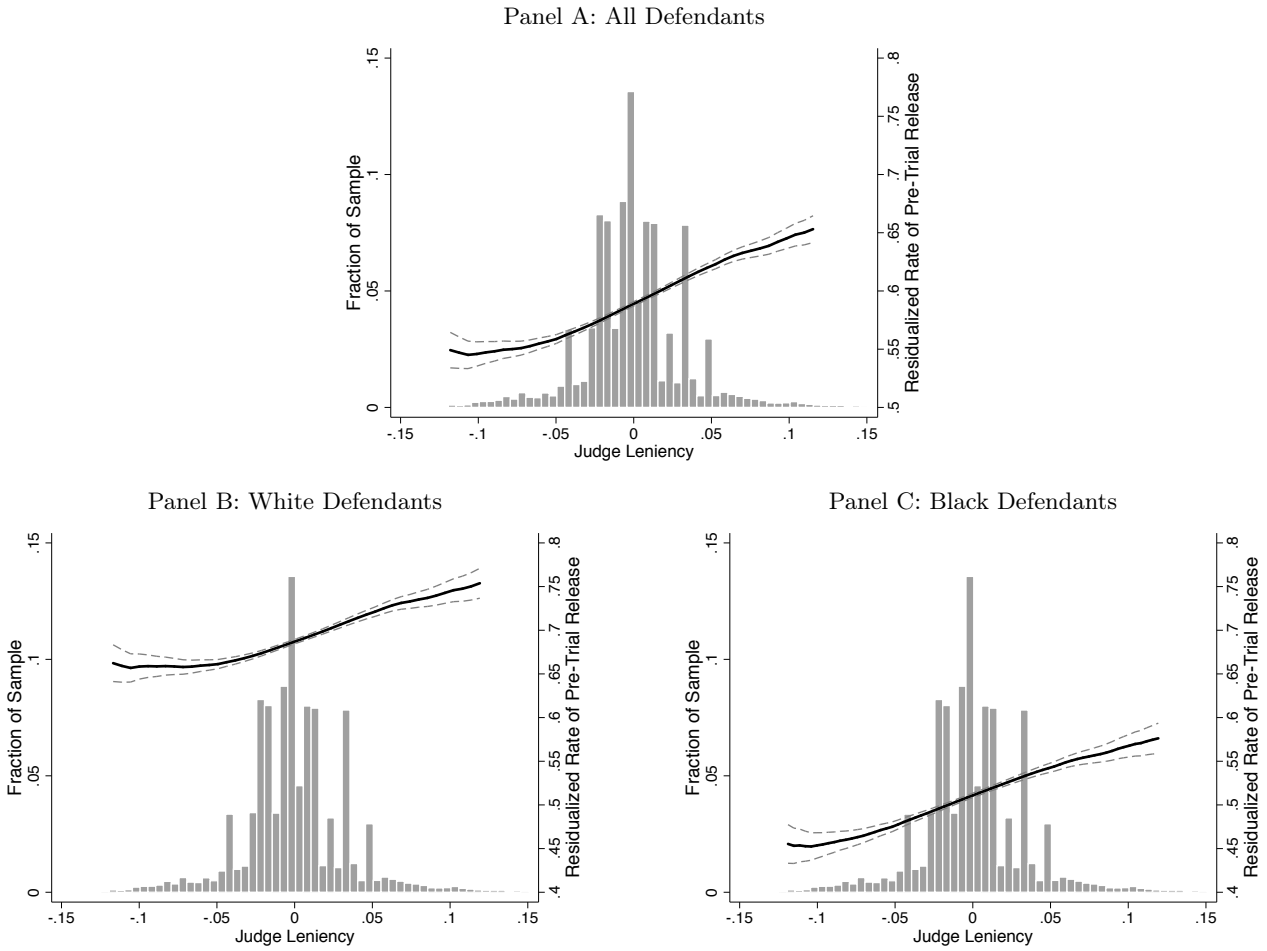
Note: This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately for high risk and low risk defendants. Risk is computed by estimating a linear regression of the probability of rearrest prior to case disposition conditional on release on the crime and demographic controls discussed in Section II. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 6: Pre-trial Release and Criminal Outcomes: The Role of Experience

	Judge Specialization			Judge Experience		
	Miami Non-Spec.	Philadelphia Specialist	Difference	Miami Low Exp.	Miami High Exp.	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)	(4)	(5)	(6)
Rearrest Prior to Disposition	0.291*	−0.014	0.305*	0.480**	0.193	0.288
	(0.161)	(0.094)	(0.182)	(0.231)	(0.160)	(0.272)
	[0.226]	[0.182]	–	[0.224]	[0.228]	–
Panel B: Rearrest by Crime Type						
Drug Crime	0.105	0.090	0.015	0.187*	0.070	0.117
	(0.065)	(0.060)	(0.085)	(0.108)	(0.080)	(0.130)
	[0.076]	[0.082]	–	[0.224]	[0.228]	–
Property Crime	0.070	−0.051	0.121	0.171	−0.028	0.198
	(0.087)	(0.045)	(0.097)	(0.142)	(0.102)	(0.169)
	[0.080]	[0.080]	–	[0.224]	[0.228]	–
Violent Crime	0.110*	0.041	0.069	0.102	0.124	−0.023
	(0.062)	(0.038)	(0.071)	(0.084)	(0.080)	(0.113)
	[0.050]	[0.047]	–	[0.224]	[0.228]	–
Observations	93,572	193,431	–	47,772	45,800	–

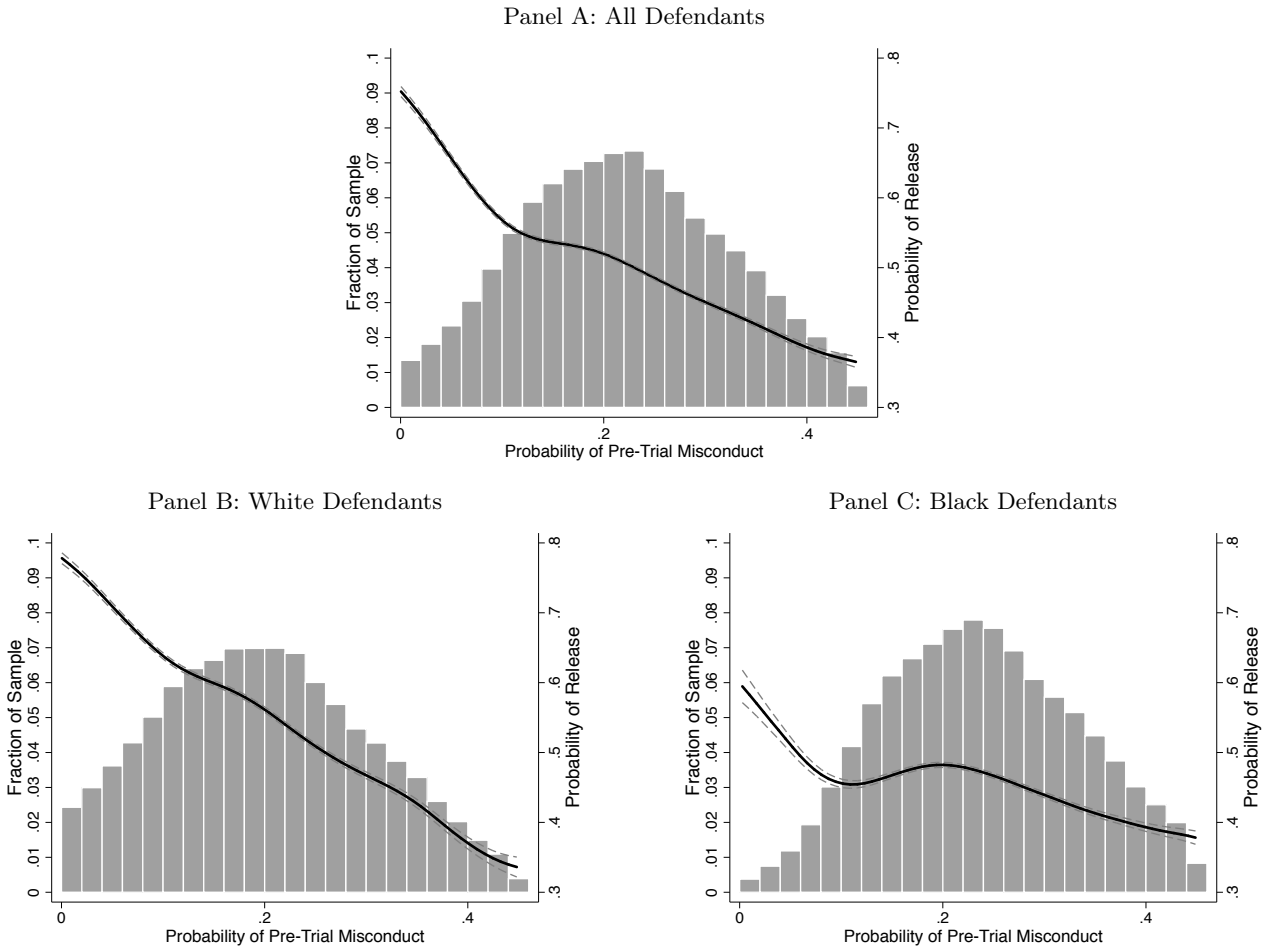
This table reports two-stage least squares results of the estimate of racial bias separately by judge specialization and judge experience. Columns 1-3 report estimates for non-specialist bail judges in Miami-Dade and specialist bail judges in Philadelphia. Columns 4-6 report estimates for non-specialist bail judges in Miami with below and above median years of experience. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Figure 1: Distribution of Judge Leniency and First Stage



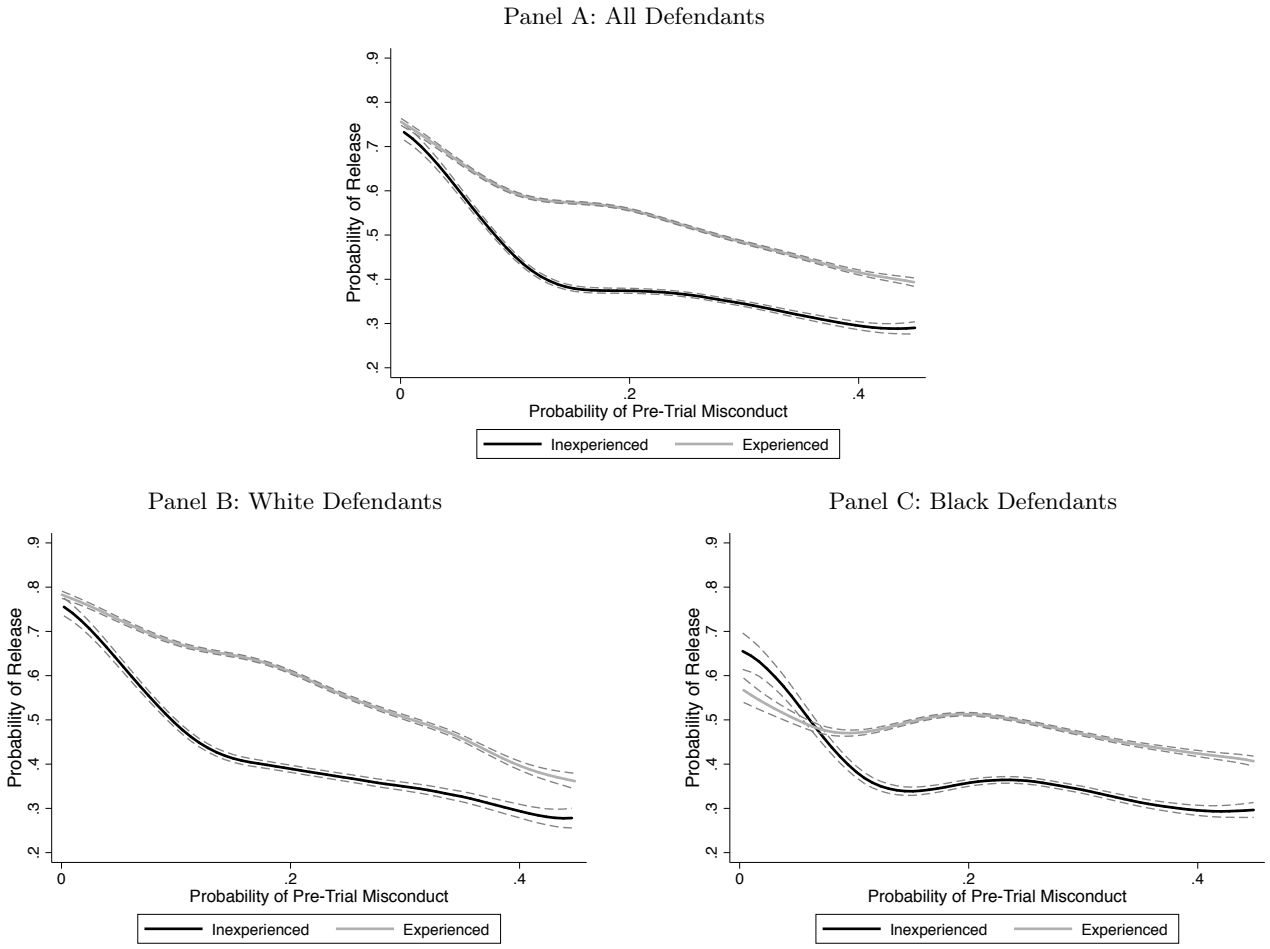
Note: This figure reports the distribution of the judge leniency measure that is estimated using data from other cases assigned to a bail judge in the same year following the procedure described in Section II.B. Panel A pools all defendants. Panel B restricts the sample to white defendants. Panel C restricts the sample to black defendants.

Figure 2: Probability of Pre-trial Misconduct and Detention Status



Note: This figure reports the distribution of the predicted probability of pre-trial misconduct and the relationship between pre-trial release and the predicted probability of pre-trial misconduct, estimated by local linear regression, separately by defendant race. The probability of pre-trial misconduct is computed by estimating a linear regression of the probability of rearrest prior to case disposition conditional on release on the crime and demographic controls discussed in Section II.C.

Figure 3: Probability of Pre-trial Misconduct and Detention Status by Experience



Note: This figure reports the relationship between pre-trial release and the predicted probability of pre-trial misconduct, estimated by local linear regression, separately by defendant race and judge experience as defined in Section IV. The probability of pre-trial misconduct is computed by estimating a linear regression of the probability of rearrest prior to case disposition conditional on release on the crime and demographic controls discussed in Section II.C.

Appendix A: Additional Results

Appendix Table A1: White-Hispanic Results

	White	Hispanic	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)
Rearrest Prior to Disposition	0.150**	0.250**	−0.099
	(0.075)	(0.119)	(0.138)
	[0.196]	[0.191]	–
Panel B: Rearrest by Crime Type			
Drug Crime	0.103**	0.054	0.049
	(0.046)	(0.055)	(0.072)
	[0.073]	[0.083]	–
Property Crime	0.080*	0.001	0.079
	(0.045)	(0.080)	(0.091)
	[0.061]	[0.056]	–
Violent Crime	0.000	0.102**	−0.101*
	(0.033)	(0.049)	(0.058)
	[0.027]	[0.030]	–
Observations	35,468	78,554	–

This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by race. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A2: Characteristics of Compliers by Race

	White			Black		
	P[X=x]	P[X=x|complier]	$\frac{P[X=x\|complier]}{P[X=x]}$	P[X=x]	P[X=x|complier]	$\frac{P[X=x\|complier]}{P[X=x]}$
Drug	0.313	0.388	1.241	0.348	0.335	0.963
	(0.001)	(0.026)	(0.082)	(0.001)	(0.022)	(0.064)
NonDrug	0.687	0.612	0.890	0.652	0.665	1.019
	(0.001)	(0.026)	(0.037)	(0.001)	(0.022)	(0.034)
Violent	0.192	0.006	0.029	0.204	0.016	0.080
	(0.001)	(0.019)	(0.101)	(0.001)	(0.019)	(0.095)
NonViolent	0.808	0.994	1.230	0.796	0.984	1.235
	(0.001)	(0.019)	(0.024)	(0.001)	(0.019)	(0.024)
Felony	0.433	0.249	0.574	0.529	0.375	0.707
	(0.001)	(0.024)	(0.056)	(0.001)	(0.023)	(0.042)
NonFelony	0.567	0.751	1.326	0.471	0.625	1.329
	(0.001)	(0.024)	(0.043)	(0.001)	(0.023)	(0.048)
Prior	0.273	0.333	1.223	0.328	0.384	1.168
	(0.001)	(0.021)	(0.076)	(0.001)	(0.019)	(0.058)
NonPrior	0.727	0.667	0.916	0.672	0.616	0.918
	(0.001)	(0.021)	(0.029)	(0.001)	(0.019)	(0.028)

Note: This table presents the sample distribution, complier distribution, and relative likelihood for different subgroups by race. Bootstrapped standard errors in parentheses are obtained using 500 replications.

Appendix Table A3: First Stage Results by Case Characteristics

	Crime Severity		Crime Type			Defendant Type	
	Misd.	Felony	Property	Drug	Violent	Prior	No Prior
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Pre-trial Release	0.793***	0.383***	0.744***	0.615***	0.068	0.703***	0.533***
	(0.046)	(0.043)	(0.055)	(0.052)	(0.057)	(0.049)	(0.039)
	[0.643]	[0.360]	[0.371]	[0.598]	[0.475]	[0.369]	[0.369]
Court x Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crime Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	148,269	138,734	95,012	71,113	56,791	86,552	200,451

Note: This table reports first stage subsample results. The regressions are estimated on the sample as described in the notes to Table 1. Judge leniency is estimated using data from other cases assigned to a bail judge in the same year following the procedure described in Section II.B. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A4: Non-Hispanic White-Black Results

	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)
Rearrest Prior to Disposition	0.147*	0.006	0.141
	(0.076)	(0.057)	(0.091)
	[0.196]	[0.216]	–
Panel B: Rearrest by Crime Type			
Drug Crime	0.101**	−0.020	0.121**
	(0.047)	(0.037)	(0.057)
	[0.073]	[0.096]	–
Property Crime	0.080*	−0.001	0.080
	(0.045)	(0.033)	(0.053)
	[0.061]	[0.059]	–
Violent Crime	−0.003	−0.038	0.035
	(0.033)	(0.027)	(0.042)
	[0.027]	[0.042]	–
Observations	35,468	149,408	–

This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately for blacks and non-Hispanic whites. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A5: Results for Drug Offenders vs. Other Crimes

	Drug			Other Crime		
	White	Black	Difference	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)	(4)	(5)	(6)
Rearrest Prior to Disposition	0.264**	−0.096	0.360**	0.156*	0.053	0.103
	(0.113)	(0.120)	(0.160)	(0.080)	(0.062)	(0.102)
	[0.221]	[0.268]	–	[0.153]	[0.189]	–
Panel B: Rearrest by Crime Type						
Drug Crime	0.258**	−0.093	0.351**	0.009	0.003	0.007
	(0.112)	(0.120)	(0.160)	(0.011)	(0.009)	(0.015)
	[0.220]	[0.267]	–	[0.006]	[0.005]	–
Property Crime	0.002	−0.002	0.004	0.038	0.008	0.031
	(0.012)	(0.010)	(0.016)	(0.061)	(0.046)	(0.074)
	[0.001]	[0.001]	–	[0.074]	[0.089]	–
Violent Crime	−0.003	−0.001	−0.001	0.064	−0.052	0.116**
	(0.007)	(0.006)	(0.009)	(0.040)	(0.038)	(0.056)
	[0.001]	[0.001]	–	[0.038]	[0.063]	–
Observations	43,057	51,955	–	94,538	97,453	–

Note: This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by crime type. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. Subgroup-specific means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A6: Results for Prior Offenders vs. No Prior Offenders

	Prior Offender			No Prior Offender		
	White	Black	Difference	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)	(4)	(5)	(6)
Rearrest Prior to Disposition	0.285***	−0.027	0.311**	0.118	0.032	0.086
	(0.097)	(0.081)	(0.126)	(0.083)	(0.073)	(0.110)
	[0.256]	[0.285]	–	[0.144]	[0.183]	–
Panel B: Rearrest by Crime Type						
Drug Crime	0.095*	−0.040	0.135*	0.073*	0.007	0.066
	(0.058)	(0.052)	(0.075)	(0.040)	(0.043)	(0.058)
	[0.109]	[0.130]	–	[0.059]	[0.080]	–
Property Crime	0.057	−0.013	0.070	0.005	0.009	−0.004
	(0.055)	(0.046)	(0.067)	(0.059)	(0.042)	(0.071)
	[0.079]	[0.080]	–	[0.041]	[0.048]	–
Violent Crime	0.034	−0.040	0.074	0.042	−0.043	0.084*
	(0.042)	(0.038)	(0.056)	(0.035)	(0.035)	(0.050)
	[0.032]	[0.049]	–	[0.024]	[0.038]	–
Observations	37,506	49,046	–	100,089	100,362	–

Note: This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by whether the defendant is a prior offender. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A7: Results for Felonies vs. Misdemeanors

	Felony			Misdemeanor		
	White	Black	Difference	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)	(4)	(5)	(6)
Rearrest Prior to Disposition	0.594**	−0.095	0.689**	0.034	0.074	−0.040
	(0.246)	(0.129)	(0.272)	(0.045)	(0.046)	(0.065)
	[0.217]	[0.249]	–	[0.141]	[0.180]	–
Panel B: Rearrest by Crime Type						
Drug Crime	0.275**	−0.009	0.284**	0.002	−0.025	0.027
	(0.113)	(0.088)	(0.136)	(0.023)	(0.025)	(0.034)
	[0.095]	[0.112]	–	[0.056]	[0.078]	–
Property Crime	0.134	−0.017	0.151	−0.008	0.016	−0.024
	(0.157)	(0.078)	(0.170)	(0.022)	(0.023)	(0.032)
	[0.086]	[0.085]	–	[0.025]	[0.030]	–
Violent Crime	0.118	−0.078	0.196*	0.017	−0.012	0.029
	(0.088)	(0.065)	(0.108)	(0.019)	(0.018)	(0.026)
	[0.040]	[0.058]	–	[0.016]	[0.023]	–
Observations	59,628	79,106	–	77,967	70,302	–

Note: This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by crime severity. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A8: Results for High-Income vs. Low-Income Offenders

	Above Median Income			Below Median Income		
	White	Black	Difference	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)	(4)	(5)	(6)
Rearrest Prior to Disposition	−0.070	0.017	−0.086	0.257**	0.011	0.246*
	(0.091)	(0.115)	(0.145)	(0.103)	(0.074)	(0.130)
	[0.154]	[0.211]	–	[0.181]	[0.213]	–
Panel B: Rearrest by Crime Type						
Drug Crime	0.055	−0.062	0.116	0.082	−0.002	0.084
	(0.048)	(0.064)	(0.079)	(0.050)	(0.045)	(0.068)
	[0.058]	[0.085]	–	[0.082]	[0.097]	–
Property Crime	−0.112*	0.043	−0.155*	0.093	−0.032	0.125
	(0.064)	(0.068)	(0.091)	(0.068)	(0.043)	(0.080)
	[0.047]	[0.066]	–	[0.049]	[0.056]	–
Violent Crime	−0.014	−0.047	0.033	0.009	−0.026	0.036
	(0.040)	(0.049)	(0.060)	(0.043)	(0.037)	(0.057)
	[0.023]	[0.040]	–	[0.027]	[0.041]	–
Observations	33,766	17,526	–	85,960	110,992	–

Note: This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately for high-income and low-income defendants. A defendant is classified as high-income if the zip code of has an average level of income greater than the median level of income in the city, while low-income is defined as coming from a zip code with an average level of income less than the median level of income in the city. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A9: Results Weighting by Case and Defendant Characteristics

	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)
Rearrest Prior to Disposition	0.191***	0.032	0.159*
	(0.066)	(0.054)	(0.085)
	[0.174]	[0.216]	–
Panel B: Rearrest by Crime Type			
Drug Crime	0.082**	−0.019	0.101**
	(0.033)	(0.032)	(0.045)
	[0.073]	[0.096]	–
Property Crime	0.028	0.008	0.020
	(0.044)	(0.033)	(0.054)
	[0.051]	[0.059]	–
Violent Crime	0.045*	−0.021	0.066*
	(0.027)	(0.024)	(0.036)
	[0.026]	[0.042]	–
Observations	137,595	149,408	–

Note: This table reports weighted two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by race. Results are re-weighted with the weights chosen to match the distribution of observable characteristics by race. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A10: Robustness to Clustering at the Judge Level

	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)
Rearrest Prior to Disposition	0.185**	0.005	0.180*
	(0.082)	(0.059)	(0.098)
	[0.174]	[0.216]	–
Panel B: Rearrest by Crime Type			
Drug Crime	0.077*	−0.020	0.097*
	(0.045)	(0.036)	(0.051)
	[0.073]	[0.096]	–
Property Crime	0.029	−0.001	0.030
	(0.043)	(0.031)	(0.055)
	[0.051]	[0.059]	–
Violent Crime	0.044	−0.038	0.082**
	(0.031)	(0.026)	(0.038)
	[0.026]	[0.042]	–
Observations	137,595	149,408	–

This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by race. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A11: Robustness to Race-Specific Leniency Measures

	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)
Rearrest Prior to Disposition	0.181*	−0.032	0.213*
	(0.105)	(0.068)	(0.123)
	[0.196]	[0.216]	–
Panel B: Rearrest by Crime Type			
Drug Crime	0.129*	−0.050	0.179**
	(0.067)	(0.044)	(0.079)
	[0.073]	[0.096]	–
Property Crime	0.097	−0.006	0.103
	(0.062)	(0.037)	(0.070)
	[0.061]	[0.059]	–
Violent Crime	−0.038	−0.039	0.001
	(0.046)	(0.031)	(0.057)
	[0.027]	[0.042]	–
Observations	35,468	149,408	–

This table reports two-stage least squares results of the impact of pre-trial release on the probability of pre-trial misconduct separately by race with judge leniency computed separately by race. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A12: OLS Results

	White	Black	Difference
Panel A: Rearrest for All Crimes	(1)	(2)	(3)
Rearrest Prior to Disposition	0.042***	0.037***	0.005
	(0.003)	(0.002)	(0.003)
	[0.174]	[0.216]	–
Panel B: Rearrest by Crime Type			
Drug Crime	0.023***	0.029***	−0.006***
	(0.001)	(0.002)	(0.002)
	[0.073]	[0.096]	–
Property Crime	0.003*	0.002	0.001
	(0.001)	(0.001)	(0.002)
	[0.051]	[0.059]	–
Violent Crime	−0.008***	−0.009***	0.002
	(0.001)	(0.001)	(0.001)
	[0.026]	[0.042]	–
Observations	137,595	149,408	–

Note: This table replicates the Knowles et al. (2001) test. The table reports OLS results of the impact of pre-trial release on the probability of pre-trial misconduct separately by race. The regressions are estimated on the sample as described in the notes to Table 1. The dependent variable is listed in each row. Robust standard errors two-way clustered at the individual and judge-by-shift level are reported in parentheses. The sample means of the dependent variables are reported in brackets. All specifications control for court-by-time fixed effects as well as the demographic and crime controls discussed in Section II.C. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Appendix Table A13: Pre-Trial Release and Pre-Trial Misconduct by Judge and Defendant Race

	Race of Judge	
	White	Black
Panel A: Release Rates	(1)	(2)
White	0.345	0.339
	(0.475)	(0.474)
Black	0.311	0.318
	(0.463)	(0.466)
Panel B: Pre-Trial Rearrest Rates		
White	0.175	0.174
	(0.380)	(0.379)
Black	0.253	0.273
	(0.435)	(0.446)

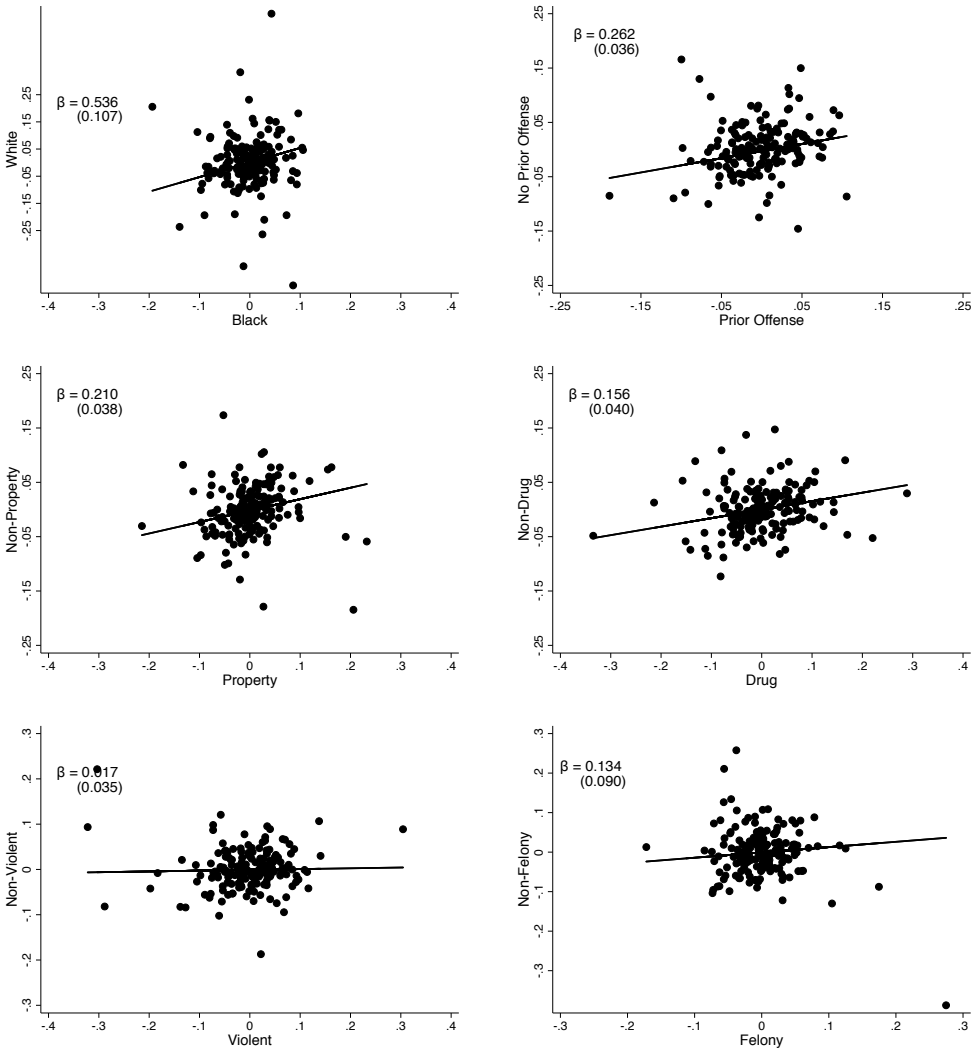
Note: This table presents average rates of pre-trial release and pre-trial misconduct conditional on release by defendant and judge race in Miami. The means are calculated using the Miami sample reported in Table 1. See text for additional details.

Appendix Table A14: p-values from Tests of Relative Racial Prejudice

	p-value
	(1)
Pre-Trial Release	0.364
Pre-Trial Rearrest	0.412

Note: This table replicates the Anwar and Fang (2006) test for pre-trial release rates and pre-trial misconduct rates. This table presents bootstrapped p-values testing for relative racial bias. The null hypothesis is rejected if white judges are more lenient on white defendants, and black judges are more lenient on black defendants.

Appendix Figure A1: Judge Leniency by Defendant and Case Characteristics



Note: These figures show the correlation between our residualized measure of judge leniency for different groups of defendants. We take the average leniency for each group over all available years of data. The solid line shows the best linear fit estimated using OLS relating each judge leniency measure.

Appendix B: Proofs of Propositions

A. Proof of Proposition 2

In this proposition, we assume there is no racial bias and bound the maximum size of bias (difference in the treatment effect for marginal white and marginal black defendants) that could theoretically be estimated by our IV estimator. The direction of the bias is ambiguous and so for illustrative purposes we bound the maximum bias against black defendants. The proof is exactly analogous for bounding the maximum bias against white defendants.

As noted in the text, Imbens and Angrist (1994) show that the IV estimator converges to:

$$\alpha_r = \sum_{j=1}^{J} \lambda_r^j \alpha_r^{j,j-1} \tag{B.1}$$

where $\alpha_r^{j,j-1} = \mathbb{E}[Y_i(1) - Y_i(0) | R_i(z_j) - R_i(z_{j-1}) = 1, r_i = r]$ and the weights, λ_r^j, are non-negative and sum up to one. Our estimator for racial bias is given by:

$$D = \alpha_W - \alpha_B \tag{B.2}$$

If $D > 0$, then we conclude there is racial bias against black defendants in the judicial system.

In general, our empirical test may identify racial bias against blacks for two reasons: (1) if the weights, λ_r^j, vary by defendant race and (2) if treatment effects among the complier population, $\alpha_r^{j,j-1}$, vary by defendant race. Both of these reasons could arise due to differences in the underlying distribution of treatment effects by defendant race.

Specifically, Imbens and Angrist (1994) show that the weights in the IV estimator are given by the formula:

$$\lambda_r^j = \frac{(P(Released = 1|z_j, r) - P(Released = 1|z_{j-1}, r)) \cdot \sum_{l=j}^{J} \pi_r^l (g(z_l) - \mathbb{E}[g(Z)])}{\sum_{m=1}^{J} (P(Released = 1|z_m, r) - P(Released = 1|z_{m-1}, r)) \cdot \sum_{l=m}^{J} \pi_r^l (g(z_l) - \mathbb{E}[g(Z)])} \tag{B.3}$$

where $g(Z)$ is the instrumental variable and π_r^j is the probability of being assigned judge j for defendant race r. In our setting, we use our measure of judge leniency as the instrument, such that $g(z_l) = z_l$. Given the linear specification for the first stage equation (Assumption 4), the following statement can be re-expressed as:

$$P(Released = 1|z_j, r) - P(Released = 1|z_{j-1}, r) = \gamma_r (z_j - z_{j-1})$$

Plugging this statement back into the formula for our IV estimator and simplifying yields:

$$\lambda_r^j = \frac{z_j - z_{j-1} \cdot \sum_{l=j}^{J} \pi_r^l (z_l - \mathbb{E}[Z])}{\sum_{m=1}^{J} (z_m - z_{m-1}) \cdot \sum_{l=m}^{J} \pi_r^l (z_l - \mathbb{E}[Z])} \tag{B.4}$$

Under this expression, it is straightforward to see that if the probability of being assigned to

judge j does not differ by race, such that $\pi_B^j = \pi_W^j$ for all j, these weights λ_r^j will be equal for black defendants and white defendants. Under the exclusion restriction (Assumption 2), the probability of being assigned to any particular judge should not differ by defendant race (consistent with the patterns observed in Figure 1). Therefore, for illustrative purposes, we drop the r subscript on the weights λ^j.

Recall that in our setting, $\alpha_r^{j,j-1}$ is the average treatment effect for all individuals who are released if assigned to judge j but detained if assigned to judge $j-1$, in other words compliers. Specifically, compliers in this scenario are defendants of race r for whom $t_r^{j-1}(\mathbf{V}_i) < \mathbb{E}[\alpha_i] \leq t_r^j(\mathbf{V}_i)$. Let the smallest treatment effect in this interval be given by $\alpha_{min}^{j,j-1}$ and the largest treatment effect be given by $\alpha_{max}^{j,j-1}$. Given that judges make release decision based on the expected treatment effect, the largest treatment effect in the interval is the treatment effect of the individual marginal to the more lenient judge j, while the smallest treatment effect is the treatment effect of the first individual detained by the stricter judge $j-1$. Under the null of no racial bias, our empirical test will be most biased towards finding racial bias against black defendants if we assume that:

$$\begin{aligned} \alpha_W^{j,j-1} &= \alpha_{max}^{j,j-1} \\ \alpha_B^{j,j-1} &= \alpha_{min}^{j,j-1} \end{aligned} \tag{B.5}$$

In other words, for every interval between judge j and $j-1$, our estimator is most biased if we assume that the treatment effect for all white defendants is equal to the maximum treatment effect in the interval, and we assume that the treatment effect for all black defendants is equal to the minimum treatment effect in the interval. This assumption implies:

$$D \leq \sum_{j=1}^{J} \lambda^j \cdot \alpha_{max}^{j,j-1} - \sum_{j=1}^{J} \lambda^j \cdot \alpha_{min}^{j,j-1} \tag{B.6}$$

Also, note that:

$$\alpha_{max}^{j,j-1} < \alpha_{min}^{j+1,j} \tag{B.7}$$

Given that $\alpha_{max}^{j,j-1}$ is the treatment effect for the individual marginal to judge j, any defendant released by judge $j+1$, but detained by judge j must have a larger treatment effect than the defendant marginal to judge j. Therefore, replacing $\alpha_{max}^{j,j-1}$ with $\alpha_{min}^{j+1,j}$ in equation (B.6) allows us to bound the level of racial bias by:

$$D < \sum_{j=1}^{J} \lambda^j \cdot \alpha_{min}^{j+1,j} - \sum_{j=1}^{J} \lambda^j \cdot \alpha_{min}^{j,j-1} \tag{B.8}$$

Relabeling $\alpha_{min}^{j,j-1} = \alpha^{j-1}$ for simplicity and rearranging yields:

$$D < \sum_{j=1}^{J} \lambda^j \cdot (\alpha^j - \alpha^{j-1}) \tag{B.9}$$

Now consider the first two elements of the above sum:

$$\lambda^1(\alpha^1 - \alpha^0) + \lambda^2(\alpha^2 - \alpha^1)$$

Note that in our model $\alpha^{j+1} - \alpha^j > 0$ for all j (the more lenient the judge the higher the threshold for release and therefore the larger the treatment effect). Without loss of generality, assume $\lambda^1 \geq \lambda^2$. Then:

$$\lambda^1(\alpha^1 - \alpha^0) + \lambda^2(\alpha^2 - \alpha^1) \leq \lambda^1(\alpha^2 - \alpha^0)$$

Without loss of generality, assume $\lambda^1 \geq \lambda^j$ for all j. Noting that $\alpha^0 = \alpha_{min}$ (the smallest treatment effect is associated with the most strict judge) and $\alpha^J = \alpha_{max}$ (the largest treatment effect is associated with the most lenient judge), then:

$$D < \lambda^1(\alpha_{max} - \alpha_{min}) \tag{B.10}$$

Therefore, under the null of no racial bias, the maximum potential bias against blacks that our estimator could generate is bounded by $\max_j(\lambda^j)(\alpha_{max} - \alpha_{min})$. □

Estimating Maximum Bias under Null of No Racial Bias: We now illustrate how we empirically estimate the maximum potential bias of our IV estimator under the null of no racial bias by using the formula in Proposition 2. Because we cannot observe $\alpha_{max} - \alpha_{min}$, we take the most conservative approach and assume that this value is equal to 1. In other words, we assume that there are defendants who are rearrested with probability 1 if released but never rearrested if detained such that $\alpha_{max} = 1$, and also that there are defendants whose rearrest probability is unaffected by release status such that $\alpha_{min} = 0$.

We then calculate the weights using the following formula:

$$\lambda^j = \frac{z_j - z_{j-1} \cdot \sum_{l=j}^{J} \pi^l (z_l - \mathbb{E}[Z])}{\sum_{m=1}^{J} (z_m - z_{m-1}) \cdot \sum_{l=m}^{J} \pi^l (z_l - \mathbb{E}[Z])} \tag{B.11}$$

We replace π^j and $\mathbb{E}[Z]$ with their empirical counterparts:

$$\hat{\pi}^j = \sum_{i=1}^{N} \frac{\mathbb{1}\{Z_i = z_j\}}{N} \tag{B.12}$$

$$\mathbb{E}[Z] = \frac{1}{N} \sum_{i=1}^{N} Z_i \tag{B.13}$$

Plugging these quantities into the formula for the weights yields an estimate of the weight attached to each pairwise LATE. We then take the maximum of our weights, multiply it by $\alpha_{max} - \alpha_{min}$ which we conservatively assume to be 1, and interpret this result as the maximum bias that our IV estimator could yield under the null of no racial bias.

This procedure yields a maximum bias of 0.005. Therefore, under the null of no racial bias, the estimated difference between the causal effect of pre-trial release on rearrest between marginal white and marginal black defendants could be no larger than 0.5 percentage points. We also take a more conservative approach and split our instrument into 100 quantiles, greatly reducing the number of distinct values the instrument takes on. With this alternative specification, the maximum bias our estimator could yield is approximately 1.5 percentage points.

B. Proof of Proposition 3

Let the weights for all white defendants be equal to 1. We construct the weights for a black defendant with observables equal to $\mathbf{X}_i$ as:

$$\Psi(\mathbf{X}_i) = \frac{Pr(W|\mathbf{X}_i)Pr(B)}{Pr(B|\mathbf{X}_i)Pr(W)} \tag{B.14}$$

Where $Pr(W|\mathbf{X}_i)$ is the probability of being white given observables $\mathbf{X}_i$, $Pr(B|\mathbf{X}_i)$ is the probability of being black given observables $\mathbf{X}_i$, and $Pr(B)$ is the unconditional probability of being black. Next, we show that under the null of no racial bias, the instrumental variables estimates of the effect of pre-trial release on misconduct will be equal between races.

For illustrative purposes, assume that the pairwise LATEs are equal by race and so we drop the race subscripts on $\alpha_r^{j,j-1}$ and λ_r^j. For the re-weighted sample, our IV estimate for black defendants will converge to:

$$\alpha_B^{\omega} = \sum_{x \in X} \sum_{j=1}^{J} \lambda^j(x)\alpha^{j,j-1}(x)Pr(\mathbf{X}_i = x|B)\Psi(x) \tag{B.15}$$

The pairwise LATEs now depend on x due to the threshold of release varying by other observables. As discussed in the text, this may cause the estimated treatment effects to differ by race due to the composition of crimes, rather than direct racial bias.

Note that Bayes rule implies:

$$Pr(\mathbf{X}_i = x|B) = \frac{Pr(B|\mathbf{X}_i = x)Pr(\mathbf{X}_i = x)}{Pr(B)} \tag{B.16}$$

Expanding the weights and simplifying yields:

$$\alpha_B^{\omega} = \sum_{x \in X} \sum_{j=1}^{J} \lambda^j(x)\alpha^{j,j-1}(x)Pr(\mathbf{X}_i = x|W) = \alpha_W^{\omega} \tag{B.17}$$

given that the weights for all white defendants are equal to 1. □

Appendix C: Alternative Model of Bail Setting

Model Setup: In this section, we present a model in which judges make racially biased prediction errors. We show that a model motivated by racially-biased prediction errors can generate the same predictions as a model of taste-based discrimination.

As in the main text, let i denote defendants and $\mathbf{V}_i$ denote all case and defendant characteristics considered by the bail judge, excluding defendant race r_i. The benefit of releasing defendant i assigned to judge j is $t(\mathbf{V}_i)$, which does not vary by judge.

The expected cost of release for defendant i conditional on observable characteristics $\mathbf{V}_i$ is equal to the expected probability of pre-trial misconduct, $\mathbb{E}^j[\alpha_i|\mathbf{V}_i, r_i]$, which varies across judge. We can write the expected cost of release as:

$$\mathbb{E}^j[\alpha_i|\mathbf{V}_i] = \mathbb{E}[\alpha_i|\mathbf{V}_i, r_i = r] + \tau_r^j(\mathbf{V}_i) \tag{C.1}$$

where τ_r^j is a prediction error that is allowed to vary by race and judge. To simplify our notation, we let the true probability of pre-trial misconduct conditional on all variables observed by the judge be denoted by $\mathbb{E}[\alpha_i|r_i]$.

Definition C.1. We define judge j as racially biased against black defendants if $\tau_B^j(\mathbf{V}_i) > \tau_W^j(\mathbf{V}_i)$. Thus, racially biased judges systematically over-estimate the cost of release for black defendants relative to white defendants.

Finally, we assume that bail judges are risk neutral and maximize the net benefit of pre-trial release. Thus, bail judge j will release defendant i if and only if the benefit of pre-trial release is greater than the expected cost of release:

$$\mathbb{E}^j[\alpha_i|\mathbf{V}_i, r_i] = \mathbb{E}[\alpha_i|r_i] + \tau_r^j(\mathbf{V}_i) \leq t(\mathbf{V}_i) \tag{C.2}$$

If we relabel $t(\mathbf{V}_i) - \tau_r^j(\mathbf{V}_i) = t_r^j(\mathbf{V}_i)$, then this model reduces to the formulation in the main text. With this relabeling, the model motivated by racially-biased prediction errors outlined here can generate the same predictions as a model of taste-based discrimination outlined in the main text.

Interpretation of IV Estimates: While the same theoretical predictions can be generated by a model based on racially-biased prediction errors and a model based on taste-based discrimination, the interpretation of the judge IV estimates does differ depending on the model. In a model based on racially-biased prediction errors, a judge IV estimator identifies the impact of pre-trial detention due to judge-specific differences in prediction errors, not judge-specific differences in preferences for release.

To see this, it is helpful to consider the two-judge case. Let bail judge 1's and bail judge 2's prediction errors be given by $\tau_r^1(\mathbf{V}_i)$ and $\tau_r^2(\mathbf{V}_i)$ with $\tau_r^1(\mathbf{V}_i) > \tau_r^2(\mathbf{V}_i)$. In this framework, judge 1 overestimates the probability of pre-trial misconduct relative to judge 2, and judge 1 will therefore be less lenient than judge 2. Let z_1 and z_2 measure the judge-specific differences in predictions

across judge 1 and judge 2. We can express the resulting instrumental variables estimates of α_W and α_B as follows:

$$\begin{aligned} \alpha_W &= \mathbb{E}\big[\alpha_i | R_i(z_2) - R_i(z_1) = 1, r_i = W] \\ \alpha_B &= \mathbb{E}\big[\alpha_i | R_i(z_2) - R_i(z_1) = 1, r_i = B] \end{aligned} \tag{C.3}$$

In this scenario, compliers are defendants of race r for whom $t(\mathbf{V}_i) - \tau_r^1(\mathbf{V}_i) < \mathbb{E}[\alpha_i | r_i = r] \leq t(\mathbf{V}_i) - \tau_r^2(\mathbf{V}_i)$. This is due to the fact that judge j releases defendant i of race r if and only if $\mathbb{E}[\alpha_i | r_i = r] \leq t(\mathbf{V}_i) - \tau_r^j(\mathbf{V}_i)$.

Finally, let $\tau_r^1 \to \tau_r^2$, such that the marginal defendant $\alpha_r \to t(\mathbf{V}_i) - \tau_r$. Then, our estimator for discrimination $D = \alpha_W - \alpha_B$ converges to:

$$D = t(\mathbf{V}_i) - \tau_W - (t(\mathbf{V}_i) - \tau_B) = \tau_B - \tau_W \tag{C.4}$$

where if $D > 0$, that indicates that judges systematically overestimate the cost of release for marginal black defendants relative to marginal white defendants. In contrast, in a model based on taste-based discrimination, an estimate of $D > 0$ indicates that judges value the freedom of marginal black defendants less than the freedom of marginal white defendants.

Appendix D: Data Appendix

Judge Leniency: We calculate judge leniency as the leave-one-out mean residualized pre-trial release decisions of the assigned judge within a bail year. We use the residual pre-trial release decision after removing court-by-time fixed effects. In our main results, we define pre-trial release based on whether a defendant was released within the first three days after the bail hearing.

Release on Recognizance: An indicator for whether the defendant was released on recognizance (ROR), where the defendant secures release on the promise to return to court for his next scheduled hearing. ROR is used for defendants who show minimal risk of flight, no history of failure to appear for court proceedings, and pose no apparent threat of harm to the public.

Non-Monetary Bail: An indicator for whether the defendant was released on non-monetary bail, also known as conditional release. Non-monetary conditions include monitoring, supervision, halfway houses, and treatments of various sorts, among other options.

Monetary Bail: An indicator for whether the defendant was assigned monetary bail. Under monetary bail, a defendant is generally required to post a bail payment to secure release, typically 10 percent of the bail amount, which can be posted directly by the defendant or by sureties such as bail bondsman.

Bail Amount: Assigned monetary bail amount in thousands, set equal to missing for defendants who receive non-monetary bail or ROR.

Race: Information on defendant race is missing for the Philadelphia data prior to 2010.

Hispanic: We match the surnames in our data to census genealogical records of surnames. If the probability a given surname is Hispanic is greater than 80 percent, we label the defendant as Hispanic.

Prior Offense in Past Year: An indicator for whether the defendant had been charged for a prior offense in the past year of the bail hearing within the same county, set to missing for defendants who we cannot observe for a full year prior to their bail hearing.

Number of Offenses: Total number of charged offenses.

Felony Offense: An indicator for whether the defendant is charged with a felony offense.

Misdemeanor Offense: An indicator for whether the defendant is charged with only misdemeanor offenses.

Rearrest: An indicator for whether the defendant was rearrested for a new crime prior to case disposition.

Race: We collect information on judge race from court directories and conversations with court officials. All judges in Philadelphia are white. Information on judge race in Miami is missing for two of the 170 judges in our sample.

Experience: We use historical court records back to 1999 to compute experience, which we define as the difference between bail year and start year (earliest 1999). In our sample, years of experience range from zero to 15 years.

Appendix E: Institutional Details Appendix

Philadelphia County: Immediately following arrest in Philadelphia County, defendants are brought to one of six police stations around the city where they are interviewed by the city's Pre-Trial Services Bail Unit. The Bail Unit operates 24 hours a day, seven days a week, and interviews all adults charged with offenses in Philadelphia through videoconference, collecting information on the arrested individual's charge severity, personal and financial history, family or community ties, and criminal history. The Bail Unit then uses this information to calculate a release recommendation based on a four-by-ten grid of bail guidelines that is presented to the bail judge. However, these bail guidelines are only followed by the bail judge about half the time, with judges often imposing monetary bail instead of the recommended non-monetary options (Shubik-Richards and Stemen 2010).

After the Pre-Trial Services interview is completed and the charges are approved by the Philadelphia District Attorney's Office, the defendant is brought in for a bail hearing. Since the mid-1990s, bail hearings have been conducted through videoconference by the bail judge on duty, with representatives from the district attorney and local public defender's offices (or private defense counsel) also present. However, while a defense lawyer is present at the bail hearing, there is no real opportunity for defendants to speak with the attorney prior to the hearing. At the hearing itself, the bail judge reads the charges against the defendant, informs the defendant of his right to counsel, sets bail after hearing from representatives from the prosecutor's office and the defendant's counsel, and schedules the next court date. After the bail hearing, the defendant has an opportunity to post bail, secure counsel, and notify others of the arrest. If the defendant is unable to post bail, he is detained but has the opportunity to petition for bail modification in subsequent court proceedings.

Miami-Dade County: The Miami-Dade bail system follows a similar procedure, with one important exception. As opposed to Philadelphia where all defendants are required to have a bail hearing, most defendants in Miami-Dade can avoid a bail hearing and be immediately released following arrest and booking by posting an amount designated by a standard bail schedule. The bail schedule ranks offenses according to their seriousness and assigns an amount of bond that must be posted to permit a defendant's release. Critics have argued that this kind of standardized bail schedule discriminates against poor defendants by setting a fixed price for release according to the charged offense rather than taking into account a defendant's ability to pay, or propensity to flee or commit a new crime. Approximately 30 percent of all defendants in Miami-Dade are released prior to a bail hearing, with the other 70 percent attending a bail hearing (Goldkamp and Gottfredson 1988).

If a defendant is unable to post bail immediately in Miami-Dade, there is a bail hearing within 24 hours of arrest where defendants can argue for a reduced bail amount. Miami-Dade conducts separate daily hearings for felony and misdemeanor cases through videoconference by the bail judge on duty. At the bail hearing, the court will determine whether or not there is sufficient probable cause to detain the arrestee and if so, the appropriate bail conditions. The bail amount may be lowered, raised, or remain the same as the scheduled bail amount depending on the case situation

and the arguments made by defense counsel and the prosecutor. While monetary bail amounts at this stage often follow the standard bail schedule, the choice between monetary versus non-monetary bail conditions varies widely across judges in Miami-Dade (Goldkamp and Gottfredson 1988).

Institutional Features Relevant to Empirical Design: Our empirical strategy exploits variation in the pre-trial release tendencies of the assigned bail judge. There are three features of the Philadelphia and Miami-Dade bail systems that make them an appropriate setting for our research design. First, there are multiple bail judges serving simultaneously, allowing us to measure variation in bail decisions across judges. At any point in time, Philadelphia has six bail judges that only make bail decisions. In Miami-Dade, weekday cases are handled by a single bail judge, but weekend cases are handled by approximately 60 different judges on a rotating basis. These weekend bail judges are trial court judges from the misdemeanor and felony courts in Miami-Dade that assist the bail court with weekend cases.

Second, the assignment of judges is based on rotation systems, providing quasi-random variation in which bail judge a defendant is assigned to. In Philadelphia, the six bail judges serve rotating eight-hour shifts in order to balance caseloads. Three judges serve together every five days, with one bail judge serving the morning shift (7:30AM-3:30PM), another serving the afternoon shift (3:30PM-11:30PM), and the final judge serving the night shift (11:30PM-7:30AM). While it may be endogenous whether a defendant is arrested in the morning or at night or on a specific day of the week, the fact that these six bail judges rotate through all shifts and all days of the week allows us to isolate the independent effect of the judge from day-of-week and time-of-day effects. In Miami-Dade, the weekend bail judges rotate through the felony and misdemeanor bail hearings each weekend to ensure balanced caseloads during the year. Every Saturday and Sunday beginning at 9:00AM, one judge works the misdemeanor shift and another judge works the felony shift. Because of the large number of judges in Miami-Dade, any given judge works a bail shift approximately once or twice a year.

Third, there is very limited scope for influencing which bail judge will hear the case, as most individuals are brought for a bail hearing shortly following the arrest. In Philadelphia, all adults arrested and charged with a felony or misdemeanor appear before a bail judge for a formal bail hearing, which is usually scheduled within 24 hours of arrest. A defendant is automatically assigned to the bail judge on duty. There is also limited room for influencing which bail judge will hear the case in Miami-Dade, as arrested felony and misdemeanor defendants are brought in for their hearing within 24 hours following arrest to the bail judge on duty. However, given that defendants can post bail immediately following arrest in Miami-Dade without having a bail hearing, there is the possibility that defendants may selectively post bail depending on the identity of the assigned bail judge. It is also theoretically possible that a defendant may self-surrender to the police in order to strategically time their bail hearing to a particular bail judge. As a partial check on this important assumption of random assignment, we test the relationship between observable characteristics and bail judge assignment.