

Understanding Race Disparities in Criminal Court Outcomes



SHAWN BUSHWAY, ANDREW JORDAN, DEREK NEAL, AND STEVEN RAPHAEL

We construct a framework that defines optimal outcomes in criminal courts, and we use this framework to interpret and organize the existing literature on racial disparities in pretrial detention, sentencing, and community corrections outcomes. Existing research indicates that some actors within courts and within the agencies that implement the sentences that courts impose make decisions that are contaminated by racial animus or racially biased assessments of the recidivism risks posed by some offenders. However, the most important sources of racial disparities in case outcomes are numerous practices, regulations, and laws that are too punitive—that is, their social costs are likely greater than any derived social benefits. Since minorities, especially Blacks, face arrest at much higher rates than Whites, they bear large disparate impacts from such policies.

Keywords: courts, racial disparity, sentencing goals, disparate impact, structural racism, crime control, social costs and benefits

Racial disparities in many criminal justice outcomes are prevalent and pronounced throughout the United States. Black people are more likely than Whites to be stopped by the police (NASEM 2018), to be treated harshly during police stops (Lofstrom et al. 2021), to be arrested (Snyder 2011), to be incarcerated at any given point in time (Carson and Sabol 2015), and to

be under some form of community supervision (Kaeble 2023). Similar results hold for Native Americans, and to a lesser extent for Hispanics (NASEM 2022). In addition, a recent National Academies of Sciences consensus panel documents racial disparities in the most serious forms of criminal victimization as well as disparities in the likelihoods that victims see

Shawn Bushway is professor in the Department of Public Administration and Policy and the School of Criminal Justice at SUNY Albany, United States. **Andrew Jordan** is assistant professor of economics at Washington University in St. Louis, United States. **Derek Neal** is William C. Norby Professor of Economics at University of Chicago, United States. **Steven Raphael** is the James D. Marver Professor of Public Policy at University of California, Berkeley, United States.

© 2025 Russell Sage Foundation. Bushway, Shawn, Andrew Jordan, Derek Neal, and Steven Raphael. 2025. "Understanding Race Disparities in Criminal Court Outcomes." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 11(3): 86–135. <https://doi.org/10.7758/RSF.2025.11.3.03>. We thank Daniel Brudney, Magne Mogstad, Jack Mountjoy, Isaac Neal, and Evan Rose for useful discussions. We thank Jennifer Dohm and Jordan Boulger for their insights concerning the coding of court dispositions in Cook County, Illinois. We appreciate input from other contributors to the Russell Sage special issue "Racial Bias in Law Enforcement, Criminal Justice, and Incarceration." We also appreciate comments from participants in the 2024 Summer Research Workshop of the Institute for Research on Poverty. Direct correspondence to: Derek Neal, at n9na@uchicago.edu, 1126 East 59th Street, Chicago, IL 60637.

Open Access Policy: *RSF: The Russell Sage Foundation Journal of the Social Sciences* is an open access journal. This article is published under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License.

crimes against them cleared by an arrest (NASEM 2022).

These disparities likely reflect the combined effects of many factors. To begin, persons of different races may experience different levels of involvement in criminal activity. Further, among persons engaged in comparable levels of criminal activity, members of different groups may have heterogeneous interactions with police in terms of relative frequency, the nature of interactions, or some combination of the two. Police may be more likely to search or detain persons that belong to certain racial groups, and holding the quality of evidence constant, police may be more likely to file charges against citizens from particular groups.

We analyze racial disparities in case outcomes among the sample of persons who already face criminal charges, and many factors may contribute to these disparities. To begin, distributions of case and defendant characteristics may differ by race. Further, courts may produce different distributions of outcomes for defendants of different races even among sets of defendants who face similar charges and are observationally similar on numerous dimensions. These differential outcomes may reflect racial animus among court actors or various forms of statistical discrimination. Finally, even in settings where courts produce comparable distributions of case outcomes for comparable defendants regardless of the race of defendants, the structure of sentencing rules or guidelines may generate disparate impacts on some racial groups that cannot be justified by reasonable concerns for public safety.¹

Existing work on criminal justice outcomes contains at least two different definitions of racial disparity. As an example, if the incarceration rate among Black males is fifteen hundred per one hundred thousand persons and the incarceration rate among non-Black males is three hundred per one hundred thousand persons, the relative racial disparity in incarceration rates is 5 to 1, but the absolute disparity is twelve hundred persons per one hundred thousand. Suppose that policymakers adopt a sen-

tencing reform plan that improves public safety while reducing both incarceration rates by one half. The relative disparity in incarceration rates would remain 5 to 1 (750 to 150), but the absolute disparity would fall by one half to six hundred.

Analysts who focus on relative disparities may claim that our hypothetical reform did not reduce racial disparities in incarceration since, before and after the reform, incarceration rates among Black males are five times higher than rates among non-Black males. We disagree. Movements in absolute disparity measures capture changes in the impacts of the criminal justice system on different communities. The fifteen hundred and three hundred per one hundred thousand figures in our example are reasonably close to the actual adult imprisonment rates of Blacks versus Whites respectively in 2019.² If the actual rates had been fifteen and three per one hundred thousand, the relative disparity would have still been 5 to 1, but we doubt that we would be writing a paper for a special issue on racial bias in the criminal justice system.

In the next three sections, we provide a framework for thinking about different sources of intergroup disparities in criminal case dispositions. We begin by defining key concepts. We then flesh out the utilitarian framework that we adopt and discuss the social welfare function at the center of the problem that courts solve. Given this objective function, we define several types of discrimination as well as structural racism, and we then discuss the information and empirical methods researchers need to identify discriminatory outcomes. Armed with these tools, we devote several sections to reviews of the research on racial disparities in criminal case processing, pretrial detention decisions, case dispositions, and post-sentence supervision outcomes among those on probation and parole. We conclude by presenting simulation results. We examine the potential impacts of reforms to pretrial detention, sentencing, and community supervision that reduce punitive treatment for offenders in

1. Policies that are too lenient may also create systemic inequalities since victimization rates have historically been higher in minority communities.

2. See table 6 in Carson (2020). The rates are 1,446 and 263 respectively.

ways that are not likely to create significant harms to public safety. The results suggest that such reforms may be effective tools for reducing racial disparities in overall incarceration rates.

DEFINING BIAS

We begin by defining racial bias in the context of court proceedings. The outcome of a criminal case is racially biased if the outcome is not socially optimal and also contributes to racial disparities in case outcomes. We define the socially optimal outcome for a given case as the outcome that equates the marginal costs and marginal benefits of punishment, and we assume that society values these costs and benefits the same way regardless of the racial identities of those who bear punishment costs or benefit from lower victimization rates. The marginal costs of punishment include funds that the government allocates to impose punishment, the psychic and financial costs that an offender pays, and the costs borne by families and friends of the offender. The marginal benefits include reductions in victimization, an improved sense of safety and well-being among local residents, lower crime abatement costs, and so forth.³

These types of costs and benefits are hard to measure, and we recognize that placing social weights on changes in these outcomes requires subjective judgments. However, the vast majority of the conclusions we reach hold given any social welfare function that we or many other scholars would deem reasonable. Our key conclusion is that a number of existing practices, rules, and laws generate punishment costs for many defendants that almost certainly exceed any social benefits derived from incapacitation or deterrence. Further, since minority citizens, especially Blacks, face criminal charges at

much higher than average rates, they bear disparate impacts from these excessively punitive policies.

We will not address how deontological concerns like proportionality inform judgments about racial disparities in court outcomes. If a researcher gave policymakers solid evidence that when courts choose not to incapacitate a robber who physically punches a victim who initially resists, the same robber will almost certainly maim or kill a future victim who fights back after being punched, our utilitarian framework dictates that this evidence is a potential justification for laws that mandate prison time for robbers who assault their victims. We are not interested in arguments about whether a prison sentence is a proportionate punishment for punching someone in the face and taking his phone. We care about how the social costs of punishment compare to the social value of the expected reduction in victimization.

If our framework also incorporated concerns about proportionality or related deontological matters, our analyses would lack discipline. Since Black citizens face criminal charges at much high rates than other groups, we could label almost any punitive law or practice “racist” by making a subjective argument that the practice in question seemed disproportionate to us. In contrast, our utilitarian approach provides a concrete definition of disproportionate or nonoptimal punishment. If the social costs generated by a punishment exceed the expected benefits to society, the punishment is too severe. If the marginal costs are less than the marginal benefits, the punishment is too lenient.⁴ In our framework, there are no trump cards like “fairness” or “proportionality” that take precedence over all other concerns. Optimal policy making always involves balancing marginal costs and benefits.⁵

3. Mark Cohen (2000) notes that, historically, many criminology studies that address the costs and benefits of changes in criminal justice policy did not consider punishment costs incurred by the offender. We do not consider any financial or psychic gains that criminals derive from engaging in crime as benefits to society, but we do consider the losses that criminals suffer from punishment as social costs.

4. We maintain the assumptions that, within the relevant set of policy alternatives, marginal costs of punishment are increasing, and the marginal benefits are decreasing.

5. We realize that other researchers, for example, Kevin Reitz and Cecelia Klingele (2019), blend deontological concerns with approaches that are largely utilitarian. We do not, in part, because the concept of a target outcome

Since our definition of racial bias relies on the concept of an optimal or target outcome for each case, classical statistics offers a useful starting point for understanding our approach. In the nineteenth and early twentieth centuries, researchers developed a formal set of tools for analyzing the outcomes of experiments and testing the hypotheses that motivated these experiments. Thus, all beginning students in probability and statistics learn that random variables are rules that assign numbers to the outcomes of experiments, and estimators are random variables because estimators take data produced by experiments as inputs and return numbers or vectors of numbers.

Consider researchers who seek to estimate the expected value of the impact of an experimental program. If the researchers gather data from a treatment group and a control group and then report the mean outcome difference between the treatment and control group, they are applying a specific estimator to the results of an experiment. Given certain assumptions about the experimental protocol, this estimator is unbiased—that is, if the researchers repeated the experiment hundreds of times and plotted the results of all trials in a histogram, the histogram would be centered on the true expected impact of treatment. On the other hand, if the researchers employ a biased estimator, the histogram will not be properly centered. When an estimator is biased, it produces results that, on average, miss a given target parameter.

If we think of court cases as experiments, we can assign numbers to various case outcomes, such as dismissals, verdicts, and sentences. As a result, we can view courts as random variables that assign numbers to the outcomes of experiments. Given this framework, we say that a court or a court system is biased if, on average, it systematically produces outcomes that are off target. To be clear, we are not saying that unbiased courts never make mistakes. We are saying that, on average, unbiased courts produce socially optimal outcomes—that is, for each type of case and among all types of defendants, unbiased courts produce verdicts that are not too punitive or too lenient on average.

If a court does systematically make mistakes by imposing sentences for certain types of cases that are not socially optimal, and if these mistakes impose greater costs on particular race groups, we say that the court produces racially biased outcomes.

Given this framework, we identify three sets of racially biased outcomes. To begin, consider a set of cases that involve a group of comparable defendants who face the same charge and assume that the socially optimal sentence for each defendant is the same. Further, suppose that the law permits judges to assign this socially optimal sentence to each defendant. In this setting, if researchers observe that the court assigns different average outcomes to defendants of different races, the court's decisions, taken as a whole, cannot be socially optimal. The existence of different average sentences among two groups that should receive the same average sentence implies that at least one group is receiving punishment that is, on average, either excessive or lenient relative to our socially optimal benchmark. Further, in this scenario, the fact that two groups that differ by race receive different average case outcomes implies that this failure to sentence all groups optimally generates racial disparities in case outcomes among comparable defendants.

However, we must also consider cases where the law does not permit socially optimal decisions. The law may require courts to treat two groups of defendants differently even though the socially optimal sentence for both groups is the same. In this second scenario, if defendants of one race are overrepresented in cases where the law requires excessive punishment, court actors produce inefficient racial disparities, even as they treat all defendants equally under the law. These disparities are forms of structural racism, since the structure of the sentencing laws requires different average treatment of defendants by race among sets of defendants who should receive the same expected sentences.

As an example, Michelle Alexander (2011) provides a detailed account of how federal sen-

for each case is central to our approach. In our framework, the optimal or target outcome for a case is unique and well defined because it is the outcome that maximizes a single utilitarian objective.

tencing guidelines in the 1980s required judges to treat defendants convicted of selling crack cocaine much more harshly than those convicted of selling powder cocaine, even though the former can be made from the latter. These guidelines were not optimal. Either the sentencing guidelines for powder cocaine offenses were too lenient or the guidelines for crack offenses were too severe. Further, the resulting racial disparities in sentences were a form of structural racism because Black defendants were much more likely than White defendants to face charges related to crack cocaine.

Finally, among broad categories of cases, the law may require courts to assign outcomes to all defendants that are systematically too lenient or too severe relative to some set of socially optimal target outcomes. Since racial minorities are more likely to appear in court, courts that routinely produce outcomes that are independent of defendant race but too severe produce racially disparate impacts, and the excessive severity built into the system is therefore a form of structural racism.⁶ As we will explain, many of the changes in sentencing laws that drove the prison boom that began in the late 1970s are difficult to justify on efficiency grounds, and while courts appear to have applied these new laws in a color-blind way, Black males bore a large disparate impact from these punitive policies because they face criminal charges at much higher rates than other groups.

In our framework, racial bias in court outcomes is always rooted in departures from social optimality. When comparable defendants of different races do not receive equal treatment under the law, the resulting disparities are rooted in the fact that the treatment assigned to at least one race group departs from our socially optimal benchmark. When the law requires excessive punishment for offenses that are particularly common among one racial group, the application of socially inefficient punishment again creates racial disparity, and

when the law requires excessive punishment for a broad range of offenses, racial groups that face higher overall arrests suffer disproportionate harm.

ANALYTICAL FRAMEWORK

We use the following notation to define outcomes and various key information sets available to policymakers and researchers. Each variable describes a characteristic or outcome of a particular case j that involves a particular defendant i . Yet, for ease of presentation, we suppress the $_{ij}$ subscripts on each variable.:

- D is a vector of outcomes or decisions that apply to a given case that involves a particular defendant, such as verdict, sentence, conditions of supervision, and so forth.
- X^* is the set of case and defendant characteristics that determine the target outcomes for the case—that is, these characteristics determine the expected costs and benefits of various outcomes that a court may produce for a given case.
- X^{law} is the set of case and defendant characteristics that the law requires or encourages the court to consider when deciding the case.
- \tilde{X} is the set of case and defendant characteristics that a researcher observes.
- R is a categorical variable that codes the race of the defendant j in case i .
- $F(D|X^*)$ is the distribution of outcomes produced by a court system conditional on the relevant cases and defendant characteristics.

In our framework, racial bias in the form of structural racism may exist even when comparable defendants facing comparable charges receive similar sentences regardless of race. However, most empirical work on racial bias in courts focuses on the possibility that similar

6. Policies that are color-blind and too lenient may also be a problem. For example, if violent crime victimization rates are higher among minority residents and these residents are victims of intra-race violence, lenient sentencing rules that failed to use incapacitation as a means of creating valuable protection for minority citizens could also represent structural racism. Given current criminal justice policies, this possibility is not our main concern, but as a matter of logic, leniency can also become a form of structural racism.

defendants facing the same charges may receive different sentences depending on their race. Thus, most empirical work on bias in courts focuses on group differences in conditional average case outcomes.

These studies seek to test some form of the following condition:

$$E(D|X^*) = E(D|X^*, R) \quad (1)$$

This condition requires that, given all relevant case and defendant characteristics, X^* , the expected outcome—for example, dismissal, conviction, length of sentence to probation or prison, and so forth—is the same regardless of defendant race. Judges may make mistakes. Courts may make mistakes when processing information. So, we should expect a distribution of case outcomes given any conditioning set, X^* , that determines the optimal outcomes for a case. However, a defendant's race is not an element of X^* , and conditional on X^* , the expected outcomes produced by a socially optimal court system should not depend on defendant race.

Conditioning Sets and Measured Disparities

In the analyses below, we often assume that prosecutors, judges, and other participants in courts have full information—that is, they see all the elements of X^* and X^{law} . Further, in many cases, we assume that $\tilde{X} \in X^* \cup X^{law}$ —that is, researchers see a subset of the information available to court officials. Since researchers can only condition on information they possess, the empirical literature on racial bias in court outcomes focuses on studies that test null hypotheses of the form:

$$H_0 : E(D|\tilde{X}) = E(D|\tilde{X}, R) \quad (2)$$

For now, we proceed under the assumption that $\tilde{X} = X^{law}$, since researchers who attempt to test whether defendants of different races receive equal treatment under the law often construct \tilde{X} with this goal in mind.

Given this assumption, there are at least two reasons that researchers may reject condition 2. First, even if we are willing to assume that conditioning on \tilde{X} is at least approximately equivalent to conditioning on X^* , conditional case outcomes may differ by race because court officials harbor explicit or implicit animus toward minority defendants, and this animus may produce more severe punishments for minority defendants compared to White defendants who pose similar risks to public safety, face comparable prospects for rehabilitation, and face comparable charges.

Second, conditioning on \tilde{X} may not be a good approximation for conditioning on X^* . Here, court officials produce expected outcomes that vary with the race of defendants conditional on \tilde{X} , even though these expected outcomes do not vary with race conditional on X^* . For example, consider a judge who learns that a convicted defendant lives in an area controlled by a violent street gang. This information is potentially relevant for sentencing if we assume the judge can use the information to form a better estimate of the likelihood that the defendant will commit future crimes (that is, the information is part of X^*); not measured by researchers (that is, the information is not part of \tilde{X}); and correlated with race. If our hypothetical judge weighs the defendant's residence in a violent neighborhood as an aggravating factor when sentencing the defendant, the judge is statistically discriminating against the defendant on the basis of a factor that is both correlated with race and not identified in the law as a mitigating or aggravating factor.

The consequences of this type of statistical discrimination are not obvious. If the judge weighs this factor properly when making decisions, the practice creates racial differences in the treatment of defendants, given $\tilde{X} = X^{law}$, but the judge is behaving optimally with respect to the welfare consequences of the decision. If the judge places improper weight on unmeasured factors that are correlated with race, the judge creates racial disparities in court outcomes that are socially harmful.⁷

7. Here, we assume that factors like the level of violence in the defendant's neighborhood are not factors that the law requires judges to consider. However, we assume that most sentencing statutes grant judges discretion to make subjective judgments about each defendant's propensity to commit future crimes according to a range of

It seems reasonable to expect that judges often possess relevant information that researchers can never recover from datasets. As a result, researchers who discover racial differences in average case outcomes given \tilde{X} often offer conjectures concerning whether such differences would remain if they were able to condition on X^* . Given this reality, some researchers argue that tests for racial bias should involve tests for racial differences in the distribution of future behaviors among defendants who face the same charges and receive the same verdict and sentence. The rationale behind these tests is straightforward. Consider a set of defendants who face the same charge, receive the same verdict, and receive the exact same sentence. If average future criminal justice outcomes among this set of defendants vary with the race of the defendant, then courts are assigning the same sentences to different groups of defendants who differ both by race and by their propensities to commit new crimes. To many, this pattern suggests that courts are making racially biased decisions. However, as we explain in the section “Outcome Tests,” researchers require strong assumptions in order to infer racial bias in judicial decision making from racial differences in post-sentencing outcomes among defendants who receive the same sentence.

Taken as a whole, existing research has shown that it is difficult to conduct compelling tests of the null hypothesis that comparable defendants receive comparable case outcomes regardless of their race. However, as we noted, even courts that treat all citizens equally under the law may produce racially biased outcomes for structural reasons. We have already mentioned Alexander (2011) and her treatment of differences in mandatory sentences for persons convicted of selling crack versus powder cocaine. In the section “Pretrial Custody,” we note that, as an institution, cash bail is likely a form of structural racism. Scant evidence links pretrial release decisions to the capacity and willingness of defendants to post bonds is an effective way to retain high-risk defendants or compel good behavior among released defen-

dants, but linking pretrial release to cash bonds does generate disparate harm for groups who have been harmed economically by past discrimination.

In addition, structural racism exists when the overall structure of court practices, sentencing laws, and rules governing community supervision create environments where courts routinely impose overly punitive sentences on all offenders in many offense categories. Since minority citizens, especially Black citizens, are much more likely to face arrest for a broad range of offenses, they suffer disparate harm from the ubiquitous assignment of excessive punishment to defendants. In the section “Sentencing Rules and Guidelines,” we discuss how a wave of punitive sentencing reforms drove the prison boom that began around 1980, and in the section “Incarceration and Public Safety,” we argue that existing evidence suggests that these reforms likely went too far. At the margin, the social costs of such large prison populations likely exceed any social benefits for the public, and since Black citizens are arrested at much higher rates than other groups, they suffer disparate harm from the sentencing laws that produce large prison populations.

In the section “Community Supervision,” we discuss the common practice of allowing parole officers to initiate the re-incarceration of parolees who violate technical conditions of their supervision, such as moving to a new residence without prior approval. There is little evidence that this practice generates important improvements in public safety, and because minority citizens are much more likely to be under parole supervision, they suffer disparate harm from this practice.

In the section “Simulating the Impacts of Reforms,” we present the results of several simulation exercises that attempt to assess how much policymakers can reduce absolute racial disparities in incarceration rates by making the practices and policies that govern pretrial release, sentencing, and parole supervision more efficient. The results suggest that it may be possible to significantly reduce racial disparities in

factors that the judge observes that are not recorded in court files. As we will note, when the law requires judges to consider irrelevant factors or forces judges to ignore relevant factors, we classify the resulting welfare losses as the result of structural problems.

incarceration rates by simply improving the rules that guide and constrain the decisions of judges, parole boards, and parole officers.

Alternative Taxonomies

Some may point to incomplete information as a potential source of bias in court outcomes. If judges and prosecutors do not observe the full set of defendant and case characteristics that are either socially or legally relevant for filing charges, conducting trials, negotiating plea bargains, and assigning sentences—that is, court officials produce case outcomes without full knowledge of X^* , we expect them to produce nonoptimal decisions, and since rates of criminal justice involvement are much higher in minority communities, these nonoptimal decisions may disparately impact these communities.

We do not disagree, but we note that the justice system in the US is adversarial. The job of defense attorneys is to present information to courts that increases the likelihood that courts will produce more favorable case outcomes for the defendants they represent. Thus, any racial disparities in case outcomes that arise because judges often do not possess relevant exculpatory or mitigating information or because judges are less likely to possess such information when defendants are Black or Hispanic are disparities that arise from structural racism. As an example, low funding levels for public defenders may represent a form structural racism because historic patterns of discrimination against Blacks imply that, on average, Black defendants have less financial capacity to fund their own defense.

We take a similar stance on the topic of implicit bias. A significant literature documents that actors in the judicial system implicitly associate negative personality traits with membership in racial minority groups. Scholars debate the meaning of these associations, but we have found no compelling evidence that implicit bias generates racial bias in actual judicial decisions. Further, if any study were to establish a link between implicit bias and racial bias in judicial outcomes, such a result would simply be evidence that the legal system is structurally flawed. We contend that decisions that are uninformed, ill-considered, or both are more likely to be either too punitive or too lenient relative to socially optimal benchmarks, and therefore legal systems should contain guardrails that prevent rushed and impulsive decisions.⁸

OUR TASK

We seek to understand how courts as well as the probation and parole authorities who enforce the sentences generated by courts contribute to racial disparities in criminal justice outcomes like felony convictions, the incidence of incarceration, and lengths of incarceration spells. Further, we examine the impacts of legislators who determine the laws and rules that govern the operation of courts.

While we focus on racial disparities among criminal defendants, we also note that court outcomes shape racial disparities in criminal victimization. Black citizens are not only more likely to face criminal charges but also more likely to be victims of crime.⁹ If courts are too punitive, they likely create disparate impacts

8. In NASEM (2021), Jennifer Eberhardt recommends “building friction” into police decisions and cites one program that required citizens reporting suspicious behaviors to go through a checklist to reduce the chances that the report was not motivated by racial stereotypes. If implicit bias does impact judicial decisions, researchers and policy analysts may need to consider how to build similar frictions in the process of reviewing charges, generating plea deals, and formulating sentences. See Jeffrey Rachlinski and colleagues (2008) and Bernice Donald and colleagues (2020) for discussions of the difference between harboring implicit associations and acting on them as well as how courtroom norms and procedures may shape this difference.

9. The National Academies of Sciences (NASEM 2022) documents large racial disparities in victimization, especially for violent offenses. In an analysis of data from the National Crime Victimization Survey for the years 2012 through 2019, the study finds property crime victimization rates per 10,000 households of 127 for Black Households and 112 for White households. Violent victimization rates not including homicides are also elevated for Black people, with a rate per 8.9 for Black people versus 6.8 for White people. The largest disparities are observed for homicide victimization. In 2015, the age-adjusted homicide rates (homicides per 100,000) among

on Black defendants. If courts are too lenient, they likely create disparate impacts on Black victims of crime.

However, we are not trying to understand all sources of racial disparity in criminal justice outcomes. Specifically, we take as given the cases filed in court and do not address racial bias in policing. While the National Academies of Sciences (NASEM 2022) documents that the racial composition of perpetrators by offense as described by crime victims is broadly consistent with racial disparities in arrest rates for serious person offenses, there is also evidence that both cross-jurisdiction heterogeneity in police practices and differential treatment of individuals citizens by race within specific agencies contribute to racial disparities in the composition of cases that police bring to courts. Steven Raphael and Sandra Rozo (2019) find that police are more likely to book juvenile arrests involving Black youth in comparison with White and Hispanic youth, and this difference is driven in part by harsher treatment for Black youth within specific jurisdictions as well as harsher average treatment for all youth in jurisdictions with large Black populations. Nonetheless, we will not explore police charging decisions here. Our goal is to examine the contribution of courts to racial disparities given the set of cases that police bring to the courts.

In the following sections, we discuss the screening and charging decisions made by prosecutors, pretrial custody decisions, trial outcomes, and the community corrections agencies that supervise persons on probation or parole. At each stage, we will examine evidence that actors in various court systems make systematically biased decisions, but we also pay close attention to the legal rules and institutions that constrain their choices in ways that create disparate impacts on defendants by race.

CHARGING DECISIONS

When police file charges, prosecutors must determine whether to proceed with a case. Pros-

ecutors quickly drop many charges because they conclude that the police have not presented evidence that supports the charges. Prosecutors also divert many cases to special programs that provide services for defendants and offer some defendants the opportunity to avoid traditional court sanctions by successfully completing program requirements. Finally, preliminary hearing judges dismiss many cases on the grounds that prosecutors have not established probable cause.

Prosecutors have near-absolute discretion in the selection of criminal cases. They have the power to reject new cases or drop existing cases, usually without providing a reason and without the decision being subject to review (Sklansky 2018). This has inspired concern that prosecutors may exercise biases, implicit or explicit, when they screen potential criminal cases.

Andrew Jordan (2022) finds that the screening function of prosecutors is important. Over half of Chicago defendants who are arrested but not convicted have their cases dropped at the felony review screening step. Chicago prosecutors also vary considerably in their apparent ability to identify and reject cases that will not produce convictions. However, there is little evidence that felony review prosecutors contribute to racially biased court outcomes. The set of cases offered by police features more low-quality cases against black suspects, but felony review does not exacerbate this disparity.

Other studies reach similar conclusions. Besiki Kutateladze and colleagues (2014) study multiple decisions made by New York City prosecutors. They find no significant effect of initial case screening because most cases are accepted. Turning their attention to subsequent dismissals, they find that prosecutors were actually more likely to dismiss a case when the defendant was Black or Latino. Alex Chohlas-Wood and colleagues (2021) use an algorithm to remove information from case reports that may signal the race of the suspect. They find that this has little effect on the charging decisions of prosecutors. They attribute

males are 35.4 for Black males and 3.6 for White males. The comparable figures for Black and White females are 4.9 and 1.7.

this to the absence of racial bias prior to their intervention.

The studies considered here all focus on large, urban prosecutors' offices. Prosecutors in less populous areas may make less benign use of their discretion, and if they do, researchers may find it difficult to detect these abuses of discretion. In smaller prosecutors' offices, a single prosecutor often handles each case from beginning to end, and this arrangement does not produce data that researchers can readily use to isolate the impacts of screening decisions.

PRETRIAL CUSTODY

A significant and growing body of research examines racial disparities in pretrial detention rates. In most jurisdictions, persons who are arrested and charged with crimes soon appear before a judge who explains the charges, informs defendants about the date of their next court appearance, and then decides whether and under what conditions the defendants may leave jail while their cases go through the courts. The options available to a bond court judge vary from case to case and by jurisdiction. Judges typically decide whether to release a defendant with no conditions, release a defendant subject to conditions, or deny bail.

The possible conditions vary by jurisdiction. Some courts allow judges to require defendants to post a bond that equals a portion of a nominal bail amount. If the defendant posts the bond and then appears at all future court dates without committing a new crime, the court may return all or a portion of the bond to the defendant. If the defendant fails to appear in court or commits a new crime, the defendant typically forfeits the bond, and the court may try to collect all or a portion of the remaining bail amount. Some courts allow defendants to borrow bond payments from private bail-bond companies. If defendants work with a bail bondsman to gain release from jail, the bondsman is liable for some or all of the financial penalties that arise when defendants fail to appear for court dates or commit new crimes. Some courts also allow judges to require defendants who leave jail to submit to specific monitoring programs or wear electronic bracelets

that allow the sheriff to monitor their movements electronically.

Here, we review findings in recent research on racial disparities in pretrial detention. We then discuss reform proposals that are informed by these results. We highlight proposals that eliminate cash bail because we conclude that, as an institutional practice, cash bail likely creates racial disparities in pretrial release rates that reflect racial differences in financial capacity rather than racial differences in the risks of pretrial violations given release. Thus, it may be possible to reduce the racial disparities in pretrial detention generated by cash bail without harming public safety.

Before we review the existing work on pretrial release decisions, it is important to recall the key elements of our framework. In equation 1, we provide a necessary condition for the absence of racial bias, and we argue that racial bias exists when the expected decisions made by a particular group of court actors, conditional on a set of relevant case and defendant characteristics, depends on the race of the defendants who appear in court. Much of the literature on bond court involves empirical work that tests the following condition

$$E(D|\tilde{X}, R = r) = E(D|\tilde{X}, R = r')$$

Here, \tilde{X} is the set of defendant and case characteristics that researchers measure, and D is an outcome of the bond court proceedings, such as the bail amount set by the judge, an indicator for whether the defendant is released within a week, or a measure of the number of days that a defendant spends in jail while the defendant's case is being resolved. Finally, r and r' are two race groups, for example, Black versus non-Black.

Conditional Disparities in Release Rates

A number of studies find racial disparities in pretrial release rates conditional on various sets of controls for case characteristics and the criminal history of the defendant. Stephen Demuth and Darrell Steffensmeier (2004) examine pretrial release decisions in the seventy-five counties covered by the State Courts Processing Statistics (SCPS) program in its 1990, 1992,

1994, and 1996 surveys. They find that Black and Hispanic defendants are much less likely than comparable White defendants to obtain pretrial release. However, this result appears to be driven by the inability of these defendants to post bail. Black defendants are no more likely to be denied bail than comparable White defendants, and among those not denied bail, bail amounts for Black defendants are comparable to those for White defendants. Hispanic defendants who are not denied bail face slightly higher bail amounts than other comparable defendants, but these differences could never account for the high pretrial detention rates that Demuth and Steffensmeier (2004) document among Hispanics. Whites are detained pretrial in 23.5 percent of cases. The corresponding rates for Blacks and Hispanics are 28.4 and 34.7 percent respectively, and the largest driver of these racial differences is the inability of Black and especially Hispanic defendants to post bail.

Using data from Detroit, Charles Katz and Cassia Spohn (1995) also find that Black defendants are less likely to obtain pretrial release than comparable non-Black defendants, but they find little evidence of racial differences in average bail amounts given case and defendant characteristics. It appears that Blacks in Detroit are more likely to remain in jail pretrial because they have less financial capacity to post bail.

Table 1 presents original results from the main felony bond court in Chicago, Illinois. We examine data from 2011 through the summer of 2017. We highlight racial differences in several bond court dispositions. First, we explore racial differences in the bail amounts set by bond court judges. For these analyses, we create a single bail index that assigns a bail amount to all cases, even those that do not involve bail explicitly. If the bond court judge releases the defendant without requiring the defendant to post bail, we code the bail amount as zero dollars. If the bond court judge denies bail, we code the bail amount as \$500,000.¹⁰

Judges rarely assign bail greater than \$1.5 million dollars, and over this period, judges in this bond court seldom deny bail. It appears that many judges use a \$500,000 bail to signal that they want the defendant to remain in jail.

For male defendants, we run a median regression of bail amount on an indicator for Black, indicators for class of crime, an indicator for a charge involving a violent crime, an indicator for having a prior felony arraignment, and an indicator for time period that captures changes in the administration of bond court. We find a median conditional racial gap in bail of \$5,000. We begin with median regression because these results are robust to the specific bail imputation rules we use for defendants who are either released with no bail or denied bail. When we run ordinary least squares OLS regressions of our bail index on defendant and case characteristics, we find that judges set bail almost \$1,025 dollars higher for Black male defendants.

When we conduct similar analyses for female defendants, we find slightly different results. The conditional median Black versus non-Black gap is \$0, and we find that, on average, Black female defendants face bail amounts that are almost \$1,870 lower than comparable non-Black defendants.¹¹

To place these results in context, \$5,000 is the default increment for bail decisions. Among the most common bail amounts in our data, the next lowest bail amount is \$5,000 less, and the next highest is \$5,000 more. Thus, among both men and women, both the mean and median gaps in conditional bail amounts are within one step on the schedule of the most common bail amounts we observe.

We also examine racial gaps in the likelihood that the court releases a defendant without requiring him to post any bail. Here, we see large Black versus non-Black gaps. In almost 19 percent of cases involving non-Black, male defendants, the court grants release without bail, and regression-adjusted release rates for Black males are almost 6 percentage points lower

10. The use of electronic monitoring is another source of variation in custody conditions in Cook County. In this exercise, we ignore the imposition of electronic monitoring as an additional release condition.

11. The *p*-value on this result is .067. The mean bail is just over \$54,000 among men and just over \$29,000 among women. Both means include our imputations of zero for defendants who left jail without paying any bail.

Table 1. Bond Court Outcomes: Cook County, Illinois

| Panel A: Racial Differences in Bond Court Decisions, Men | | | | | |
|--|-----------------------------|------------------|---------------------------|-----------------------|---------------------------------------|
| | No Conditions OLS (1) | Bail | | | Bail > 0 Median Regress. (5) |
| | | OLS (2) | Median Regress. (3) | Bail w/ OLS (4) | |
| Black | -.059 (.002) | 1,024 (506) | 5,000 (333) | -3,056 (561) | 0 (212) |
| N | 111,097 | 110,524 | 110,524 | 96,432 | 96,432 |
| Panel B: Racial Differences in Bond Court Decisions, Women | | | | | |
| | No Conditions OLS (1) | Bail | | | Median Regress. (5) |
| | | OLS (2) | Median Regress. (3) | OLS (4) | |
| Black | -.04 (.008) | -1,868 (1021) | 0 (693) | -4,165 (1333) | 0 (1246) |
| N | 14,870 | 14,689 | 14,689 | 10,974 | 10,974 |

Source: Clerk of Court of Cook County.

Note: This table presents results from five regressions. Each observation is a case in the main felony bond court in Chicago, Illinois. The cases cover 2011 through August of 2017. Each OLS regression contains controls for the year of case assignment, an indicator for a charge involving a violent crime, indicators for the class of the most serious charge, an indicator prior felony charges, and indicators for defendant age. Each median regression contains similar controls, but the controls for year of assignment are collapsed into an indicator for an early and late period. The dependent variable in regression column 1 is an indicator for whether the judge granted the defendant release with no conditions. In column 2 the dependent variable is the actual bail amount set by the judge or an imputed bail amount. We impute zero for persons released with no bail requirement. We impute \$500,000 for defendants denied bail or assigned a bail amount greater than \$500,000. We present median regression results in column 3, since median regression results are robust to the precise imputed values used for extreme cases. Columns 4 and 5 report similar regression results, but here we eliminate persons who were released without any bail requirement. The sample sizes are largest for the regressions in column a because there are records where the judge assigned bail, but the Clerk of Court did not record the bail amount.

than the rates for comparable non-Black defendants. Among females, we see a non-Black release rate of 32 percentage points and regression-adjusted rates that are 4 percentage points lower than rates for comparable Blacks.

This disparity in direct-release rates appears to be a key source of racial inequality in historical data from Chicago’s main bond court. Among defendants who are not released without bail, we see no racial gap in conditional median bail amounts among men or women. Fur-

ther, among this sample of defendants, Black men and women face bail amounts that are respectively \$3,056 and \$4,166 lower, on average, than the bail amounts facing comparable non-Black defendants.

Even though Black defendants do not face higher bail amounts than comparable White defendants who are also required to post bond as a condition of release, racial differences in the likelihood that defendants post bail in a timely manner are another key driver of overall

racial differences in pretrial custody outcomes. Here, we examine racial differences in the number of days that defendants spend in jail between their bond court hearings and the dates that their cases are resolved, and we make these comparisons among defendants whom the court requires to post the same bail. The results show that racial differences in the likelihood that defendants post even modest bail amounts contribute significantly to racial differences in jail time served.

Figure 1 presents regression results for male defendants. We regress jail time on the controls in our bail regressions plus a quartic in the bail amount and a separate quartic interacted with an indicator for Black. We do not include cases where defendants were released without bail or denied bail, and we further restrict the sample to defendants facing a bail amount between \$1,000 and \$250,000.¹²

The figure presents two plots of predicted days spent in jail as a function of the bail set in bond court. These plots describe results for two hypothetical defendants. One is Black, and one is not, but they both are less than twenty years old, have no prior convictions, and face a similar charge. We plot results for all bail amounts that appear in our data at least two hundred times for both Black and non-Black males.¹³ The graph starts at \$5,000 bail. Here, the Black versus non-Black gap in predicted jail time is less than four days, but the gap grows with bail. At \$50,000, the gap is just over seventeen days. Among those facing a bail amount of \$250,000 the gap is thirty-nine days.¹⁴

While the racial gaps in expected days spent in jail are small given low bail amounts, racial differences in the likelihood that a defendant

leaves jail quickly are not trivial. When we examine conditional racial gaps in the probability that a defendant leaves jail within ten days of his bond court hearing, we find that, given a bond amount of \$5,000, non-Black males are almost 19 percentage points (72.2 versus 53.5) more likely than comparable Black males to leave jail within ten days of their bond hearings. Note that defendants facing a \$5,000 bail requirement must post only \$500 to gain release. Thus, even modest posting requirements generate significant racial differences in the prevalence of prompt exits from jail.¹⁵

In Chicago, Black defendants are significantly less likely to gain pretrial release with no conditions in bond court than non-Black defendants who are observationally similar. However, in Detroit and SCPS counties, bond court judges typically make similar custody rulings for comparable defendants of different races. Nonetheless, in all three cases, we see that the practice of requiring defendants to post bonds to gain release creates racial disparities in pretrial release rates, even among defendants who receive the same bail disposition. We conjecture that racial disparities in the economic circumstances of defendants drive this pattern.

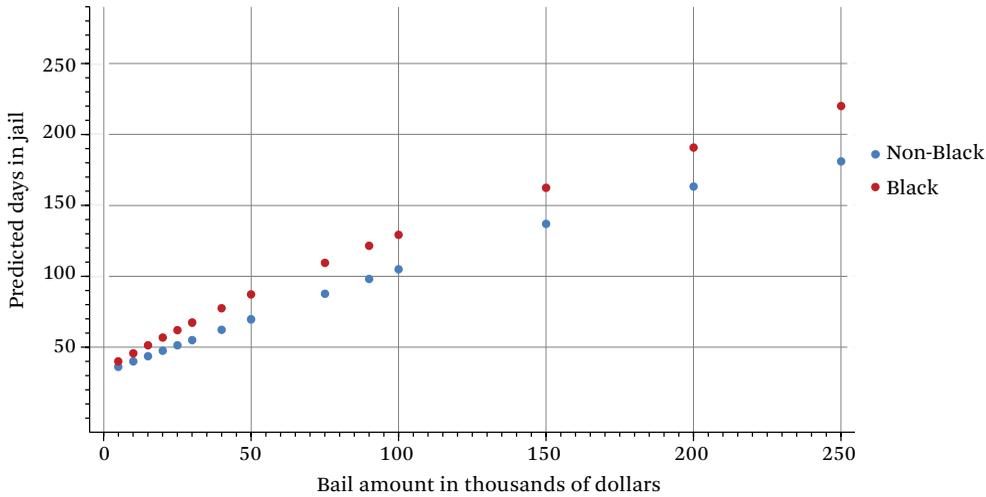
These patterns are consistent with the view that, as an institution, cash bail is a form of structural racism. Aurelie Ouss and Megan Stevenson (2023) examine a 2018 reform in Philadelphia that eliminated cash bail for defendants charged with a specific set of offenses. They compare defendants before and after the reform who are charged with offenses that were or were not affected by the 2018 reform, and they find no evidence that the monetary incentives created by cash bail reduce pretrial mis-

12. Among defendants assigned cash bail, less than 1 percent receive an amount less than \$1,000, and less than 5 percent receive more than \$250,000. By far, the most common dollar amount greater than \$250,000 is \$500,000, and some judges appear to use this as a substitute for bail denied.

13. Even though bail amounts of \$500,000 appear more than two hundred times, we do not include this amount. As we note, this amount seems to signal that judges intend to hold the defendant in jail. Bail amounts between \$250,000 and \$500,000 or in excess of \$500,000 dollars are rare.

14. These gaps follow the same pattern we observe in the raw data if we simply plot average days in jail by race within samples of defendants facing the same bail amount. The raw gaps increases from eight days to forty-six days as we move from \$5,000 to \$250,000.

15. For reasons that we do not fully comprehend, Black versus non-Black gaps in expected time in jail conditional on required bail are much smaller among female defendants.

Figure 1. Days in Jail, Cook County, Men

Source: Clerk of Court of Cook County.

Note: This figure presents two sets of predicted values. Each observation in our data is a case in the main felony bond court in Chicago, Illinois. The cases cover 2011 through August of 2017. We restrict the sample to men who are assigned bail amounts between \$1,000 and \$250,000. We then regress time spent in jail on an indicator for Black, controls for the year of case assignment, an indicator for a charge involving a violent crime, indicators for the class of the most serious charge, an indicator prior felony charges, indicators for defendant age, a quartic in the bail amount set, and a complete set of interactions between this quartic and the indicator for Black. We define time spent in jail as the number of days spent in jail during the interval between the bond court hearing and the date a case is dropped, dismissed, or decided. Here, we plot the race-specific expected days in jail for a Black versus non-Black defendant who is comparable with respect to the controls other than bail amount set.

conduct. Further, the literature contains no evidence that, if case characteristics and defendants' criminal histories are held constant, the financial capacity of defendants is highly correlated with their propensity to avoid pretrial misconduct given release. Finally, in many jurisdictions, there is no reason to expect cash bail to work well as an incentive scheme. Economically disadvantaged defendants are often incapable of paying the court cost charges they already face, and they are also judgment proof. As a result, any bond that they can post may well be applied to their court costs even if they incur no violations, and if they do earn violations, the court may find it hard to collect the resulting financial penalties that the law requires.

In sum, there is little evidence that, if case and defendant characteristics are held constant, defendants with greater financial capacity are less risky. Further, cash bail systems of-

ten provide weak ex post incentives for better defendant behavior while increasing expected jail time among groups that are economically disadvantaged. These facts imply that cash bail systems are inefficient. Further, since many forms of historical discrimination have reduced the current financial capacity of minority defendants, the disparate harm they suffer from the use of cash bail is a form of structural racism.

Consequences of Disparities in Release Rates

If cash bail systems are a form of structural racism, what are the consequences of the racial differences in conditional release rates that cash bail systems generate? The most obvious consequence is that a disproportionate share of the defendants who forfeit their liberty while waiting for their cases to resolve are non-White. In a case where the prosecutor has overwhelming evidence against the defendant and a guilty

verdict triggers a mandatory term of incarceration, this denial of liberty does not typically involve extra time behind bars since the defendant gains credit for time served in jail that is subtracted from the mandatory prison sentence.¹⁶ However, many cases are not open and shut, and a number of studies conclude that gaining pretrial release reduces the likelihood that a defendant will be convicted. Pretrial release makes it easier for defendants to coordinate with counsel concerning their defense, and some researchers claim that holding defendants in jail gives prosecutors more leverage in plea bargain negotiations. As a result, some defendants may plead guilty rather than attempt to prove their innocence at trial just to get out of jail sooner.¹⁷

Outcome Tests

Another literature looks for evidence of racial bias in bond courts by comparing racial differences in pretrial misconduct among those released before their cases are resolved. Ian Ayres and Joel Waldfogel (1993) find that, if bail amounts are held constant, third-party bail bondsmen in New Haven, Connecticut, charge lower co-pay rates to minority defendants, a finding the authors interpret as evidence that released Black defendants are less risky than released White defendants. Using data from Miami and Philadelphia, David Arnold and colleagues (2018) find that, among defendants released pretrial whom they infer to be marginal candidates for release, White defendants are significantly more likely to be rearrested than

Black defendants. They interpret this result as evidence that, on average, judges hold Black defendants to a higher standard when making release decisions.

However, the conclusions that scholars draw from tests for racial discrimination in court decisions based on racial differences in post-adjudication outcomes typically rest on strong assumptions. Researchers have long understood that they cannot draw firm conclusions concerning whether judges set different standards for release depending on the race of defendants solely on the basis of average outcome differences by race among all defendants who are released. Standard models of discrimination make predictions about differences in outcomes among marginal defendants in different groups. Thus, researchers must isolate marginal cases to conduct valid outcome tests.¹⁸ In addition, recent work shows that even in settings where researchers can isolate marginal cases, more robust models of decision-making raise new concerns about the use of outcomes tests as tools for identifying racial bias in judicial decisions.

In a recent article, Ivan Canay and colleagues (2024) invite readers to consider bond court judges who make decisions using the following threshold rule: grant pretrial release if

$$E(\Delta|r, \nu) \leq \tau(z, r, \nu)$$

where

$$\tau(z, r, \nu) = c(z, r, \nu) + \lambda(z, r, \nu) + \beta(z, r, \nu)$$

16. If researchers had a way to measure case strength, it would be interesting to know whether defendants facing strong cases against them and mandatory prison sentences following conviction are less likely to post bonds given the opportunity.

17. While pretrial detention may also make it more difficult to mount a defense, in many instances the urgency associated with being incarcerated as well as the fact that, for less serious offenses, pretrial detention periods may be greater than the actual sanction if convicted create incentives for defendants to accept guilty pleas. See Will Dobbie and colleagues (2018) for estimates of the causal effect of pretrial detention on conviction rates. See Johanna Lacoe and colleagues (2023) for evidence that pre-arraignment legal services reduce both pretrial detention as well as the likelihood of conviction through plea bargains.

18. This challenge is often called the infra-marginality problem. Imagine a world where each defendant i is characterized by a recidivism rate, δ_i , and all judges observe each defendant's rate. Next, assume that all judges follow a simple rule when sentencing all convicted offenders: assign defendant i to probation instead of incarceration if $\delta_i < \bar{\delta}$. Given data from such a court system, we may still observe that the rates of recidivism events among those assigned to probation nonetheless differ by race because the distribution of δ given $\delta < \bar{\delta}$ differs by race.

For simplicity, we assume that all defendants are the same on dimensions that researchers measure—that is, the following results are conditional on \tilde{X} .

The left-hand side is the expected cost of releasing the defendant, Δ is the violation cost society incurs from releasing the defendant, r is the race of the defendant, v is a defendant or case characteristic that the judge observes that is not observed by researchers, and z indexes judges.

On the right-hand side, $\tau(z, r, v)$ is judge z 's cost of detaining the defendant. The function $c(z, r, v)$ is judge z 's assessment of the social cost of detaining the defendant. For our purposes, $\lambda(z, r, v)$ represents the mistake that judge z makes when she forecasts Δ , and $\beta(z, r, v)$ represents judge z 's personal distaste for detaining the defendant.

In this framework, judge z makes racially unbiased decisions if $\tau(z, r, v) = \tau(z, v)$ for all r and v , but Canay and colleagues (2024) demonstrate that, even when judges are unbiased, post-verdict data on defendant outcomes may contain significant racial differences in average post-sentencing outcomes among marginal defendants.¹⁹ Further, the logic of their argument holds when we assume that the distribution of v is independent of race.

We use a simple example to illustrate the key insight that drives the reasoning of Canay and colleagues (2024). Consider two defendants. The first defendant faces multiple charges, and if convicted, faces a mandatory two-year prison sentence. Further, the police have a high-quality video that clearly establishes the defendant's guilt. The second defendant faces similar charges, but this defendant has no prior criminal record, and the state's case is so weak that a trial judge will likely dismiss the case in the weeks to come. In the first case, each day that the defendant spends in jail awaiting trial earns him a credit for time served that reduces the time he will almost surely serve in prison. In the second case, each day in jail is likely just a day of lost liberty. Since the social costs of detaining the first defendant, $\tau(z, r, v)$, are

lower, an optimal judge should be more reluctant to release him. Put differently, if detaining the defendant has little impact on the total expected punishment he receives, an optimal judge will only release him if $E(\Delta|r, v)$, the expected violation cost associated with releasing the defendant, is small.

How does this observation impact how we think about the content of marginal outcome tests for racial bias in judicial decisions? Let v be an index of the expected time served for the defendant, and let z be an unbiased judge who makes socially optimal release decisions. Next, consider a set of cases decided by z and assume that the defendants in this set are identical on all traits that z observes except race and v . Next, recall that $\tau(z, v)$ is declining in v and assume that $E(\Delta|r, v)$ is an increasing function of v for all defendants of all races. Further suppose that, for any fixed v , expected violation rates are lower for Black defendants. Since $\tau(z, v)$ is a decreasing function of v , the value v^* that defines expected time served for the marginal defendant, that is, $E(\Delta|r, v^*) = \tau(z, v^*)$, will be greater in the sample of Black defendants, and the expected violation risk, $E(\Delta|r, v^*)$, will therefore be lower for marginal Black defendants than marginal non-Black defendants. In this example, average violation rates differ by race among marginal cases even when a court makes racially unbiased and socially optimal release decisions.

This is not a far-fetched scenario. If past discrimination by police inflates the criminal records of Black defendants relative to non-Black defendants who have engaged in the same illegal activities, it is reasonable to expect that, among defendants facing the same expected time served, Black defendants have lower expected violation rates given release. Further, this is just one of many plausible scenarios that raise concerns about marginal outcomes tests as tests for racial bias in release decisions. Any factor v that causes $\tau(z, v)$ to slope up or down is a potential source of another example.²⁰

Our second concern about the interpretation of outcome tests in this setting is less for-

19. They also demonstrate that biased judges may generate data that do not fail these tests.

20. Suppose that expected violation risk is the sum of a constant that differs by race and a monotonic function of how long the defendant has been without work, v . Here, the constant captures the effect of employment

mal and focused more on the details of bond court operations. The current literature proceeds as if the bond court judge determines the release outcome. However, in many cases, bond court judges do not determine release outcomes. They set release conditions, and defendants either meet or do not meet these conditions by posting required bonds, arranging for housing that satisfies requirements for electronic monitoring, and so forth. If judges can perfectly predict how defendants will respond to the conditions they set, this distinction does not matter. However, if judges have limited information concerning how defendants are likely to respond to different potential release conditions, releases outcomes in bond court are typically not decisions that judges make but decisions that defendants make in response to constraints that judges choose.

Historically, many bond courts have required judges to make decisions on bond and other release conditions in settings that put clear pressure on these judges to make quick rulings without access to formal risk assessments or information about the financial capacity of defendants.²¹ In recent years, some jurisdictions have provided bond court judges with detailed risk assessments of each defendant as well as information about the financial capacity of defendants,²² but much research on pretrial release outcomes comes from bond courts that asked judges to set release conditions with scant information about each defendant's violation risk or financial capacity. It is possible that, given these conditions, bond court judges adopt rules of thumb that assign higher bail amounts to defendants who face

more serious charges and have longer criminal histories. If this is the case, racial differences in violation rates given release may tell us nothing about how racial bias impacts the decisions made by bond court judges. Rather, these differences may arise primarily because the institutional practice of asking bond court judges to set bail without knowing anything about the financial capacity of defendants leaves many defendants who are low risk but poor in jail pretrial, and the prevalence of these defendants may well differ by race.

Fairness Tests

While tests for conditional differences in mean release rates or differential violation rates given release are the most common tests for bias in bond court outcomes, some studies test alternative conditions that, given our framework, do not make sense as tests for racial bias. These studies test what many researchers call fairness conditions. These conditions typically take the following form:

$$\begin{aligned} E(S|\tilde{X}, R = r, Y(1) = 1) \\ = E(S|\tilde{X}, R = r', Y(1) = 1) \end{aligned} \quad (3)$$

$$\begin{aligned} E(D|\tilde{X}, R = r, Y(1) = 1) \\ = E(D|\tilde{X}, R = r', Y(1) = 1) \end{aligned} \quad (4)$$

where S is a risk score or some other metric that the court uses to assess how risky it would be to grant pretrial release to the defendant. D is the release decision, and $Y(1)$ is a potential outcome. $Y(1) = 1$ implies that, if released, the defendant engages in misconduct before his

discrimination by race. If an unbiased judge, z , incorrectly assesses how violation risk varies with nonemployment, v , average outcomes among the defendants that she deems marginal candidates for pretrial release will differ by race, even if these mistakes, $\lambda(z, v)$, are independent of race. Both David Chan and colleagues (2022) and Janet Currie and colleagues (2025) explore related themes in their work on the systematic and frequent diagnosis and treatment mistakes made by many doctors.

21. See Steve Bogira (2006) for a description of the main felony bond court in Chicago, Illinois, prior to reforms that began in the 2010s.

22. See www.psapretrial.org for more on risk assessment tools. In September of 2017, Chief Judge Timothy Evans instructed bond court judges in Cook County, Illinois, to ask questions about the financial capacity of defendants and to set bail amounts after considering each defendant's ability to pay. The wording of the order implicitly asserted that, prior to this order, most bond court judges had been setting bail amounts without much concern for or information about the defendant's economic circumstances.

case is resolved. $Y(1) = 0$ implies that, if released, the defendant does not violate the conditions of his pretrial release. For simplicity, we assume that defendants cannot generate violations while in custody. So, we set $Y(0) = 0$ for all defendants.²³

Conditions 3 and 4 state that, among defendants who share a common \tilde{X} and would also earn violations if released, the expected risk score and the expected release outcome do not vary with race. Some analysts associate the concept of racial fairness with one or both of these conditions, and Arnold and colleagues (2022) construct a test for disparate impact around condition 4.

Tests of these conditions must overcome estimation challenges since $Y(1)$ is never observed for defendants who are not released, ($D = 0$). However, we argue that, even in a hypothetical setting where researchers could magically observe $Y(1)$ for all defendants, researchers should not use these conditions to test for racial bias in pretrial procedures. We use a simple example to illustrate this point.

Generating Data from an Optimal Bond Court

Imagine the following data-generating process for a bond court. When an arrested offender is charged, the offender is examined by a risk-assessment team. Above, we define S as a random variable that is the result of the assessment process. The assessment team observes the complete set of characteristics that determine violation risk—that is, $\tilde{X} = X^*$, and we assume that, conditional on X^* , the probability of misconduct given release is independent of race. We also assume that the assessment team weighs all factors in X^* optimally. As a result, $S(X^*) = P(Y(1) = 1|X^*) = P(Y(1) = 1|X^*, R)$. This means that, for defendants of all races, the score produced by the assessment team equals the true probability that $Y(1) = 1$ given X^* . Finally, we assume that the court uses a cutoff rule to make custody decisions—that is, let x_i^*

denote the relevant characteristics for defendant i , and let $s_i = P(Y(1) = 1|X^* = x_i^*)$. Then, the court releases defendant i iff $s_i < \bar{s}$. For simplicity, we assume that \bar{s} is a socially optimal cutoff rule that does not vary with defendant characteristics.²⁴

Our hypothetical bond court cannot know with certainty whether a given defendant will earn a violation if released pretrial, but the court does make color-blind decisions that are socially efficient. Nonetheless, as long as the distribution of S varies with race, the data that this optimal court produces will always violate the fairness conditions stated in equations 3 and 4 above.

This result is quite general, but we illustrate the result here with a parametric example. To further ease exposition, we suppress explicit conditioning on defendant characteristics. Consider two populations of defendants, r and r' . Among defendants of race r , $S \sim U(p_l, p_h)$, where $0 < p_l < p_h < 1$. Among defendants in r' , $S \sim U(cp_l, cp_h)$, where $c > 1$ is a constant that shifts the distribution of scores to the right and also increases the spread in scores, while maintaining a region of overlap with the distribution of scores for group r .²⁵ We assume that \bar{s} falls in this region of overlap—that is, $cp_l < \bar{s} < p_h$.

If these two populations of defendants appear before our optimal bond court, the court will generate optimal risk scores and optimal release decisions on the basis of these risk scores. Nonetheless, it is straightforward to show that, given these optimal decisions:

$$E(S|R = r', Y(1) = 1) = c * E(S|R = r, Y(1) = 1)$$

The expected risk score among the r' defendants who fail given release, $Y(1) = 1$, is c times greater than the expected score among the parallel group of r defendants. Further, it is straightforward to show that among those with $Y(1) = 1$, the optimal bond court releases a smaller fraction of group r' .

23. See Shira Mitchell and colleagues (2021) for a discussion of fairness concerns versus predictive performance in the context of algorithmic scoring rules.

24. In the section “Outcome Tests,” we point out that such cutoff rules may not be optimal. Here, we adopt this simple rule to ease exposition.

25. We require that $1 < c < \min(\frac{p_h}{p_l}, \frac{1}{p_h})$.

If we start by assuming that it is not possible to perfectly predict defendant behavior given release, then if the distributions of recidivism risks differ by race, optimal scoring rules and optimal release rules that are free of racial bias will always generate data that fail the fairness tests presented in equations 3 and 4. It is also possible to create biased decision rules that produce outcomes that violate fairness conditions, but in any context where distributions of violation risks differ by race, *all* unbiased decision rules always produce outcomes that violate fairness. So, when researchers employ data from the US and conclude that a set of risk-assessment scores or bond court release decisions fail a fairness test, they have learned little about whether the assessment scores or the release decisions were contaminated by racial bias.²⁶

Fairness concerns may loom large in a scenario where health professionals employ a racially unbiased set of diagnostic tools to detect disease (for example, early stage cancers) and recommend early treatment on the basis of these test results. In this setting, an unbiased diagnostic test may generate lower treatment levels among patients in racial groups with below-average disease prevalence.²⁷ Since unbiased diagnostic tests may become more fair if they become more accurate,²⁸ violations of fairness criteria in medical testing motivate efforts to improve the accuracy of diagnostic tests. However, medical testing differs from bond court in one key dimension. Each tumor that a medical team tests is either malignant or benign before it is tested. So, perfect classification is possible in theory. In contrast, defendants who stand before a bond court judge are neither violators nor non-violators. They are human beings, and if released, they will exercise personal agency by responding to the challenges and opportunities that confront them after release. In our framework, optimal bond courts do not attempt to classify defendants as violators or non-violators. Optimal bond courts

try to predict the expected future costs of releasing a defendant and weigh these costs against the social costs of detention.

In an influential recent paper, Arnold and colleagues (2022) propose a sophisticated test for disparate impact in New York City bond courts that is a variation on the fairness test presented in 4 above. In our notation, they define two forms of disparate impact:

$$\Delta_0 = E(D|R = \text{White}, Y(1) = 0) - E(D|R = \text{Black}, Y(1) = 0)$$

$$\Delta_1 = E(D|R = \text{White}, Y(1) = 1) - E(D|R = \text{Black}, Y(1) = 1)$$

and develop tests for disparate impact by estimating weighted sums of these components. However, we have already shown that, even if the New York City bond courts are free of racial bias, we do not expect $\Delta_1 = 0$, and a similar argument holds for Δ_0 .²⁹

In addition, according to our reading of the literature on disparate impact, we contend that, any plaintiff who tried to use the methods of Arnold and colleagues (2022) to convince a court that the New York City bond courts produce a disparate impact on Black defendants would fail. Historically, tests for disparate impact in employment cases have involved comparisons among applicants who, given the information available at the time of application, have the same expected productivity in a specific position. So, we anticipate that disparate impact cases concerning bond court outcomes would hinge on comparisons among defendants with the same expected violation risk and not defendants who share a common potential future outcome, $Y(1) = 1$ or $Y(1) = 0$. Further, the traditional disparate impact standard does not require that a bond court employ optimal methods to calculate expected violation risks. It simply states that the court cannot screen on metrics that do not bear a *demonstrable relationship* to true violation risk.

26. See Kleinberg et al. 2017 and Mitchell et al. 2021.

27. See Paulus and Kent 2020 for a discussion of these issues.

28. Tests that are unbiased and 100 percent accurate are always fair.

29. We know of no results that allow us to determine what the sum these biases should be in different contexts.

The test that Arnold and colleagues (2022) propose is novel and interesting, and other researchers may find fruitful uses for the methods they develop. Our concern is that, as a general rule, the fairness tests presented in equations 3 and 4 do not provide information that is relevant for debates about optimal pretrial procedures. Standard tests for disparate impact do not take this form, and our optimal bond court example demonstrates that researchers should expect the Arnold and colleagues' (2022) method to detect racial differences in release rates given $Y(1) = 1$ or $Y(1) = 0$, even in settings where bond courts are making release decisions that are socially optimal and unbiased.

Recent Reform Efforts

Given the growing evidence that cash bail systems are inefficient and that minority defendants suffer disparate harm from their design flaws, several states, cities, and counties have enacted bail reform measures in the last decade, and many more are considering them. Isabella Jorgensen and Susan Smith (2021) document several bail reforms across the country. They highlight recent comprehensive reforms enacted by New Jersey in 2017 and Illinois in 2023. The New Jersey reform did not eliminate cash bail entirely, but it took substantial steps to establish a presumption of release and reduce reliance on financial conditions. This reform substantially reduced jail populations in New Jersey without an attendant increase in rates of failure to appear or pretrial misconduct (Grant 2019).

The Illinois reform is the first statewide law to flatly eliminate cash bail, though it retains supervised release via electronic monitoring and mechanisms for prosecutors to request that defendants be detained prior to trial. There is not yet sufficient follow-up data to study the effects of this new regime.

Other reforms are more local and hence more tenuous. The district attorney of Philadelphia stopped requesting cash bail in a wide

range of cases beginning in 2018. Ouss and Stevenson (2023) evaluate this policy change and find limited impacts. Bail judges were still able to give cash bail even when the district attorney did not request it. Rates of cash bail did fall, but rates of pretrial detainment remained constant. This suggests that the main effect of the policy was to remove cash bail in cases where defendants would have otherwise still been able to pay. Ouss and Stevenson (2023) do not find any change in misconduct rates among these affected defendants.

On the other hand, both Georgia and Texas recently passed state laws that strengthen cash bail as an institution.³⁰ Further, Jorgensen and Smith (2021) note that all state legislatures have the power to adopt laws that prevent local courts from pursuing reform efforts.

As debates over bail reform proceed, we note that successful reforms may still become unpopular. Imagine a reform that eliminates cash bail, improves risk-scoring procedures, and enhances the quality of pretrial monitoring services. Now, suppose the reform reduces jail populations by one half while fostering a 10 percent increase in the total number of pretrial violations. Even if the social costs associated with this 10 percent increase in violations are small relative to the large reduction in incarceration costs, opponents of the reform can generate headlines in local media that read "bail reform leads to increase in crimes committed by defendants waiting for verdicts." Further, since forty-five of the fifty states allow bail bondsmen to make money by providing loans to persons who do not have the funds required to post bail, there are firms in most states that have a financial interest in preserving cash bail systems.³¹

CASE OUTCOMES

Now, we turn to the outcomes of cases brought by police that survive preliminary review by prosecutors. Here, we focus on the nominal outcomes specified in the verdicts and sentences produced by courts. The influential Na-

30. See McCullough 2021 and Amy 2024.

31. Only Illinois, Kentucky, Massachusetts, Oregon, and Wisconsin do not allow bail bondsmen. See www.ncls.org for more details.

tional Academies Press volume on sentencing (Blumstein et al. 1983) correctly notes that actors in parole and probation systems make numerous decisions that create variation in the actual punishments experienced by offenders who receive the same nominal sentence. We deal with these issues in the next section.

Charging Decisions and Plea Bargains

In the section “Charging Decision,” we discuss the screening procedures that determine whether charges filed by police evolve into cases that courts decide. Some charges do not withstand basic scrutiny, and some defendants are diverted into programs that offer rehabilitation services or special programs. However, among the cases that remain, prosecutors have considerable discretion over the specific charges they file against defendants and over the terms of plea bargains that they offer.

Charging decisions and plea bargains are intrinsically linked because the charges filed against a defendant determine both the likelihood that prosecutor will be able to secure a conviction if the case goes to trial and the severity of the expected sentence if the prosecutor does secure a conviction. Prosecutors realize that the vast majority of defendants are risk averse and are therefore able to use charges that carry mandatory prison sentences as leverage in negotiations over plea deals.

Since most criminal cases do not go to trial but end when the defendant accepts a plea agreement, researchers have explored the impacts of the practice of settling cases through plea bargains. Many findings support the idea that prosecutors leverage the risk aversion of defendants in plea bargaining negotiations, and in jurisdictions where the same group of prosecutors handle charging decisions and plea negotiations, initial charging decisions are an important component of these efforts.

Within our framework, two questions about plea bargaining are immediate. Do prosecutors file charges and negotiate plea deals in a manner that exhibits bias against minority defen-

dants? Does the practice of settling so many cases through plea bargains rather than a mechanism like bench trials produce inefficient outcomes that create disparate impacts on minority defendants?

Marit Rehavi and Sonja Starr (2014) demonstrate how bias may operate through charging decisions. They examine data from federal courts and find noteworthy racial differences in the distribution of time served in prison among comparable defendants in the same initial arrest category. The median Black defendant can expect to receive a sentence that requires him to serve roughly 9 percent more time in prison than a similar White defendant, and Rehavi and Starr (2014) show that this result is driven almost entirely by the choices of prosecutors. Holding arrest records and defendant characteristics constant, these researchers find that prosecutors are much more likely to charge Black defendants with crimes linked to mandatory minimum prison sentences. Further, Rehavi and Starr (2014) conclude that, if federal prosecutors did not discriminate in this manner, the steady-state population of Black inmates in federal prison would be about 10 percent smaller.

It is logically possible that racial differences in the nature of actual crimes committed exist within the 134 arrest categories that Rehavi and Starr (2014) employ, and that these differences justify the resulting racial differences in sentence lengths, but we have no evidence to support this claim, and we do know that, if Rehavi and Starr (2014) had introduced controls for the final charges facing each defendant, they would have found insignificant racial disparities in case outcomes. Rehavi and Starr (2014) illustrate that including variables in \tilde{X} that are chosen by court officials may produce misleading results concerning the existence of bias in the courts. Rehavi and Starr (2014) argue that conditioning on the charges filed by prosecutors is equivalent to conditioning on the mechanism that prosecutors employ to exercise racial bias.³²

32. Similar debates surround the use of controls for the criminal history of the defendant. These controls often shrink racial disparities in sentencing outcomes, but why? Ben Grunwald (2023) provides evidence that Black defendants may have a higher probability of being arrested conditional on criminal activity, and this would imply,

David Mustard (2001) also finds significant racial disparities in the federal sentencing outcomes conditional on a significant set of controls. He concludes that a key driver of more punitive sentencing outcomes for racial minorities is that prosecutors are more willing to ask judges to sentence defendants below the relevant sentencing guidelines as a reward for cooperating with prosecutors when dealing with White defendants. Given the framework we develop in the section “Analytical Framework,” we are reluctant to offer a definitive interpretation of these results. Do Mustard’s findings demonstrate that federal prosecutors have racial animus toward non-White defendants or that they statistically discriminate against non-White defendants because they correctly or incorrectly judge that White defendants are less prone to re-offend than similar non-White defendants? Or do prosecutors find it easier to obtain valuable cooperation from White defendants, and if so, why?³³

Some readers may conclude on the basis of these results that reforms that would eliminate plea bargains and mandate bench trials as the default means of resolving cases may both reduce racial disparities in case outcomes and improve social welfare. However, Jordan (2024) points out that the welfare analysis of plea bargaining as an institution is complex. Plea bargains provide insurance for defendants, and existing evidence suggests that many defendants are quite risk averse and value this insurance. It is logically possible that restrictions on the use of plea bargains could reduce racial disparities in sentencing outcomes while harming defendants of all races.

Sentencing Convicted Offenders

A larger literature adopts a more reduced form approach and examines racial differences in

sentences among either arraigned defendants or those who receive a guilty verdict. Cassia Spohn (2015) notes that, in the early twentieth century, there were well-documented cases where courts dealt with Black defendants in biased ways that almost certainly reflected racial animus on the part of court actors.³⁴ However, she also notes that many early regression studies of racial disparities in sentencing and case outcomes did not find large racial disparities conditional on controls for relevant case and defendant characteristics. Spohn (2015, 225) notes that the 1983 report by the National Research Council’s Panel on sentencing research concluded that racial disparities in sentencing outcomes were “due primarily to factors other than racial discrimination in sentencing.”

Jaeok Kim and André Kiesel (2018) use administrative data from New York City and New York state respectively to examine the relationships between distributions of types of arrests and distributions of sentencing outcomes while including controls for the criminal histories of defendants. They find that racial disparities in the nature of offenses and the extent of criminal histories at the time of arrest account for racial disparities in sentencing outcomes and that courts in New York City may actually produce smaller disparities in sentencing outcomes than researchers would expect given the racial disparities in arrest characteristics.

Greg Ridgeway and colleagues (2020) also study New York sentencing outcomes using data on convicted offenders. They find that Black convicted offenders were about 70 percent more likely to be sentenced to prison than White convicted offenders. However, given controls for criminal history and crime characteristics at the state level, this gap shrinks to roughly 7.5 percent. This study is one of many in the criminology literature that conclude that

among defendants with the same criminal history who face the same current charge, Black defendants may be less risky and deserve more lenient sentences.

33. Jordan (2024) examines plea bargaining in Cook County, Illinois. He finds that prosecutors exploit both Black and non-Black defendants’ aversion to the risk associated with going to trial. However, he finds no evidence that prosecutors bargain harder when dealing with Black defendants who face the same charges and share comparable criminal histories.

34. See chapter 2, “Executions,” in NASEM (2022).

most of the Black-White gap in incarceration sentences given conviction can be attributed to Black-White differences in the distribution of case and defendant characteristics.

However, Ridgeway and colleagues (2020) also find that conditional racial disparities in sentencing varied significantly among counties within New York state. They find that Black defendants in one county faced a 36 percent higher chance of being sentenced to prison than comparable White defendants. Low average rates of racial disparities in sentencing at a state or national level do not rule out the possibility that courts in certain jurisdictions sentence comparable defendants of different races quite differently. On the other hand, the regression methods that researchers employ to measure conditional disparities may be less reliable when applied to data from one jurisdiction. Ridgeway and colleagues (2020) note that, in some counties, Black defendants with long criminal histories who face serious charges have few or no peers among the population of White defendants. In these scenarios, researchers cannot know whether courts treat comparable defendants the same regardless of race because some types of defendants exist only within one race group.

Chad Topaz and colleagues (2023) employ data from federal courts to examine racial disparities in the length of prison sentences. Racial differences in case characteristics account for almost 90 percent of the overall nineteen-month difference in sentence length between Black and White defendants, but once again, the sizes of the conditional racial disparities in sentence length vary significantly among district courts. In some districts, these conditional disparities are small and statistically insignificant. In others, they are quite noteworthy.

Sentencing Rules and Guidelines

Since racial disparities in sentencing outcomes shrink dramatically when researchers add controls for observed case and defendant characteristics, it is possible that, given more complete sets of controls, these disparities would be negligible. However, the key point is the following: given any reasonable conjecture about

the exact magnitude of these disparities given X^* , structural factors likely remain the key drivers of overall racial disparities in criminal justice outcomes. In this section, we discuss the structural role of sentencing laws and guidelines in creating absolute racial disparities in the punishments that convicted offenders receive. In the next section, we consider structural flaws in the design of parole and probation supervision.

We begin with some observations about the history of sentencing policies. Derek Neal and Armin Rick (forthcoming) note that from 1925 to 1975, the US incarceration rate remained roughly constant at around one hundred per one hundred thousand.³⁵ Over this period, criminal law maintained an indeterminate sentencing regime that gave judges considerable discretion when sentencing convicted offenders, and court actors viewed offender rehabilitation as an important goal of the criminal justice system. Just as judges enjoyed great discretion when sentencing offenders, parole boards enjoyed great influence over the time that offenders who received a given sentence actually spent in prison. Parole boards saw their discretion as a tool for encouraging good behavior and rehabilitation. However, rising crimes rates and growing doubts that existing practices worked as tools for promoting rehabilitation sparked a change in the laws that governed sentencing and parole, and over the period 1975 to 2007, the incarceration rate increased by more than a factor of five. The rate fell steadily through the 2010s, but in 2019, it remained more than four times greater than the rate of one hundred per one hundred thousand that prevailed for roughly half of the twentieth century.

Here, we stress several points about the era of mass incarceration that has existed since roughly 1985, when the incarceration rate per one hundred persons first hit two hundred. To begin, using a variety of methods, Steven Raphael and Michael Stoll (2013) and Neal and Rick (2016) demonstrate that broad changes in sentencing laws drove growth in state prison populations. For example, Neal and Rick (2016) employ data from the National Corrections Re-

35. Reliable data on national incarceration levels do not exist prior to 1925.

porting Program (NCRP) to estimate the likelihood that offenders arrested for specific offenses in 1985 would serve prison terms of various lengths. They use the results to simulate what the incarceration rate would have been in 2005 if these probabilities had remained fixed among subsequent cohorts of arrested offenders, and they conclude that the incarceration rate would have grown by roughly 25 to 40 percent. However, the actual incarceration rate grew by almost 145 percent. They conclude that at least 71 percent and as much as 83 percent of the observed increase in incarceration rates between 1985 and 2005 must be attributed to changes in how courts punished persons arrested for particular offenses.

Raphael and Stoll (2013) employ data from 1984 and 2004. They conduct a similar simulation exercise, but they require the prison system to be in steady state in both years. They conclude that 91 percent of the growth in incarceration rates between these two dates must be attributed to changes in how courts deal with arrested offenders.

Both Neal and Rick (2016) and Raphael and Stoll (2013) conclude that a broad-based shift toward more punitive sentencing laws drove these changes in expected time served in prison given arrest. Both review numerous changes in sentencing rules and guidelines at the state level that likely played a role. Further, Neal and Rick (2016) show that increases in the likelihood that an offender arrested for a particular crime would serve prison time are evident in all crime categories. Overall, the likelihood of serving time in prison given arrest increased by 84 percent, and these rates more than doubled for drug crimes and various forms of theft. The likelihood of serving prison time given arrest for a violent crime rose less sharply, but here the expected length of prison spells given incarceration rose sharply. In all violent crime categories, the likelihood of serving at least five years in prison given arrest grew by at least 100 percent.³⁶

Both studies find no evidence that the increase in sentencing severity that drove the state prison boom involved a differential increase in severity within the crime categories where minority defendants are most overrepresented among those arrested. However, data on federal prison incarceration rates over the same period tell quite a different story. Here, rising incarceration rates for drug crimes, weapons offenses, and immigration violations drove federal prison growth between 1980 and 2010, and these trends disproportionately impacted incarceration rates for minorities. Alexander (2011) explains how the Anti-Drug Abuse Act of 1986 contained sentencing rules for offenses involving crack cocaine that were much more severe than the rules for offenses involving powder forms of cocaine. This policy appeared to target Black drug offenders directly, and the population of Black inmates in federal prisons grew much faster than the population of White inmates.

Michael Light (2022) notes that average difference in the length of federal sentences assigned to Black versus White defendants declined sharply between 2009 and 2018, and he discusses several factors that likely contributed to this development. The Fair Sentencing Act of 2010 made it more difficult to apply mandatory minimum sentences in crack cocaine cases. Further, changes in federal law enforcement strategy likely contributed to the large reduction in the annual number of crack cocaine cases filed over this period. However, Light argues that changes in how judges sentenced Black versus White offenders conditional on their criminal history, offense severity, and relevant mandatory minimum sentencing requirements played almost no role in the dramatic decline in the Black-White sentencing gap. Regarding federal incarceration rates among Black citizens, both the disproportionate increase during the prison boom and the subsequent excess decline during the 2010s appear to have been driven by changes in sentenc-

36. John Pfaff (2017) argues that prosecutors drove these trends by pursuing convictions more aggressively. Neal and Rick (forthcoming) reanalyze the data Pfaff (2017) cites and conclude that changes in prosecutor behavior played, at most, a minor role in driving these trends. Piecing together information from different data sources, Neal and Rick (forthcoming) show that expected time served given felony conviction rose sharply during the prison boom period.

ing rules and guidelines and not by changes in the behavior of judges conditional on existing policies.³⁷

Incarceration and Public Safety

We have argued consistently that laws and practices that are excessively punitive are a form of structural racism. Black males are arrested at more than five times the rate that White males are arrested. So, when courts produce approximately color-blind outcomes that are nonetheless too punitive—that is, the social costs of punishing defendants exceed the expected value of any resulting improvements in public order and safety—the laws and practices that govern court operations represent a form of structural racism. Likewise, when courts produce outcomes that are too lenient, Black communities often suffer disparate harm because victimization rates have typically been higher in Black communities, and much crime is within race.³⁸

The prison boom involved a sharp increase in the severity of sentencing rules for both state and federal courts, but to support our contention that this increase was excessive, we need to say something about the public-safety benefits created by this shift in policy. In this section, we review the literature on the links between incarceration policies and public-safety outcomes. The sharp growth of prison populations during the 1980s, 1990s, and 2000s was a catalyst for research on the public-safety benefits of incarceration. Further, in some states, recent efforts to reduce prison populations have produced additional evidence about the relationship between incarceration policies and public safety.

We begin by discussing recent evidence concerning the relationships between incarceration sentences and future offending rates for individual offenders. We then discuss work on the relationships between crime rates and the willingness of society to use incarceration as a

sanction for offenders. Studies that provide detailed information about the impacts of incarceration on the behaviors of individual defendants may tell us little about society-level relationships between incarceration rates and crime rates. Criminals who are not incarcerated may increase their criminal activities when other criminals are incarcerated, or potential criminals may be deterred from starting a criminal career when they see existing criminals sentenced to long spells in prison. Further, the treatment impacts of incarceration among the population of currently incarcerated persons may be quite different than the distribution of impacts that we would observe if courts drastically increased or decreased their use of incarceration sentences.

Individual Treatment Impacts of Incarceration

Over the past decade or more, several studies have produced estimates of the impacts of incarceration on future offending among individuals who are marginal candidates for sentences that require time in prison as opposed to community supervision. These studies do not address replacement effects—that is, the possibility that, when the state incapacitates one offender, different offenders may increase their criminal activity. Further, these studies do not address general deterrence effects—that is, the possibility that harsh punishment of existing defendants serves as a deterrent to future offending by others.

Before reviewing empirical results from these studies, we need to define terms. Here, we follow the framework presented in Jordan and colleagues (2023). When the state incarcerates an offender, two impacts on public safety are immediate. First, the offender cannot commit new offenses while in prison. Second, the offender ages while incarcerated, and much evidence suggests that aging directly lowers rates of offending. The first effect is a direct incapacitation effect. The second is an indirect

37. It is logically possible that throughout this period, federal judges or prosecutors have been biased against minority defendants. However, the initial rise and subsequent decline of racial disparities in federal incarceration rates appear to be driven by changes in policy.

38. See chapter 2 of NASEM (2022). Hispanics and Native Americans also face higher arrest rates and victimization rates than Whites.

but still unavoidable effect of incapacitating an offender. In addition to these two incapacitation effects, the experience of prison may impact the age-specific offending rates of inmates after their release from prison. Whether time in prison increases or decreases these age-specific offending rates may hinge on how prisons operate, for example, the quality of programs that offer skills training, substance abuse counseling, behavioral therapy, and so forth. Thus, the impacts of prison experiences on post-release, age-constant offending rates are not unavoidable effects linked to incapacitation per se but rather impacts linked to the specific experiences of inmates during incarceration. If these experiences increase age-specific offending rates, time served in prison is criminogenic. If these experiences decrease age-specific offending rates, time served deters crime in the future.

To make these ideas more concrete, consider a group of comparable defendants who are all age a when sentenced. Assume that the court randomly assigns half of these defendants to serve $m = \bar{m}$ periods in prison and assigns the other half to probation, $m = 0$. Let τ be a random failure time that denotes the time of the first offense a defendant commits after sentencing, and define $F(t|a, m)$ as the probability that a defendant sentenced at age a to serve m years in prison commits a new offense during the first t periods after sentencing—that is, the probability that $\tau \leq t$ given a defendant sentenced to m periods in prison at age a .

Several things are immediate. To begin, $F(t|a, m) = 0$ for $t \leq m$, since we assume that inmates cannot offend while incarcerated. In addition, if we assume that the experience of prison does not change age-constant rates of recidivism, the full impact of $m = \bar{m}$ of incapacitation at horizon t is an absolute reduction in recidivism rates of periods

$$\text{incap}(t, a, \bar{m}) = F(t|a, 0) - F(t - \bar{m}|a + \bar{m}, 0)$$

Put differently, if the experience of prison impacts post-release offending rates only because prisoners age while incarcerated, then prison reduces recidivism rates by reducing the time that defendants are exposed to recidivism

risk from t to $t - \bar{m}$ and by shifting initial exposure to age $a + \bar{m}$ instead of a . The magnitude of both effects is determined by time-served, \bar{m} .

Jordan and colleagues (2023) present estimates of $F(t|a, 0)$ for the complier set—that is, those who may or may not receive an incarceration sentence depending on the severity of the judge assigned to their case, and they find that $F(t|a, 0)$ grows little with t over horizons greater than forty-eight months. They also find small impacts of a year of aging on future recidivism rates. Taken together, these results imply that, over horizons of five years and beyond, the direct impacts of short prison spells on recidivism rates should also be small. Reducing exposure to recidivism risk from five years to four years matters little, and one year of aging has a minor impact on recidivism rates. Since most defendants sentenced to prison serve less than two years and many serve only a few months, there is no reason to expect large direct impacts of incarceration on recidivism at long horizons.

However, both Jordan and colleagues (2023) and Manudeep Bhuller and colleagues (2020) do find lasting impacts of incarceration treatment on recidivism. Bhuller and colleagues (2020) examine data from Norway and find that incarceration treatment reduces recidivism rates by roughly 29 percentage points five years after sentencing. Jordan and colleagues (2023) examine data from Cook County, Illinois. They find that incarceration reduces five-year recidivism rates by just over 29 percentage points among defendants sentenced following their first felony charge. However, when they examine repeat offenders, they find no impact of incarceration on five-year recidivism rates. The Bhuller and colleagues (2020) results for all offenders and the Jordan and colleagues (2023) results for first offenders are difficult to reconcile with the view that prison simply reduces exposure to recidivism risk and shifts recidivism risk to older and less risky ages. These two effects are not large enough to generate large reductions in recidivism rates at long horizons.

Since repeat offenders are former first offenders who were not deterred by the punishment they received in previous cases, no one should be surprised to learn that incarceration impacts recidivism differently among first of-

fenders versus repeat offenders.³⁹ Nonetheless, in the US, most defendants sentenced to prison are repeat offenders.

Further, Bhuller and colleagues (2020) argue that the positive impacts of prison on recidivism rates in Norway are likely the result of effective training programs that may have few parallels in US prisons.

Taken as a whole, the existing literature suggests that, in the US context, most incarceration sentences incapacitate offenders and do little else. Often researchers define recidivism events as new arrests, new convictions, or new incarcerations. However, most crimes do not result in an arrest. So, some researchers have attempted to measure the incapacitation channel through inmate interviews that ask defendants about their history of past offenses. Reviews of this research, much of it conducted during the 1970s, prior to the surge in US incarceration rates, suggest the typical inmate would have committed ten to twenty serious felony offenses in a given calendar year, if he had not been incarcerated and managed to avoid arrest.⁴⁰

These results describe inmates incarcerated during a period when US incarceration rates were quite low, and the stock of prison inmates likely contained a larger fraction of career criminals. Further, Raphael and Stoll (2013) reports that the average age of prison inmates was lower in this earlier period. Both these factors suggest that the benefits of incapacitating the representative inmate fell as prison populations grew during the 1980s, 1990s, and 2000s. In a more recent analysis of the criminal activity of adults convicted of felonies, Emily Owens (2009) finds that the crime-reduction value of incapacitating an additional inmate for one year is roughly one-fifth of the reduction implied by much of the earlier research.

Equilibrium Impacts of Incarceration

So far, we have reviewed the literature that addresses how incarceration impacts individual

recidivism. However, most policymakers are interested in the relationship between the propensity of courts to incarcerate offenders and the overall level of crime in society. As we note earlier, estimates of the individual treatment impacts of incarceration may not tell us much about how policies that drive the use of incarceration sentences impact overall crime rates because spillovers may exist. If society incarcerates one offender, this may create new criminal opportunities for other offenders, or it may deter others by making the prospect of punishment more salient.

Empirical work on this topic is challenging because incarceration rates in a given society may be high precisely because crime rates are high. Yet, a number of papers have exploited large-scale policy experiments to isolate potentially exogenous changes in incarceration rates that may not be correlated with changes in the underlying propensity of citizens to commit crime.

Several studies exploit the unusual Italian practice of periodic, large, and sudden prisoner releases through collective clemencies and collective pardons. Alessandro Barbarino and Giovanni Mastrobuoni (2014) construct a panel dataset of crime and incarceration rates that vary by year and by Italian province and exploit province-level variation in pardon totals for all prisoner releases occurring between 1962 and 1995. The authors find sizable impacts of incarcerating defendants on crime. These impacts are of the same order of magnitude implied by the results of many US inmate surveys in the 1970s.⁴¹

Paolo Buonanno and Steven Raphael (2013) use relatively high-frequency crime and incarceration data at the national level as well as province-level variation to estimate the effects on crime of the August 2006 Italian mass prisoner release. The authors find that incarcerating a felon eliminates thirteen to seventeen serious offenses per year served. However, they also find that, pre-pardon crime rates held con-

39. Bhuller and colleagues (2020) work with smaller samples and do not have the power to precisely estimate separate treatment impacts for first versus repeat offenders. However, point estimates in their appendix results indicate much larger absolute impacts of incarceration on recidivism among first offenders.

40. See Marvell and Moody 1994; Spelman 1993, 2000.

41. See Marvell and Moody 1994; Spelman 1993, 2000.

stant, the reduction in offenses associated with incarcerating an offender for a year is much larger in provinces with smaller pre-pardon incarceration rates.

Ben Vollaard (2013) analyzes the impact of a sentence enhancement in the Netherlands that targeted repeat offenders, defined as those with over ten prior felony convictions. In 2001, the Netherlands enacted an enhanced sentence of two years for such offenders, first allowing a small number of municipalities to experiment with the enhancement before nationwide application in 2004. The author finds that each year of time served prevented between fifty to sixty reported thefts. However, municipalities could only assign enhancements that the central government allocated to them, and the author finds significantly smaller crime reductions per additional prison year served in municipalities that were allocated enhancements for a larger share of their eligible defendants.

The overall crime reduction benefits of incarceration in the Dutch and Italian studies are comparable to estimates derived from the surveys of US prison inmates during the 1970s. However, these studies also find that inmates differ greatly in their propensity to commit crime if not incarcerated, and this heterogeneity implies that, to the extent courts target incarceration sentences to more risky offenders, we should see diminishing returns from sentencing reforms that successively mandate incarceration for larger and larger groups of defendants.

There are several studies of the crime-prison relationship based on US panel data regressions. Using US state panel data from 1972 to 1992, Steven Levitt (1996) exploits the fact that in years when states are under a court order to relieve prisoner overcrowding, state prison populations grow at relatively low rates. Using measures of the status of prisoner overcrowding lawsuits as instruments for changes in state-level incarceration rates, Levitt finds

crime-prison elasticities that are considerably larger than comparable estimates from linear regression models. He reports a property crime-prison elasticity of -0.3 and a violent crime-prison elasticity of -0.4 .

Rucker Johnson and Steven Raphael (2012) construct an instrument for state-level incarceration rates according to the difference between a state's current incarceration rate and the state's steady-state incarceration rate implied by observable admissions and release rates. The authors justify the instrument using a theoretical model that spells out conditions under which the transitory disparity between the actual and steady-state incarceration rate provides a valid instrument for one-year lead changes in the actual incarceration rate. The authors analyze state-level panel data for two time periods: 1978 to 1990 and 1991 to 2004. The former period is characterized by a relatively low incarceration rate (186 per 100,000) while the latter period is characterized by a much a higher incarceration rate (396 per 100,000).

For the early period, an additional prison year served prevents roughly 2.5 felony violent offenses and 11.4 felony property offenses. Note that the sum of these impacts is quite close to the implied impact of the 2006 Italian pardon program on total serious offenses. In 1978, the California incarceration rate was comparable to the rate in Italy prior to the pardon program that Buonanno and Raphael (2013) examine, and while the incarceration rate rose in California during much of the first period, the average rate during this period was less than half of the prevailing rate during the latter period.⁴² Yet, during the 1991–2004 period, the California incarceration rate was often close to five hundred per one hundred thousand,⁴³ and the implied reductions in crime generated by marginal increases in California incarceration rates were much smaller. Here, Johnson and Raphael (2012) estimate that crimes prevented per prison year served were 0.3 for violent felony offenses and 2.7 for felony property offenses.

42. When expressed in terms of elasticities, the results for California during the 1978 to 1990 period are also quite close to the estimates in Levitt (1996), who examined data from several US states in the period 1972 to 1993.

43. See "California's Prison and Jail Incarceration Rates," Prison Policy Initiative, April 2024, https://www.prisonpolicy.org/graphs/jails2024/CA_incarceration_rates_1978-2022.html.

Raphael and Stoll (2013) reproduce this analysis with updated data for three time periods: 1977 through 1988, 1989 through 1999, and 2000 through 2010, with corresponding weighted-average state incarceration rates of 171, 349, and 449. They find small prison-crime effects for the latter two time periods, but much larger effects for the earliest time period. Raymond Liedka and colleagues (2006) provide similar evidence with US panel data.

In sum, results from the Dutch and Italian studies are comparable to those from US panel data studies that take data from time periods when incarceration rates in the US were more comparable to those in the Netherlands and Italy. However, results for US data gathered after the prison boom indicate that marginal increases in incarceration rates produce, at best, small reductions in crime.

The most striking evidence that public-safety gains diminish at the margin as prison populations grow comes from California. California was among the first states to adopt punitive sentencing rules that produced rapid growth in prison populations.⁴⁴ However, an adverse ruling from the US Supreme Court forced them to reduce prison overcrowding, and in April 2011, California adopted reforms that sought to reduce prison populations. The legislation eliminated the practice of returning parolees to state prison custody for technical parole violation for all but a small set of the most serious offenders. The legislation also defined a group of nonserious, nonsexual, non-violent offenders who upon conviction serve their sentences in county jails. The reforms generated almost no change in crime rates even though they caused a sharp drop in the state's prison incarceration rate with a small and far from offsetting increase in the state's jail population. The net effect was a sizable reduction in the incarcerated population that was comparable in magnitude to the reduction generated by the Italian collective clemency of 2006. In fact, on a per capita basis, the shock caused by California realignment was a little larger than

the effect of the 2006 clemency in Italy. The fact that the 2011 reform in California did not impact crime rates in the way that the 2006 program did in Italy is consistent with the hypothesis that the crime-reduction benefits associated with policies that increase incarceration rates are diminishing in the baseline incarceration rate. The key difference between the two scenarios is the much larger baseline incarceration rate in California.

Some readers may conjecture that society benefits long term from establishing policies that make the prospect of incarceration a salient deterrent for most criminals. Further, if policymakers must enforce punitive sentencing rules for many years before potential offenders revise their beliefs about the expected consequences on future convictions, researchers may find it difficult to measure these benefits. Many researchers have examined the importance of such general deterrence effects, and several high-quality literature reviews summarize the evidence from their studies.⁴⁵ There is little evidence that increases in the severity of punishment deter criminal activity. Criminal behavior among potential offenders is more responsive to their beliefs concerning the likelihood that a crime they commit will be detected, prosecuted, and punished in some manner. Increasing the severity of punishment (for example, increasing the length of a potential prison sentence for an offense that already carries a prison sentence) does not appear to increase deterrence. However, factors that increase the likelihood of punishment—for example, higher police staffing levels as demonstrated in Aaron Chalfin and Justin McCrary (2018) or the swiftness of a sanction (see the evaluation of Hawaii's HOPE Pretrial program in Janet Davidson and colleagues (2019)—do create general deterrence. These results make sense if persons who commit offenses serious enough to draw the attention of the police, prosecutors, and courts are also persons who heavily discount the future.⁴⁶

Neal and Rick (forthcoming) provide a more

44. See Raphael and Stoll 2013 and Neal and Rick 2016 for details.

45. See Chalfin and McCrary 2017 and Nagin 2013.

46. See Mastrobuoni and Rivers 2016 for evidence of present orientation among offenders.

detailed treatment of much of the evidence we review here, and they conclude that changes in sentencing policies account for most of the sharp rise in incarceration rates that began in the late 1970s and for the fact that incarceration rates remain far above any level that most scholars or policymakers could have envisioned fifty years ago. These policy changes disparately impacted minority citizens, who face arrest much more often than Whites, and there is considerable evidence that the continued push for ever more punitive sentencing rules during the 1990s and early 2000s generated, at best, marginal improvements to public order and public safety. In the section “Simulating the Impacts of Reforms,” we assess the prospects for reducing racial disparities in incarceration rates by making sentencing laws less punitive for offenders who are neither violent offenders or recent recidivists.

COMMUNITY SUPERVISION

Courts assign sentences to convicted offenders, but parole and probation authorities often make decisions that determine the consequences of these sentences for defendants. The combined populations of persons serving nominal incarceration sentences under either probation or parole supervision are larger than the population of persons in prison. In 2021, roughly 1.8 million people were incarcerated. Most prisoners were in a state or federal prison, but a sizable portion were in jail. At the same time, roughly 3.75 million people were under community supervision. Of these, 803,000 were on parole while approximately three million were under probation supervision.⁴⁷

Further, there are enormous racial disparities in community corrections supervision rates. While Black people constitute 13.6 percent of the US resident population, they account for 30 percent of those on probation and 37 percent of those on parole (Kaeble 2023). Expressed as a rate per 100,000 adults, the community supervision rate is 3,749 per 100,000 for Black adults. Yet, the rates are 1,193 per 100,000 for Hispanic adults and 1,211 per 100,000 for White adults.

For much of US history, most prison sentences were indeterminate. Judges sentenced offenders to a range of possible times served, and *ex post*, parole boards decided when and if (in the case of indeterminate life sentences) to release individual prisoners. Beginning in the 1970s, many states began to add more determinacy to sentencing. This took many forms. Some states abolished parole boards together and introduced determinate sentencing systems where a felony conviction results in a fixed sentence and time served is ultimately determined by the sentence length minus fixed credits for pretrial detention and time off for good behavior. Other states introduced a series of mandatory-minimum sentences, sentence enhancements, and truth-in-sentencing provisions that constrained the release authority of parole boards until mandatory minimums are served.

While there are currently fourteen states that scholars commonly label as determinate sentencing states (Reitz 2019), the actual degree of determinacy in these states is best characterized as being on a continuum, with some states leaving little discretion to parole boards and other states still affording considerable discretion to parole authorities in determining ultimate time served. Despite this heterogeneity, state sentencing is more structured today than in the past, and discretionary release authority has been curtailed.

Whether states grant parole boards much or little control over the time that prisoners serve given their sentence, when prisoners do leave prison, they typically remain under the supervision of parole officers for a period of time. Since prisoners rarely serve their entire sentence before being released, persons under parole supervision are legally inmates who are serving the balance of their incarceration sentence under community supervision. Likewise, the vast majority of convicted felons who do not receive incarceration sentences receive probation sentences, but again, probationers are legally persons who are serving incarceration sentences under community supervision rather than in prison.

47. The community correction supervision rate has dropped considerably over the past fifteen years, from a peak of 2,240 per 100,000 US adults in 2007 to the rate of 1,440 per 100,000 in 2021 (Carson and Kluckow 2023).

The fact that, in most cases, persons under parole or probation supervision are legally persons who are serving alternative incarceration sentences is crucial. For this reason, persons under community supervision do not possess many rights that ordinary citizens enjoy, and these lost rights may also contribute to racial disparities in the experiences of defendants who receive the same nominal sentence from a court. For example, persons under community supervision are not protected from warrantless searches by police or parole officers, and there is evidence that, at least in some jurisdictions, law enforcement officers disproportionately ask Black drivers whether they are on probation or parole, and if drivers answer in the affirmative, police are much more likely to conduct a vehicle search (Macdonald and Raphael 2020).

In addition, community corrections officers may impose numerous requirements on the persons under their supervision: for example, avoid people who are criminally active or with known criminal histories, avoid drugs and alcohol, meet regularly with a supervising officer in the field, find and maintain employment, do not leave one's county of residence without permission, wear an electronic monitor, participate in rehabilitative programming, participate in community service, and do not contact crime victims. Further, in many jurisdictions, probation and parole officers have the power to initiate proceedings that result in probation or parole revocation and a new prison admission by simply alleging that a person under their supervision violated one such condition of supervision. Most state prison systems record these events as admissions to prison for technical violations. Data from 2021 and 2022 indicate that, in each year, at least one hundred thousand persons on probation or parole faced revocation of their supervision and admission to a state prison. Further, more than half of these revocations were the result of technical violations rather than new crimes.⁴⁸

In this section, we review the existing literature on parole board hearings, parole supervision, and probation supervision. We note that some studies provide evidence that racial bias impacts the decisions made by parole boards and by community corrections officers, but our main concerns involve the design of community corrections jobs and the delegation of decision-making rights to community corrections officers. We know of no states where community corrections officers face strong incentives to increase the number of parolees or probationers under their supervision who secure stable employment or housing. However, media reporting on the criminal activities of parolees and probationers always creates public pressure to prevent recidivism among those under community supervision.

Given this pressure, community corrections officers face incentives to err on the side of caution when deciding whether to recommend that those who violate supervision rules face revocation.

These incentives matter because, as we noted earlier, many states grant probation and parole officers the power to initiate technical revocation proceedings for persons under their supervision, even if these persons have not committed a new crime. Further, critics argue that parole boards often rubber-stamp the judgments of community corrections officers who recommend technical revocations.⁴⁹ However, there is little evidence that community corrections officers employ technical revocations in ways that incapacitate offenders who pose particularly high risks to public safety, which raises concerns about the efficacy of granting probation and parole officers the power to initiate technical revocations.

Further, technical revocations have a disparate impact on Black citizens because Blacks are at least three times more likely to be under community supervision than Whites or Hispanics. While there is also some evidence that community corrections officers discriminate against Black persons under their supervision

48. See table 8 in Carson and Kluckow (2022) and CSG (2019).

49. See "Prisoners' Rights: Morales v. Monreal," Uptown People's Law Center, 2025, <https://www.uplccchicago.org/what-we-do/prison/morales-v-monreal.html>.

when making revocation recommendations, the key point is that existing policies generate many revocations that are difficult to justify in terms of their expected impacts on public safety, and Black citizens bear far more than their share of the costs created by these revocations. Thus, the laws that grant these powers to community corrections officers are likely a form of structural racism.

Here, we discuss three separate sets of results. We discuss the literature on parole board hearings that govern the initial release of some prisoners, and then we discuss technical revocations of parole that result in readmission to prison. We also discuss probation supervision and revocations of probation that generate new prison admissions.

Parole Boards

Parole supervision usually applies to people released from state prison.⁵⁰ For much of the twentieth century, prison sentences involved indeterminate sentences with minimum and maximum sentences set at trial, and parole boards determined effective time served, but beginning in the 1970s, many states adopted more determinate sentencing laws, and the vast majority adopted Truth in Sentencing statutes in response to the financial incentives provided by the 1994 Violent Crime Control and Law Enforcement Act.⁵¹

Nonetheless, Kevin Reitz and Edward Rhine (2020) note that despite the move toward deter-

minate sentences in sixteen particular states, parole board decisions still impact time served for nearly three quarters of those released from prison. Even when sentences are determinate, parole and prison authorities exercise discretion when awarding time-served credits for good behavior and program participation.⁵²

In a recent annual report, the California Committee on Revision of the Penal Code (CRPC 2021) analyzed the results of parole hearings occurring in the state during calendar years 2019 and 2020.⁵³ The state parole board controls the release date for a small fraction of inmates, but the report found no differences in the likelihood of being granted parole by race. Both the average grant rate and the average grant rate conditional on the number of recent rules violation reports did not vary by race.

There are few independent estimates of the effect of race on parole hearing outcomes. Beth Huebner and Timothy Bynum (2008) analyze parole release decisions by race and ethnicity for serious youthful offenders in an unnamed jurisdiction. The authors find that legally relevant factors, such as a parole guideline score and an indicator of institutional misconduct, predict the release decision. However, they also find that conditional on the controls they include in a hazard analysis, the time to release conditional on being eligible for parole is relatively longer for Black people and shorter for Hispanic people relative to White people. We should note, however, that the vector of con-

50. There is no parole in the federal prison system. Among the states, California adopted a 2011 reform that assigns people released from prison following convictions for less serious crimes to supervision by county probation departments. See Lofstrom and Raphael 2016.

51. See table 1 in Neal and Rick 2016.

52. Reitz and Rhine (2020) also note the enormous heterogeneity in practice across states. In some states, minimum sentences are set at a quarter or a third of the maximum, while in others, judges have the discretion to set both minimum and maximum sentence. In some states, good-time credits apply to the minimum eligible release date, effectively reducing the amount of time that one needs to serve prior to being considered for parole, while in other states, credits only apply to the maximum sentence.

53. California voters expanded the reach of the parole board to some determinate sentences through passage of Proposition 57 in 2016. Specifically, the proposition permits the early release of people convicted of less serious offenses once they have served the sentence for their primary offense. Hence, time associated with consecutive sentences, enhancements, and so forth can be disregarded subject to administrative review by the board. While the number of people eligible for release under Proposition 57 is small, the provision provides an interesting example of a complex source of indeterminacy in supposed determinate sentences, a general issue noted by Reitz and Rhine (2020).

trols used in the study is limited.⁵⁴ More granular controls might yield different findings. Again, the control set in this study may not be the full set of relevant controls. While the evidence pertaining to racial disparities in placement on community supervision is thin, stronger evidence indicates that community corrections supervision is unduly restrictive, revokes many people on probation and parole who would likely not have re-offended, and disparately impacts Black people. We turn to this evidence now.

Parole Supervision

As we have noted, in states where inmates are released from prison to parole supervision, parolees often face significant risk of returning to prison even if they commit no new crimes. In many states, parole officers frequently initiate proceedings that result in former inmates returning to prison because they have violated a rule that governs the conditions of their supervision. In 2021, roughly 27 percent of admissions to state prisons involved people on community supervision, and in many cases, these admissions were punishments for technical violations of supervision conditions.⁵⁵

The oversized impact of technical violations in generating prison admissions is readily observable in data on post-release returns to custody among people released from the Illinois Department of Corrections (IDOC). Persons released from IDOC are subjected to mandatory supervision terms (referred to as Mandatory Supervised Release [MSR]) of one, two, or three years, with lengthier terms associated with more serious offenses.⁵⁶ Like those in other parole systems across the country, people on MSR must check in with their parole officer and obey all predetermined restrictions and orders of

the MSR officer. Failure to comply and a recommendation to revoke by the MSR officer results in a hearing before the Prisoner Review Board that is external to the criminal court system in Illinois.

Figure 2 presents sixty-day moving average estimates of the proportion of people on MSR supervision in Cook County, Illinois, who experience their first return to prison as a function of time since release from prison. The data come from the period 1990–2010. The figure follows releasees for five years and provides separate estimates for persons with terms of one, two, and three years.⁵⁷

The first line in each pair of vertical lines marks one year, two years, and three years from the date of prison release. The second line in each pair is sixty days later. Beyond each of these dates, one of the estimated hazard functions is no longer impacted by supervision-induced revocations.

There are several noteworthy patterns in figure 2. First, while noisy, the hazard rates are similar over much of the first year for those assigned to one versus two years of MSR, and these two categories account for most of the sample. Second, once MSR supervision ends, we see marked declines in the re-incarceration hazard. This occurs at one year for those with a one-year MSR term (circles), at two years for those with two-year MSR terms (diamonds), and at three years for those with three-year terms (squares). Finally, as each group is released from their respective MSR terms, the re-incarceration hazard aligns closely with those of groups whose terms previously expired after shorter periods of supervision.

Individuals who commit and are convicted of new felony offenses will be punished for the new offense regardless of whether they are un-

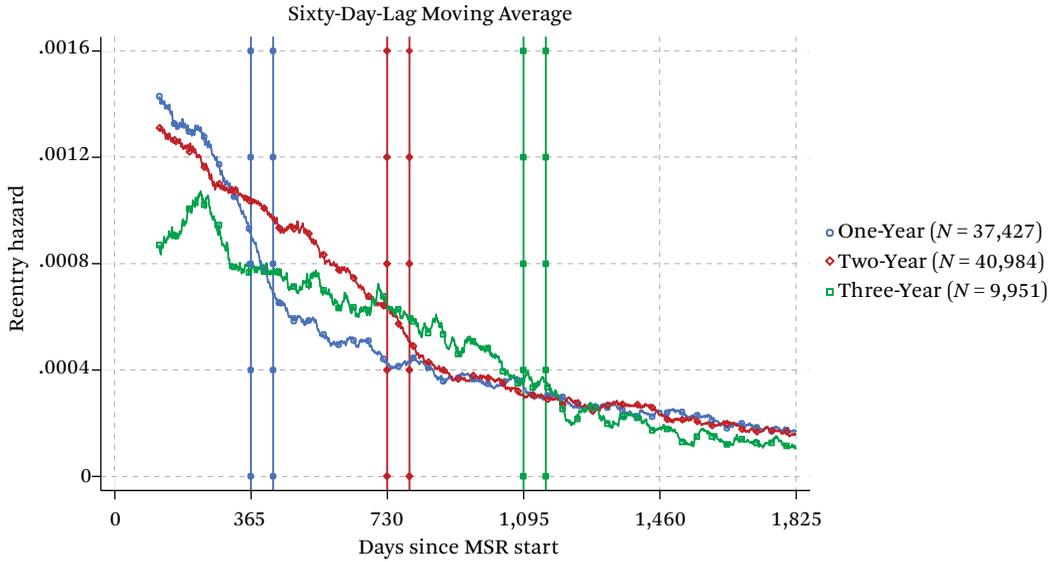
54. The controls are educational attainment, employment status at time of offense, indicator for mental health problems, gang membership, whether the offense was a serious person offense, a drug offense, an institutional misconduct indicator, the parole score, number of prior convictions, and time served.

55. See Carson 2022. Catalina Franco and colleagues (2024) point out that technical parole revocations implicitly increase the expected time served associated with any nominal prison sentence.

56. Recent reforms in Illinois have introduced more than three possible MSR terms, but the data we present come from the pre-reform period.

57. Whether released prisoners serve one, two, or three years on MSR is determined by the class of the offense that generated the conviction.

Figure 2. Empirical Re-incarceration Hazards by Mandatory Supervised Release Spell Length



Source: Clerk of Court of Cook County, Illinois; Illinois Department of Corrections.

Note: Each line presents a sixty-day moving average of the daily re-incarceration hazard for offenders recently released from state prison in Illinois. The individuals in the sample are persons sentenced to prison between 1990 and 2010 in Chicago, Illinois. Circles indicate former inmates given one year of Mandatory Supervised Release (MSR). Diamonds indicate former inmates given two years of MSR. Squares indicate former inmates given three years of MSR. The first vertical line of each type marks the end of the MSR period. The second vertical line marks sixty days after the end of MSR.

der community corrections supervision, and readmissions to prison associated with new criminal charges vary little around these cut-offs. Thus, the patterns in the figure are driven almost entirely by the fact that, when supervision ends, returns to prison for technical violations end.

Luke Brinkman and colleagues (2025) analyze more recent data on parole supervision in Illinois. In December 2022, Illinois reduced the standard MSR term for class 3 and 4 offenders from twelve months to six months. Brinkman and colleagues (2025) employ several different estimation strategies, and they find that reducing supervision terms by six months reduces the one-year rate of return to prison by roughly 7–10 percentage points or roughly one-third to one-half of the pre-reform level. This drop in returns to prison is driven by a drop in technical revocations. One year after release, the reform had no impact on the rate of revocations linked to new charges.

Brinkman and colleagues (2025) find similar results for both Black and non-Black persons who exit prison to MSR. The reform therefore had a disparate impact on Black citizens, who are overrepresented in the Illinois prison system.

Data on parole-revocation-hearing outcomes are not as widely available as data on sentencing outcomes in court cases, and we are unaware of other articles that investigate differential parole revocation rates by race while conditioning on individual case and person characteristics. However, two recent evaluations of reforms to parole practices in California and Kansas provide evidence that, historically, revocations rates have likely been inefficiently high, and Black citizens have been adversely and disparately impacted by the policies that produced these high revocation rates.

As we note earlier, California drastically revamped its community corrections practices in 2011. The reform greatly restricted the use of

technical violations as a tool for returning parolees to prison and transferred authority for people released on nonserious, nonviolent, nonsexual offenses to county probation departments.

Weekly prison admissions in the state fell from approximately 2,100 to 600, and within six months, the size of the state prison populations dropped by roughly 30 percent (Lofstrom and Raphael 2016).⁵⁸ The findings from several studies strongly indicate that, prior to this reform, California was using technical revocations to incapacitate parolees who did not pose significant risks to public safety. First, there is little evidence that the large reduction in the prison population generated by the decline in revocations increased crime rates. Magnus Lofstrom and Steven Raphael (2016) employ synthetic cohort analysis as well as analyses of cross-county variation in the impact of the reform within the state and find no evidence of an impact of the policy change on violent crime. They find evidence of a modest effect on property crime, but this is driven almost entirely by an increase in one crime category, motor vehicle theft.

Second, in an analysis of monthly release cohorts from the state's prison system, Lofstrom and colleagues (2014) show that, even though parolees experienced a 30-percentage-point decline in the likelihood of being returned to prison custody within one year of the implementation of realignment, the re-arrest rate for parolees did not change. Further, the rate of conviction given rearrest rose by only 3.1 percentage points. These findings suggest that, in the pre-alignment period, parole officials were not using technical revocations to reincarcerate parolees who were serious threats to public safety.

These results are important because the decline in the state prison population had a disparate impact on incarceration rates by race. Lofstrom and colleagues (2020) analyze the effects of a constellation of California reforms oc-

curing through 2017 on racial disparities in arrest and incarceration rates in the state. As in all states in the US, Black people in California are much more likely to be stopped by the police, to be arrested, to be under community corrections supervision, and to be incarcerated. Lofstrom and colleagues (2020) use data from the American Community Survey to estimate the proportion of California adults eighteen to fifty-five years of age who are residing in non-institutional group quarters when surveyed. In 2011, the percent of Black, Hispanic, and White men institutionalized in the state was 9.3, 2.4 and 1.5 percent, respectively. By 2014 the comparable figures declined to 7.8, 2.2, and 1.3 percent. Hence, while the percent institutionalized fell among all groups, the decline was greatest among Black men, and the Black-White gap fell from 7.8 to 6.5 percentage points.

Ryan Sakoda (2023) analyzes several reforms to the use of technical revocation for people under supervision in Kansas. In 2000, the state enacted legislation that eliminated supervision for people who were released from prison after entering prison because they violated a technical condition of probation supervision. In Kansas, the new law describes this change as eliminating probation supervision. However, for our purposes, the original Kansas program is a form of parole supervision because it governs the supervision of people after their release from prison.

The new law applied to a subset of offenders convicted of less serious offenses, and Sakoda (2023) treats the new law as a natural experiment that defines a treatment group and a control group convicted of similar crimes not included in the legislation. The author documents a 28.5-percentage-point decline in the rate of returning to custody against a baseline rate of 35 percentage points. Further, among the treated, there is no measurable increase in convictions for new crimes that result in a return to prison. Finally, Sakoda (2023) demonstrates that, among the treated, the change

58. The steady-state decline in the prison population of over 27,000 people was partially offset by an increase in the state's jail population (or roughly 8,000) driven by revocations to jail of former parolees (yet for shorter periods) and an aspect of the reform that made lower level felonies punishable with a jail spell rather than a prison sentence. Regardless, the reform caused a sharp and sustained drop in the combined prison and jail population.

in policy eliminates the 17-percentage-point Black-White disparity in the likelihood of being returned to prison that existed prior to the reform.⁵⁹

Probation Supervision

As we have noted, courts legally sentence many convicted felons to incarceration but do not require them to enter prison. Probation sentences allow convicted offenders to technically serve prison sentences in their communities under the supervision of a probation officer. If they refrain from committing new crimes and comply with all conditions of their supervision, they earn release from community supervision without ever entering prison. However, as in the case of parole supervision, many states grant probation officers the discretion to initiate proceedings that can produce probation revocations and prison admissions simply because a probationer violated a technical condition of supervision. Further, the best evidence available suggests that granting this power to probation officers is just as problematic as granting similar powers to parole officers.

Probation supervision is often applied to people convicted of misdemeanors or less serious felonies. In misdemeanor cases, failure to comply with the conditions of supervision may result in jail time. In felony cases, revocations usually result in time served in a state prison. In some jurisdictions, a period of probation supervision is appended to short jail spells, again with the threat of being returned to custody for failure to comply (Phelps 2020). Many jurisdictions also use systems of graduated sanctions and various forms of swift-and-certain compliance schemes in an attempt to avoid the imposition of lengthy prison terms.⁶⁰

Evan Rose (2021) analyzes a 2011 North Carolina reform that reduced the punishment for certain technical violations of probation, such as failure to pay fees or failing drug tests, from

revocation to prison to short jail stays and other non-incarceration sanctions. Using people on supervised probation as the treatment group and those on unsupervised probation as the control group, Rose first documents a 5.3-percentage-point relative decline in the likelihood of being revoked among those whose supervision conditions were altered by the policy. Among the treated, overall arrests increased by roughly 2 percentage points.

Among Black people on supervised probation, the decline in revocation rates is twice as large, while the increase in rearrest rates is only slightly larger than 2 percent. Thus, the new policy eliminated the Black-White gap in revocation rates while having almost no impact on the Black-White gap in rearrest rates.

Rose notes that under specific, testable assumptions concerning behavioral responses to the law change, it is possible to estimate the expected recidivism rate among persons whom probation officers would have revoked in the pre-reform period but are not allowed to revoke after the reform. Rose calls this rate a measure of probation officer accuracy. If it equals one, then probation officers are using their discretion to revoke persons who are certain to commit a new crime in the near future unless they are returned to prison. If it equals zero, probation officers are revoking persons who pose no threat to public order or safety.

While probation officers initiate technical revocations for persons with higher propensities to offend, probation officers do not possess crystal balls. Further, their accuracy rate is discreetly lower among Black probationers relative to non-Black probationers, roughly .33 versus .55 respectively. In addition, as a tag for recidivists, revocation due to technical violations fails to identify roughly 93 percent of Black and non-Black probationers who re-offend, yet falsely classifies 9 percent of Black compliers and 2.5 percent of non-Black compliers as re-

59. Sakoda's findings are consistent with experimental evidence demonstrating that more intensive supervision for high-risk probationers appears to increase incarceration and the likelihood of absconding without reducing the probability of conviction for a new offense (Hyatt and Barnes 2017). Similar experimental evidence suggests that reduced-intensity supervision among low-risk probationers does not increase the likelihood of a new offense (Barnes et al. 2012).

60. See Davidson and colleagues (2019) for a review of this literature and an application of this model to pretrial supervision.

cidivists. Finally, Rose (2021) also reports that, due to racial differences in revocation accuracy by race, the technical revocations applied to Black offenders in the pre-reform period are much less likely to have generated reductions in crime that justify the resulting increased costs of incarceration.

Rose (2021) also notes that his results do not necessarily imply that probation officers harbor racial animus against Black persons on probation. Probation officers may follow the same rules of thumb for recommending revocations regardless of the race of the person under their supervision, but members of different race groups may find it more costly, on average, to comply with the standard conditions of supervision, such as paying fines in a timely manner. If failure to comply with such conditions is more a signal of economic hardship than criminality, the Rose (2020) results may be expected.⁶¹

A Common Thread

Collectively, the results from reform efforts in Illinois, California, Kansas, and North Carolina indicate that parole and probation supervision are blunt instruments for controlling crime. There is some evidence that technical revocations reduce crime, but the overall correlation between recidivism risk and the incidence of revocation is weak. In all four states, new laws restricted the ability of officials involved in community corrections to incarcerate those under their supervision without going through the courts, and in all four instances, prison populations shrank. Further, in three of four cases, the population of Black prisoners shrank proportionally more than the population of non-Black prisoners.

The Rose (2021) results do provide direct evidence that probation officers are more likely to falsely tag Black probationers. Further, it is hard to imagine a data-generating process that could produce the Sakoda (2023) results that does not involve some form of racial bias dur-

ing the pre-reform period.⁶² Yet, neither study provides direct evidence that community corrections officers treat Black persons under their supervision differently than others. The key point is that policies that require or allow community corrections officers to send persons under their supervision to prison as punishment for violating administrative rules of conduct create hundreds of thousands of incarceration events each year that appear to do little to improve public safety. It is possible that violations of these rules are weak signals of the propensity to commit crime. If so, these bad policies are forms of structural racism, even if they are enforced in a color-blind manner, because Black persons are grossly overrepresented among those under community supervision.

SIMULATING THE IMPACTS OF REFORMS

We have examined sources of racial disparities in court outcomes in the US. In our framework, bias exists when courts systematically produce results that miss target outcomes. These target outcomes balance the social costs of punishing offenders, which include the loss of liberty that defendants suffer as well as the fiscal costs of operating jails, prisons, and community supervision programs, and the public benefits that punishment produces, which include lower rates of violent victimization, lower insurance costs, reduction in crime prevention expenditures, and so forth. Researchers have not developed perfect ways to measure many of these costs and benefits, but within our framework, courts are likely biased if they regularly assign different punishments to offenders of different races even when these defendants are comparable on dimensions that determine socially optimal punishment. If two sets of defendants of different races should be assigned the same punishment, and the court assigns different average punishments to them, at least one group is receiving the wrong average sentence, and the sentences produced by the court are there-

61. Peter Finn and Sarah Kuck (2003) argue that community corrections officers are often stressed by heavy caseloads that make careful supervision of offenders difficult. Some officers may respond by adopting rules of thumb that map violations into revocation recommendations in a mechanical way.

62. The 17-percentage-point change in the Black-White revocation rate gap is striking, and Sakoda (2023) provides no evidence that crime rates jumped among Black probationers.

fore both racially biased and socially inefficient. This scenario may arise because courts assign punishments that vary with race among persons charged with the same offense who pose comparable recidivism risks or because the law requires courts to assign different punishments for distinct types of offenses even though the socially optimal sentencing rules for all the types of offenses in question are the same. In the latter scenario, any correlation between race and the nature of the charges defendants face generates racial differences in average sentencing outcomes that are socially inefficient.

In addition, courts may produce outcomes that are independent of race conditional on relevant defendant characteristics but are systematically too punitive relative to sentencing outcomes that balance the social costs and benefits of punishment. Since Black and Hispanic citizens are arrested and charged at higher-than-average rates, they bear disparate costs when courts are too punitive.⁶³ Neal and Rick (2016) and Raphael and Stoll (2013) present evidence that courts applied the punitive sentencing reforms that drove the prison boom in a race-neutral way. In broad terms, the policies that drove the prison boom did not change relative racial disparities in punishment given arrest but instead created enormous disparate impacts on minority citizens. Black citizens suffered the greatest harms because they have faced much higher arrest rates than Whites for at least half of a century.

Taken as a whole, the studies we review do not provide clear evidence that most court actors assign wildly different average treatment to comparable defendants of different races. The Rose (2021) study of probation reform in North Carolina and the Lofstrom and colleagues (2014) study of parole reform in California provide evidence that Black persons on probation in North Carolina and parole in California faced excessive rates of technical re-

vocations prior to the reforms these articles study. However, neither study pins down the extent to which these results reflect racial bias among community corrections officers or disparate impacts of the rules that govern community supervision. For example, in North Carolina, probation revocations proceedings in the pre-reform period often began because a probationer failed to pay certain fees and fines, and many states impose conditions of supervision that are easier to meet when persons under supervision live in families and communities with more resources. If we assume that access to resources is only weakly correlated with recidivism risk, it is easy to imagine that color-blind enforcement of rules that require persons under community supervision to pay fines, find work, secure housing in low-crime areas, and so forth could have large disparate impacts on Black probationers. Rose (2021) explicitly notes that the patterns he documents may arise because probation officers enforce inefficient rules in a color-blind way.

We also note above that Rehavi and Starr (2014) conclude that eliminating the racial bias in the charging decisions made by federal prosecutors during the period 2006–2008 would have lowered the steady-state population of Black inmates in federal prisons by roughly 10 percent. However, federal prison populations have always been much smaller than the total populations in state prisons, and the federal prison population shrank by roughly one-fourth between 2008 and 2019 and has remained roughly constant since. Finally, federal prisons shrank, in part, because the federal government passed legislation that sought to reduce racial disparities in federal incarceration rates.⁶⁴

If the federal courts adopted a reform tomorrow that created a distinct unit of attorneys tasked with assigning charges but not trying cases, and this reform eliminated the bias that Rehavi and Starr (2014) identify, the stock of

63. See NASEM (2022) and Allen Beck (2021) for information on racial differences in offense and arrest rates.

64. See “Generate Inmate Populations Report,” n.d., https://www.bop.gov/about/statistics/population_statistics.jsp#pop_report_cont; Fair Sentencing Act of 2010, Public Law 111–220, August 3, 2010, <https://www.congress.gov/111/plaws/publ220/PLAW-111publ220.pdf>; and “First Step Act,” Federal Bureau of Prisons, n.d., <https://www.bop.gov/inmates/fsa/>. Rose (2021) identifies potential racial bias among probation officers in North Carolina, but this is not a large driver of overall racial disparities in incarceration rates.

Black prisoners in federal custody would decline by less than ten thousand. However, the simulation results we will present demonstrate that it may be possible to reduce the number of Black citizens in jails and state prisons by well over three hundred thousand if states implement reforms that scale back punitive policies that provide marginal benefits in terms of crime reduction and public safety.

In sum, we contend that implementing the world's most effective antibias training for prosecutors, judges, and community supervision officers would produce, at best, relatively small reductions in observed racial disparities in court outcomes. In contrast, reforms that change the practices, policies, and laws that shape or even constrain the decisions that prosecutors, judges, and community supervision officers make are more powerful tools for reducing racial disparities in court outcomes.

Here, we conduct a simulation exercise that allows us to assess the quantitative significance of three specific reforms: elimination of pretrial detention for nonviolent offenders who are not on probation or parole, elimination of technical revocations of parole that lead to re-incarceration, and reduction of prison admission rates for convictions on nonviolent offenses by one half. We first discuss the motivations for these three simulations. We then discuss the simulation methods and results.

Three Simple Changes

We have argued that policymakers could replace cash bail systems with pretrial detention systems that produce less disparate harm for minority communities while sacrificing little in terms of public safety. Since cash bail rations pretrial release according to a defendant's financial capacity rather than the risk the defendant poses to the community, it should be possible to lower jail populations without harming public safety by tying release decisions more closely to violation risks rather than the financial capacity of defendants.

Here, we provide a rough sense of the potential impacts of such reforms by considering the impacts of a policy that mandates pretrial release for all persons charged with nonviolent crimes who are not on probation or parole at the time of their arrest. Since this ban does not apply to violent offenders or those with recent convictions, this policy experiment provides some insight into the potential gains of linking pretrial detention outcomes to public-safety risks rather than the financial capacity of defendants.

We also simulate the elimination of technical parole revocations. To date, no empirical studies provide evidence that parole officers effectively target offenders who pose relatively high risks to public safety when they initiate technical parole revocations that result in parolees returning to prison. Further, Lofstrom and colleagues (2014) show that the elimination of technical revocations in California reduced prison populations without serious harms to public safety.

Finally, we simulate the impacts of reducing prison admission rates for nonviolent offenders by one-half. Neal and Rick (2016) show that, during the prison boom, the likelihood that persons arrested for nonviolent crimes would serve prison terms of short, medium, and long lengths rose sharply, often by a factor of two or more.⁶⁵ Further, in the section "Case Outcomes," we provide considerable evidence that the crime reduction benefits of increasing incarceration rates diminish with the size of prison populations.⁶⁶

Before presenting the results of these three simulations and a fourth simulation that captures the simultaneous implementation of all three reforms, we note that none of these hypothetical reforms impact how courts respond to violent crime. These reforms address jail time and prison time for nonviolent offenses as well as returns to prison as punishment for behaviors that violate no criminal laws. Our goal is to determine how much policymakers

65. Neal and Rick (forthcoming) provide evidence that these trends were driven largely by changes in the sentences assigned to convicted offenders rather than changes in conviction rates given arrest.

66. Incarceration rates did decline after 2008, but in 2019, the incarceration rate remained roughly four times greater than the rate in 1975. Prison populations shrank somewhat during COVID, but it is too soon to know whether incarceration rates will rebound to pre-pandemic levels.

can shrink prison and jail populations through reforms that may well have limited public safety costs. We conclude that policymakers could likely use such reforms to reduce the number of persons in jail or state prison by more than 40 percent.

Methods

We use steady-state simulation methods to evaluate the potential impacts of our three criminal justice reforms. The appendix details our methodology. Here, we provide an overview.

We define thirty-two states that a person may occupy. Three states involve persons who are neither facing a felony charge nor under a sentence for a felony conviction.⁶⁷ These persons may be under probation supervision, parole supervision, or no supervision. We next define twenty-four types of defendants who are facing felony charges that are being adjudicated. Twelve of these types are detained in jail while courts resolve their cases. The other twelve types are granted release. Each type within each set of twelve types is defined by the interaction of four charge categories and indicators for whether the defendant faces a charge while on probation, on parole, or under no community supervision.⁶⁸ Finally, we define five types of prison inmates: persons in prison as the result of technical violations of parole and persons in prison following convictions on charges specific to one of our four crime categories.⁶⁹

We employ data from the State Court Processing Statistics, the National Corrections Reporting Program, and other sources to calculate transition probabilities between these thirty-two states. We set many of the possible transition probabilities to zero to respect the legal rules that govern these transitions. For example, persons cannot go from the state of no community supervision to prison without first

going through one of the twenty-four states that describes the operation of courts.⁷⁰

Given an estimate of the transition matrix that generates movements among our thirty-two states, we can calculate the steady-state fractions of the population that occupy each state. We then multiply these fractions by the total US population in 2019 to produce baseline estimates of the counts in each state. We calculate four alternative steady states that capture the results of implementing each of our three potential reforms in isolation as well as the results of implementing all three simultaneously. Table 2 presents four key steady-state populations for the four scenarios: those in jail on pretrial detention, on probation, on parole, and in prison. The table also presents five subpopulations of prison inmates defined by our four charge categories and the possibility of entering prison following a technical parole violation.

Results

Table 2 presents our results. Column 2 presents the steady-state stock levels implied by the transition matrix that we construct. These levels are often reasonably close to the actual values for 2019. We chose 2019 because the impacts of COVID likely moved the various stocks in the criminal justice system far away from steady-state values, and our counterfactual exercises involve comparing different steady states.

Column 3 presents the implied steady-state stocks if we set pretrial detention rates to zero for all nonviolent offenders who are not under probation or parole supervision while holding other transition rates constant. This change shrinks the stock of persons on pretrial detention by roughly 58 percent and also shrinks the stock of prisoners and persons under probation supervision by 4 percent and 7 percent re-

67. We do not model events related to misdemeanor charges.

68. The four charge categories are violent, property, drug, and public order.

69. We are implicitly treating persons sent to prison for technical probation violations as persons entering prison for the charge that generated the probation sentence.

70. Some restrictions are more obvious than others. For example, in our model, persons cannot leave prison and directly enter probation supervision. This constraint holds in many states but not all because of variation among states in the definition of probation versus parole supervision.

Table 2. Simulated Steady-State Stocks: Three Possible Reforms

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---------------------|-----------|-----------|-------------|--------------|---------------|-----------|
| | 2019 | 2019 | No Pre- | No Technical | .5* Admission | All Three |
| | Actual | Simulated | Detention | Revocations | Rate | Reforms |
| | | | Nonviolent, | | Nonviolent | |
| | | | Non-Sex | | | |
| Pretrial Detention | 482,235 | 681,052 | 291,090 | 687,102 | 672,828 | 287,430 |
| Probation | 1,816,308 | 1,447,432 | 1,347,577 | 1,447,431 | 1,753,427 | 1,627,374 |
| Parole | 878,900 | 775,168 | 728,266 | 850,960 | 515,134 | 538,844 |
| Prison Violent | 639,899 | 604,783 | 601,198 | 607,165 | 603,077 | 601,158 |
| Prison Property | 146,037 | 123,237 | 112,684 | 123,864 | 61,421 | 56,379 |
| Prison Drug | 138,879 | 121,677 | 111,910 | 122,418 | 60,355 | 55,749 |
| Prison Public Order | 128,708 | 125,219 | 114,323 | 125,848 | 62,272 | 57,054 |
| Prison Violation | 202,291 | 155,637 | 146,220 | 0 | 103,428 | 0 |
| Prison Total | 1,255,815 | 1,130,554 | 1,086,335 | 979,296 | 890,553 | 770,339 |

Source: Authors' calculations.

Note: This table presents results from four steady-state policy simulations. Appendix A describes how we calculate steady-state populations. Column 2 is our baseline simulation for 2019. Column 3 calculates a new steady-state under the assumption that all defendants who are not under community supervision and not charged with a violent offense are released pretrial. Column 4 presents results when we eliminate reincarceration for technical parole violations. Column 5 reports the results of halving the prison admission rate among those convicted of nonviolent offenses. Column 6 reports the results of implementing all three reforms simultaneously. The prison population numbers are estimates of the populations in state prisons. We are not able to run similar simulations for the federal system.

spectively. These latter reductions reflect the fact that recent research indicates that conviction rates are lower when defendants are not detained pretrial.⁷¹

Column 4 presents the implied stocks if we set the rate of technical parole revocations to zero. This change shrinks the prison population by just over 150,000 persons or roughly 13 percent of baseline simulated stock for 2019. This change also increases the number of persons under parole supervision by roughly 75,000.

Column 5 presents the implied stocks if we cut prison admission rates in half for those convicted of nonviolent crimes. The new steady-state population in state prisons is more than 20 percent lower, and the number of persons under probation supervision grows by over 300,000.

In the final column, we implement all three potential reforms at once. While each reform generates modest impacts when implemented alone, the results of implementing all three at once are quite striking. In the simulated 2019 baseline, there are more than 1.8 million persons behind bars: 681 thousand in jail and more than 1.13 million in state prisons. In our final simulation, the total incarcerated population is less than 1.06 million. Put differently, these three reforms, when implemented simultaneously, could reduce the number of persons incarcerated in jails and state prisons by about 40 percent. Further, the total number of persons under community supervision changes little since the increase in probationers is more than offset by a decline in parolees.⁷²

We do not possess the data required to calculate steady-state populations by race, but if

71. See, for example, Dobbie et al. (2018).

72. We take most parameters for these simulations from published reports or research papers. The key parameter that we estimate is the fraction of parole revocations that are the result of technical violations. We reran

these reforms were adopted, there is no reason to believe that Black or Hispanic defendants would see smaller than average reductions in their incarceration rates. Given existing evidence on racial disparities in income and wealth, reducing the number of persons detained pretrial, often because they cannot pay relatively small cash bail amounts, may benefit minority defendants more than White defendants. Further, we know that some prior reforms to the use of technical revocations greatly narrowed racial disparities in prison admissions through this channel.⁷³

We have examined data from the NCRP to assess the impact of curtailing technical parole violations in California on racial disparities in prison admissions. From 2008 through 2011 (the four-year period mostly preceding the reform), we estimate average annual admissions per 100,000 for technical parole violations in California equal to 1,049 for Black people, 191 for Hispanic people, and 174 for White people. For the first four years following the realignment reform implemented in October 2011, the comparable average annual admission rates are 78 for Black people, 17 for Hispanic people, and 10 for White people. Hence, the Black-White differential in annual admissions of this type shrank from 875 per 100,000 to 68 per 100,000 (a 92 percent decrease).

Further, even if these reforms generated the same proportional reductions in incarceration rates for all racial groups, the absolute reductions would be greater for Black and Hispanic citizens than for Whites, and much greater for Black citizens. At the end of 2019, the number of White persons in prison was less than the number of Black inmates. However, the White population was more than five times greater

than the Black population. Our results in column 6 imply that implementing all three reforms at once would reduce state prison populations by 32 percent. Given 2019 prison populations as a starting point, a 32 percent reduction in total incarceration rates for all race groups implies that more than 460 additional adults per 100,000 Black adult citizens would be living outside prison, but among Whites, the comparable figure is less than 85.⁷⁴

The difference between columns 2 and 6 in counts of persons in jail waiting for a verdict is larger than the difference in total state prison populations, even though the baseline pretrial detention population in column 2 is much smaller than the baseline state prison population. We have no reliable data on racial differences in pretrial detention rates, but the overall jail incarceration rate in 2019 among Blacks was more than three times greater than the rate among Whites.⁷⁵ The change in pretrial detention populations between columns 2 and 6 likely implies a drop in jail incarceration rates among Blacks that is more than 200 per 100,000, but the comparable figure among Whites is far less than 100.⁷⁶

CONCLUSION

We have explored the sources of racial disparities in criminal court case outcomes within a utilitarian framework. In this framework, optimal case outcomes do not depend on the race of the defendant, so the existence of studies that report different average outcomes by race among comparable defendants raises the possibility that biased decision-making is a key source of racial disparities in case outcomes. However, we conclude that structural factors that generate socially inefficient punishment

these simulations using an independent and lower estimate of this fraction from the Council of State Governments. The total prison populations in columns 4 and 6 increased by roughly 19,000 and 8,000 respectively. The latter change is just over 1 percent of the corresponding entry in table 2. See the appendix for details.

73. See our earlier discussion of Rose 2021 and Sakoda 2023.

74. See tables 3 and 6 in Carson (2020).

75. See Zeng and Minton 2021, table 2. Jail incarceration rates for Hispanics are actually slightly lower than rates for Whites.

76. The BJS reports that 65 percent of jail inmates are being held pretrial (Zeng and Minton 2021) In 2019, the same report tabulates the average daily jail population in the United States to be 741,900.

for many types of defendants are more consequential as drivers of these racial disparities.

The results in table 2 highlight the potential efficacy of simple structural reforms as tools for reducing racial disparities in incarceration. Nonetheless, large racial differences in steady-state stocks remain even when we simulate the implementation of all three reforms simultaneously, and there is likely no set of reforms to courts or sentencing laws that can completely eliminate these differences without producing significant harms to public safety. Blacks citizens are arrested for homicide and other violent crimes at much higher rates than Whites, and victimization surveys support the hypothesis that these differences in arrest rates are largely driven by differences in offense rates. Further, except for robbery, most violent crime victimization is within race. Thus, reforms that drastically reduce the incarceration of violent offenders would possibly visit harm on the millions of citizens in minority communities who are not engaged in criminal activity.⁷⁷

Our charge does not require that we explore the various reforms to education, social, and economic policies that could mitigate or eliminate racial differences in violent or nonviolent offending rates or arrest rates. The reform ideas we present here offer ways to dramatically reduce absolute racial disparities in court outcomes if we hold constant racial disparities in arrest rates. Further, a common thread ties these ideas together. Policymakers can both improve social welfare and mitigate racial disparities in court outcomes by reducing the use of punitive sanctions that produce limited reductions in crime and small improvements in public safety. Within our framework, steps in this direction are the policy-reform equivalent of picking the low-hanging fruit, and our simulation results suggest that a bountiful harvest is there for the taking.

APPENDIX: SIMULATION DETAILS

We use data from the 2019 NCRP and various BJS publications to calibrate a steady-state model of US correctional populations, with a focus on the populations in pretrial detention, on felony probation, in state prisons, and on

parole. We do not model the federal courts or federal prison populations. We use the NCRP data to model the transition probabilities between thirty-two possible states over the course of a year, simulate the steady state given these empirical transition probabilities, and then simulate the effects of various change in correctional practices.

States and Potential Transitions

The thirty-two states are listed in table A.1. There are three states for people who are not currently in a state prison and who are not felony defendants in an ongoing court case. These states are not on probation or parole, on probation, on parole. We define twenty-four states that characterize types of persons who are in the process of being adjudicated for a felony arrest. The second column lists the twelve types who are detained pretrial. The third column lists the twelve types who are released pretrial. The states in these two columns are defined by all the possible combinations of four offense categories (violent, property, drug, public order), three status categories (not on probation or parole, on probation, on parole), and two pretrial detention categories (detained pretrial, released pretrial). The sum of the populations in the second column gives the number of jail inmates who are being detained pretrial on a felony case. Finally, we define five prison states. These include categories that distinguish prisoners according to why they were admitted to prison. We divide new court commitments into convictions for violent, property, drug, or public-order crimes. The final category is technical parole revocations.

With thirty-two categories, the transition probability matrix that describes the likelihood of transitioning between any two states involves 1,024 probabilities. Many of these are zero by construction or by assumption. For example, among persons who are not felony defendants, not under supervision, and not incarcerated, there are positive transition probabilities to many states that involve facing charges, but no person can transition from freedom to prison or community supervision

77. See chapter 2 of NASEM (2022) for more on racial differences in reported victimization.

Table A.1. Simulation Details

| Description of Subpopulation Modelled Using the Steady-State Population Shares Implied by the Empirical Transition Probabilities | | | |
|--|---|---|--|
| Non-Felony Defendants | Felony Defendants Detained Pretrial ^a | Felony Defendants Released Pretrial | Prison Inmate Populations ^b |
| (1) Not on probation or parole | (1) not on probation/parole, violent arrest | (1) not on probation/parole, violent arrest | (1) Violent offense |
| (2) On probation | (2) On probation, violent arrest | (2) On probation, violent arrest | (2) Property offense |
| (3) On parole | (3) On parole, violent arrest | (3) On parole, violent arrest | (3) Drug offense |
| | (4) not on probation/parole, property arrest | (4) not on probation/parole, property arrest | (4) Public order offense |
| | (5) On probation, property arrest | (5) On probation, property arrest | (5) technical parole violation |
| | (6) On parole, property arrest | (6) On parole, property arrest | |
| | (7) not on probation/parole, drug arrest | (7) not on probation/parole, drug arrest | |
| | (8) On probation, drug arrest | (8) On probation, drug arrest | |
| | (9) On parole, drug arrest | (9) On parole, drug arrest | |
| | (10) not on probation/parole, public order arrest | (10) not on probation/parole, public order arrest | |
| | (11) On probation, public order arrest | (11) On probation, public order arrest | |
| | (12) On parole, public order arrest | (12) On parole, public order arrest | |

Source: Authors' compilation.

^a Sum of the subpopulations in this column provides the pretrial felony jail population.

^b Sum of the subpopulations in this column provides the state prison population.

without going through the courts. As another example, those in prison cannot transition directly to court. They may stay in prison, exit prison directly, or transition to parole.

For all arrested defendants, we calculate the probability that they are detained pretrial or released. Then, for each custody state, we calculate the probability of various case outcomes. For felony defendants who are not on probation or parole at the time of arrest, we calculate, for each offense category, the probabilities of being acquitted, convicted and sentenced to probation, and convicted and sent to prison. For defendants on probation or parole at the time of arrest, we estimate the probabilities

that defendants are convicted and admitted, convicted but not admitted, and neither convicted nor admitted to prison. We assume that pretrial defendants under community correction supervision at the time of a new arrest do not successfully exit supervision in the year of a new arrest. They can only transition either to prison or back to their original community corrections status.

Finally, people incarcerated in state prison can transition to being not on probation or parole (an unconditional release) or to being on parole (conditional release), or they can continue in prison. We again calculate these probabilities separately for each offense category.

Transitions into Court

From the BJS analysis of the 2009 SCPS data (Reaves 2013), we observe that among felony defendants in large urban counties, 24.9 percent were arrested for violent offenses, 29.1 percent were arrested for property offense, 32.6 percent were arrested for drug offenses, and 13.4 percent were arrested for public-order offenses. In 2019, there were 495,871 part 1 violent offenses.⁷⁸ Assuming that all part 1 violent arrests are for felony offenses, we use the percentages for 2009 felony defendants to impute total offenses for the other offense categories (that is, for property, drug, and public-order felony arrests).

We use tabulations from Brian Reaves (2013) pertaining to the proportion of arrests within the four offense categories involving persons with no criminal justice status, persons on parole, and persons on probation to apportion arrests across these categories. In addition, we use tabulations from Reaves (2013) pertaining to overall pretrial release rates by offense category and overall release rates by criminal justice status to impute pretrial release rates by offense and criminal justice status. These imputations permit us to generate estimates of the transition probabilities from each of the three states in first column of table A.1 to the twenty-four states listed in the second and third columns.

Transitions out of Court

From Reaves (2013) we also observe overall felony conviction rates by offense group as well as the difference in conviction rates between those detained pretrial and those released pretrial. We generate conviction rates for each of the twenty-four groups listed in the second and third columns of table A.1 assuming that the detained-released conviction ratio is the same within offense groups. We then apply these imputed conviction totals to arrests within groups defined by community correction status, arrests offense, and pretrial detention status. This step yields total convictions within each

offense group. Among those who are not convicted, we assume that defendants transition back to their original state.

We apportion the convicted within each offense group into those transitioning into prison and those transitioning into community corrections in the following manner. Using the 2019 admissions data from the NCRP, we first estimate the distribution of new court commitment prison admission across the four offense categories. We then apply this distribution to total admissions in 2019 to generate new court commitments for violent, property, drug, and public-order offense.⁷⁹ The ratio of total imputed convictions by offense category to admissions by offense category provides the admissions rate conditional on conviction for each of the four offense groups. We use these results to generate prison admissions totals for each of the twenty-four groups listed in the second and third columns of table A.1. We assume that convicted persons who are not admitted to prison either transition to probation (if not on probation-parole at arrest or on probation at arrest) or to parole (if on parole at time of arrest).

Transitions out of Community Supervision

We rely on the BJS analysis provided in Barbara Oudekerk and Danielle Kaeble's (2021) study to calculate values for the total probation population and the number of successful exits from probation. We assume that the remaining probation population (the overall population less the successful exits, less those that we impute being admitted to prison on a new court commitment) continues on probation into the next year.

For the parole population, we pull values for the overall population, the number of successful exits, and the number that are returned to prison from parole (Oudekerk and Kaeble 2021). From the total returned to prison, we subtract the numbers that we impute being admitted to prison on a new court commitment. The balance we treat as being admitted for a

78. See 2019 Crime in the United States, table 29, <https://ucr.fbi.gov/crime-in-the-u.s/2019/crime-in-the-u.s.-2019/topic-pages/tables/table-29>, accessed on August 9, 2024.

79. Note that the NCRP 2019 data cover only forty-six states. Hence, rather than using total admissions from the NCRP, we estimate the distribution of court commitments across offense categories and then pull total admissions from the BJS analysis of all fifty states presented in Carson (2021).

technical parole violation. We estimate that roughly 87 percent of admissions to prison from parole are the result of technical parole violations.⁸⁰

Transitions out of Prison

From Ann Carson (2021), we glean that 19.3 percent of people released from prison are released unconditionally, with the remainder presumably released to parole. We use the NCRP to estimate the distribution of releases by the following commitment offense categories: new commitment, violent; new commitment, property; new commitment, drug; new commitment, public order; and technical parole violations. We apply this distribution to total releases from state prison in 2019 to generate an estimate of each flow. We also use the 2019 NCRP to estimate the distribution of the incarcerated population across these same offense categories, and we apply this distribution to the total state prison population for that year to generate stocks. These flows plus the stock estimates provide transition probabilities from each of the five prison incarceration groupings to three possible states: free with no community corrections supervision, on parole, or still incarcerated.

Solving for the Steady-State Populations Shares

The imputations we have discussed generate estimates of the probability of transitioning between each of the thirty-two categories described in table A.1 over the course of year. We array these transition probabilities in a 32×32 matrix, T , where T_{ij} is the probability of transitioning from state i to state j over the course of a year. Note that each element in T is nonnegative and each row of T sums to one. Let T' denote the transpose of T . An eigenvector, v , of T' must satisfy:

$$T'v = \lambda v$$

where λ is the associated eigenvalue. Holder's Inequality guarantees that the largest eigen-

value of T' is always one. Therefore, v^* contains the steady-state population shares of the system defined by the laws of motion coded in T' if

$$T'v^* = v^*$$

and v^* is the eigenvector associated with $\lambda = 1$ that is normalized so that shares sum to one. We multiply the elements v^* by the US population in 2019. The aggregation of the twelve population estimates for the categories in column 2 of table A.1 gives the steady-state pretrial felony jail population. The sum of the estimates for the populations for the five prison groupings in the last column yields the estimated prison population.

REFERENCES

- Alexander, Michelle. 2011. "The New Jim Crow." *Ohio State Journal of Criminal Law* 9(1): 7-26.
- Amy, Jeff. 2024. "Georgia Could Require Cash Bail for 30 More Crimes, Including Many Misdemeanors." *Associated Press*, February 1. <https://apnews.com/article/georgia-bail-bond-misdemeanors-crimes-db44e8e147fc1b7296610535c1f21a27>.
- Arnold, David, Will Dobbie, and Peter Hull. 2022. "Measuring Racial Discrimination in Bail Decisions." *American Economic Review* 112(9): 2992-3038.
- Arnold, David, Will Dobbie, and Crystal S. Yang. 2018. "Racial Bias in Bail Decisions." *Quarterly Journal of Economics* 133(4): 1885-932.
- Ayres, Ian, and Joel Waldfogel. 1993. "A Market Test for Race Discrimination in Bail Setting." *Stanford Law Review* 46(5): 987-1047.
- Barbarino, Alessandro, and Giovanni Mastrobuoni. 2014. "The Incapacitation Effect of Incarceration: Evidence from Several Italian Collective Pardons." *American Economic Journal: Economic Policy* 6(1): 1-37.
- Barnes, Geoffrey C., Jordan M. Hyatt, Lindsay C. Ahlman, and Daniel T. L. Kent. 2012. "The Effects of Low-Intensity Supervision for Lower-Risk Probationers: Updated Results from a Randomized

80. A 2019 report from the Council of State Governments (2019), estimates that roughly 63.6 percent of prison admissions for people on parole are for technical parole violations. If we use this rate instead, our simulated total prison population, in the case where we consider all three reforms simultaneously, is 779,478 instead of 770,339. This is an increase of just over 1 percent.

- Controlled Trial." *Journal of Crime and Justice* 35(2): 200–20.
- Beck, Allen J. 2021. "Race and Ethnicity of Violent Offenders and Arrestees, 2018." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/race-and-ethnicity-violent-crime-offenders-and-arrestees-2018>.
- Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Loken, and Magne Mogstad. 2020. "Incarceration, Recidivism and Employment." *Journal of Political Economy* 128(4): 1269–324.
- Blumstein, Alfred, Jacqueline Cohen, Susan E. Martin, and Michael H. Tonry, eds. *Research on Sentencing: The Search for Reform*. The National Academies Press, 1983.
- Bogira, Steve. 2006. *Courtroom 302: A Year Behind the Scenes in an American Criminal Courthouse*. Vintage.
- Brinkman, Luke, Andrew Jordan, and Derek Neal. 2025. "The Impacts of Parole Supervision." Unpublished technical report.
- Buonanno, Paolo, and Steven Raphael. 2013. "Incarceration and Incapacitation: Evidence from the 2006 Italian Collective Pardon." *American Economic Review* 103(6): 2437–65.
- CRPC (California Committee on Revision of the Penal Code). 2021. *Annual Report and Recommendations: Committee on Revision of the Penal Code*. California Law Revision Commission.
- Canay, Ivan A., Magne Mogstad, and Jack Mountjoy. 2024. "On the Use of Outcome Tests for Detecting Bias in Decision Making." *Review of Economic Studies* 91(4): 2135–67.
- Carson, E. Ann. 2020. "Prisoners in 2019." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/prisoners-2019>.
- Carson, E. Ann. 2021. "Prisoners in 2020—Statistical Tables." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/prisoners-2020-statistical-tables>.
- Carson, E. Ann. 2022. "Prisoners in 2021—Statistical Tables." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/prisoners-2021-statistical-tables>.
- Carson, E. Ann, and Rich Kluckow. 2022. "Prisoners in 2022—Statistical Tables." Bureau of Justice Statistics.
- Carson, E. Ann, and Rich Kluckow. 2023. "Correctional Populations in the United States, 2021—Statistical Tables." Bureau of Justice Statistics. <https://bjs.ojp.gov/document/cpus21st.pdf>.
- Carson, E. Ann, and William J. Sabol. 2012. "Prisoners in 2011." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/prisoners-2011>.
- Chalfin, Aaron, and Justin McCrary. 2017. "Criminal Deterrence: A Review of the Literature." *Journal of Economic Literature* 55(1): 5–48.
- Chalfin, Aaron, and Justin McCrary. 2018. "Are US Cities Underpoliced? Theory and Evidence." *Review of Economics and Statistics* 100(1): 167–86.
- Chan, David C., Matthew Gentzkow, and Chuan Yu. 2022. "Selection with Variation in Diagnostic Skill: Evidence from Radiologists." *Quarterly Journal of Economics* 137(2): 729–83.
- Chohlas-Wood, Alex, Joe Nudell, Keniel Yao, Zhiyuan Lin, Julian Nyarko, and Sharad Goel. 2021. "Blind Justice: Algorithmically Masking Race in Charging Decisions." In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery.
- Cohen, Mark A. 2000. "Measuring the Costs and Benefits of Crime and Justice." In *Measurement and Analysis of Crime and Justice*, vol. 4 of *Criminal Justice 2000*, edited by David Duffee. National Institute of Justice.
- CSG (Council of State Governments). 2019. *Confined and Costly: How Supervision Violations are Filling Prisons and Burdening Budgets*. Council of State Governments. <https://csgjusticecenter.org/wp-content/uploads/2020/01/confined-and-costly.pdf>.
- Currie, Janet, W. Bentley MacLeod, and Kate Musen. 2025. "Doctor Decision Making and Patient Outcomes." NBER working paper 32788. National Bureau of Economic Research.
- Davidson, Janet, George King, Jens Ludwig, and Steven Raphael. 2019. "Managing Pretrial Misconduct: An Experimental Evaluation of HOPE Pretrial." Working paper. Goldman School of Public Policy.
- Demuth, Stephen, and Darrell Steffensmeier. 2004. "The Impact of Gender and Race-Ethnicity in the Pretrial Release Process." *Social Problems* 51(2): 222–42.
- Dobbie, Will, Jacob Goldin, and Crystal S. Yang. 2018. "The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges." *American Economic Review* 108(2): 201–40.
- Donald, Bernice, Jeffrey Rachlinski, and Andrew Wistrich. 2020. "Getting Explicit About Implicit Bias." *Judicature* 104(3): 75–80.

- Finn, Peter, and Sarah Kuck. 2003. "Addressing Probation and Parole Officer Stress." National Institutes of Justice. <https://www.ojp.gov/pdffiles1/nij/grants/207012.pdf>.
- Franco, Catalina, David J. Harding, Shawn D. Bushway, and Jeffrey Morenoff. 2024. "Does Re-Imprisonment for Technical Violations Prevent Crime?" *Journal of Human Resources*, ahead of print.
- Grant, Glenn A. 2019. "Report to the Governor and the Legislature." New Jersey Courts. <https://www.njcourts.gov/sites/default/files/cjrannualreport2019.pdf>.
- Grunwald, Ben. 2023. "Racial Bias in Criminal Records." *Journal of Quantitative Criminology* 40(3): 1–43.
- Huebner, Beth M., and Timothy S. Bynum. 2008. "The Role of Race and Ethnicity in Parole Decisions." *Criminology* 46(4): 907–38.
- Hyatt, Jordan M., and Geoffrey C. Barnes. 2017. "An Experimental Evaluation of the Impact of Intensive Supervision on the Recidivism of High-Risk Probationers." *Crime & Delinquency* 63(1): 3–38.
- Johnson, Rucker, and Steven Raphael. 2012. "How Much Crime Reduction Does the Marginal Prisoner Buy?" *Journal of Law and Economics* 55(2): 275–310.
- Jordan, Andrew. 2022. *Racial Patterns in Approval of Felony Charges*. SSRN Electronic Journal.
- Jordan, Andrew. 2024. "Plea Bargaining with Risk Aversion." working paper, mimeo, Washington University in St. Louis.
- Jordan, Andrew, Ezra Karger, and Derek Neal. 2023. "Heterogeneous Impacts of Sentencing Decisions." NBER working paper 31939. National Bureau of Economic Research.
- Jorgensen, Isabella, and Sandra Susan Smith. 2021. "The Current State of Bail Reform in the United States: Results of a Landscape Analysis of Bail Reforms Across All 50 States." HKS working paper RWP21-033. Harvard Kennedy School.
- Kaeble, Danielle. 2023. "Probation and Parole in the United States, 2021." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/probation-and-parole-united-states-2021>.
- Katz, Charles M., and Cassia C. Spohn. 1995. "The Effect of Race and Gender on Bail Outcomes: A Test of an Interactive Model." *American Journal of Criminal Justice* 19: 161–84.
- Kim, Jaeok, and André Kiesel. 2018. "The Long Shadow of Police Racial Treatment: Racial Disparity in Criminal Justice Processing." *Public Administration Review* 78(3): 422–31.
- Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan. 2017. "Inherent Trade-Offs in the Fair Determination of Risk Scores." In *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- Kutateladze, Besiki L., Nancy R. Andiloro, Brian D. Johnson, and Cassia C. Spohn. 2014. "Cumulative Disadvantage: Examining Racial and Ethnic Disparity in Prosecution and Sentencing." *Criminology* 52(3): 514–51.
- Lacoe, Johanna, Brett Fischer, and Steven Raphael. 2023. "The Effect of Pre-Arrest Legal Representation on Criminal Case Outcomes." NBER working paper 31289. National Bureau of Economic Research.
- Levitt, Steven D. 1996. "The Effect of Prison Population Size on Crime Rates: Evidence from Prison Overcrowding Litigation." *Quarterly Journal of Economics* 111(2): 319–51.
- Liedka, Raymond V., Anne Morrison Piehl, and Bert Useem. 2006. "The Crime-Control Effect of Incarceration: Does Scale Matter?" *Criminology & Public Policy* 5(2): 245–76.
- Light, Michael T. 2022. "The Declining Significance of Race in Criminal Sentencing: Evidence from US Federal Courts." *Social Forces* 100(3): 1110–41.
- Lofstrom, Magnus, Joseph Michael Hayes, Brandon Martin, and Deepak Premkumar. 2021. *Racial Disparities in Law Enforcement Stops*. With research support from Alexandria Gumbs. Public Policy Institute of California. <https://www.ppic.org/publication/racial-disparities-in-law-enforcement-stops/>.
- Lofstrom, Magnus, Brandon Martin, and Steven Raphael. 2020. "The Effect of Sentencing Reform on Racial and Ethnic Disparities in Involvement with the Criminal Justice System: The Case of California's Proposition 47." *Criminology & Public Policy* 19(4): 1165–207.
- Lofstrom, Magnus, and Steven Raphael. 2016. "Incarceration and Crime: Evidence from California's Public Safety Realignment Reform." *Annals of the American Academy of Political and Social Science* 664(1): 196–220.
- Lofstrom, Magnus, Steven Raphael, and Ryken Grattet. 2014. *Is Public Safety Realignment Reducing*

- Recidivism in California?* Public Policy Institute of California. <https://www.ppic.org/publication/is-public-safety-realignment-reducing-recidivism-in-california/>.
- Macdonald, John, and Steven Raphael. 2020. *An Analysis of Racial/Ethnic Disparities in Stops by Los Angeles County Sheriff's Deputies in the Antelope Valley*. NCCD. https://lasd.org/wp-content/uploads/2020/10/AVDOJ_Analysis_of_LASD_Stops_in_the_AV_January-July_2019.pdf.
- Marvell, Thomas B., and Carlisle E. Moody. 1994. "Prison Population Growth and Crime Reduction." *Journal of Quantitative Criminology* 10(2): 109–40.
- Mastrobuoni, Giovanni, and David Rivers. 2016. *Criminal Discount Factors and Deterrence*. SSRN Electronic Journal.
- McCullough, Jolie. 2021. "Texas Bill to Require Cash Bail for Those Accused of Violent Crimes Becomes Law." *Texas Tribune*, September 13. <https://www.texastribune.org/2021/09/03/texas-bail-legislation-abbott/>.
- Mitchell, Shira, Eric Potash, Solon Barocas, Alexander D'Amour, and Kristian Lum. 2021. "Algorithmic Fairness: Choices, Assumptions, and Definitions." *Annual Review of Statistics and Its Application* 8: 141–63.
- Mustard, David B. 2001. "Racial, Ethnic, and Gender Disparities in Sentencing: Evidence from the US Federal Courts." *Journal of Law and Economics* 44(1): 285–314.
- Nagin, Daniel S. 2013. "Deterrence in the 21st Century: A Review of the Evidence." *Crime and Justice* 42(1): 199–263.
- NASEM (National Academies of Sciences, Engineering, and Medicine). 2018. *Proactive Policing: Effects on Crime and Communities*. The National Academies Press.
- NASEM (National Academies of Sciences, Engineering, and Medicine). 2021. *The Science of Implicit Bias: Implications for Law and Policy; Proceedings of a Workshop-in-Brief*. The National Academies Press.
- NASEM (National Academies of Sciences, Engineering, and Medicine). 2022. *Reducing Racial Inequality in Crime and Justice: Science, Practice, and Policy*. The National Academies Press.
- Neal, Derek, and Rick, Armin. 2016. "The Prison Boom and Sentencing Policy." *Journal of Legal Studies* 45(1): 1–41.
- Neal, Derek, and Armin Rick. Forthcoming. "The Role of Policy in Prison Growth and Decline." In *Oxford Handbook of Sentencing*, edited by Michael T. Light and Ryan D. King. Oxford University Press.
- Oudekerk, Barbara, and Danielle Kaebler. 2021. "Probation and Parole in the United States, 2019." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/probation-and-parole-united-states-2019>.
- Ouss, Aurelie, and Megan Stevenson. 2023. "Does Cash Bail Deter Misconduct?" *American Economic Journal: Applied Economics* 15(3): 150–82.
- Owens, Emily G. 2009. "More Time, Less Crime? Estimating the Incapacitative Effect of Sentence Enhancements." *Journal of Law and Economics* 52(3): 551–79.
- Paulus, Jessica K., and David M. Kent. 2020. "Predictably Unequal: Understanding and Addressing Concerns That Algorithmic Clinical Prediction May Increase Health Disparities." *npj Digital Medicine* 3(1): 99.
- Pfaff, John. 2017. *Locked In: The True Causes of Mass Incarceration—and How to Achieve Real Reform*. Basic Books.
- Phelps, Michelle S. 2020. "Mass Probation from Micro to Macro: Tracing the Expansion and Consequences of Community Supervision." *Annual Review of Criminology* 3: 261–79.
- Rachlinski, Jeffrey J., Sheri Lynn Johnson, Andrew J. Wistrich, and Chris Guthrie. 2008. "Does Unconscious Racial Bias Affect Trial Judges." *Notre Dame Law Review* 84(3): 1195–246.
- Raphael, Steven, and Sandra V. Roza. 2019. "Racial Disparities in the Acquisition of Juvenile Arrest Records." *Journal of Labor Economics* 37(S1): S125–S159.
- Raphael, Steven, and Michael A. Stoll. 2013. *Why Are So Many Americans in Prison?* Russell Sage Foundation.
- Reaves, Brian A. 2013. "Felony Defendants in Large Urban Counties, 2009—Statistical Tables." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/felony-defendants-large-urban-counties-2009-statistical-tables>.
- Rehavi, M. Marit, and Sonja B. Starr. 2014. "Racial Disparity in Federal Criminal Sentences." *Journal of Political Economy* 122(6): 1320–54.
- Reitz, Kevin R. 2019. "Prison-Release Reform and American Decarceration." *Minnesota Law Review* 104: 2741–79.
- Reitz, Kevin R., and Cecelia M. Klingele. 2019.

- 'Model Penal Code: Sentencing—Workable Limits on Mass Punishment.'" *Crime and Justice* 48(1): 255–311.
- Reitz, Kevin R., and Edward E. Rhine. 2020. "Parole Release and Supervision: Critical Drivers of American Prison Policy." *Annual Review of Criminology* 3: 281–98.
- Ridgeway, Greg, Ruth A. Moyer, and Shawn D. Bushway. 2020. "Sentencing Scorecards: Reducing Racial Disparities in Prison Sentences at Their Source." *Criminology & Public Policy* 19(4): 1113–38.
- Rose, Evan K. 2021. "Who Gets a Second Chance? Effectiveness and Equity in Supervision of Criminal Offenders." *Quarterly Journal of Economics* 136(2): 1199–253.
- Sakoda, Ryan T. 2023. *Abolish or Reform? An Analysis of Post-Release Supervision for Low-Level Offenders*. SSRN Electronic Journal.
- Sklansky, David Alan. 2018. "The Problems with Prosecutors." *Annual Review of Criminology* 1: 451–69.
- Snyder, Howard N. 2011. "Arrest in the United States, 1980–2009." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/arrest-united-states-1980-2009>.
- Spelman, William. 1993. *Criminal Incapacitation*. Springer Science & Business Media.
- Spelman, William. 2000. "What Recent Studies Do (and Don't) Tell Us About Imprisonment and Crime." *Crime and Justice* 27: 419–94.
- Spohn, Cassia. 2015. "Evolution of Sentencing Research." *Criminology & Public Policy* 14(2): 225–32.
- Topaz, Chad M., Shaoyang Ning, Maria-Veronica Cicocanel, and Shawn D. Bushway. 2023. "Federal Criminal Sentencing: Race-Based Disparate Impact and Differential Treatment in Judicial Districts." *Humanities and Social Sciences Communications* 10(1): 1–10.
- Vollaard, Ben. 2013. "Preventing Crime Through Selective Incapacitation." *Economic Journal* 123(567): 262–84.
- Zeng, Zhen, and Todd D. Minton. 2021. "Jail Inmates in 2019." Bureau of Justice Statistics. <https://bjs.ojp.gov/library/publications/jail-inmates-2019>.