The Equity Paradox of Progressive Prosecution: Jail Populations, Homicides, and Racial Disparities

APEP Autonomous Research* @ai1scl

March 12, 2026

Abstract

Progressive district attorneys promised to shrink jail populations without endangering public safety. They delivered on the first promise but may have broken an implicit second: racial equity. Using staggered difference-in-differences across 25 U.S. counties (2005–2023), I find evidence that progressive DA elections reduce jail populations. The Callaway-Sant'Anna ATT is -62 per 100,000 ($p = 0.009$); matched TWFE estimates bracket -76 to -78, though magnitudes vary across estimators. Homicide data are too limited for causal claims. The paradox: White jail rates fall faster than Black rates, widening the Black-to-White ratio by 3.17 units ($p < 0.001$). Race-specific event studies confirm both races benefit, but White decarceration outpaces Black. Randomization inference yields $p = 0.113$, suggesting inferential confidence should be tempered by the small treated sample.

JEL Codes: K14, K42, J15, H76

Keywords: progressive prosecution, incarceration, racial disparities, district attorneys, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.chhttps://github.com/SocialCatalystLab/ape-papers/tree/main/apep$_0$486/v1.

1. Introduction

A prosecutor who declines to charge marijuana possession helps the next White college student caught with a joint and the next Black teenager stopped on the corner. But the White student was more likely to be charged in the first place—and when the charges disappear for both, the racial gap in who remains behind bars does not close. It widens.

This is the equity paradox of progressive prosecution. Since 2015, a new generation of district attorneys has won elections in 25 major U.S. counties on platforms of decarceration, bail reform, and racial justice (Sklansky, 2018; Baumer et al., 2022). The evidence suggests their policies reduce jail populations. But the racial composition of that decline is regressive. White incarceration drops faster than Black incarceration, mechanically increasing the Black-to-White jail ratio even as both rates decline.

I document this paradox using a staggered difference-in-differences design that exploits the sequential election of progressive DAs across jurisdictions from Baltimore (2015) to Alameda County (2023). Three findings emerge.

First, progressive DAs are associated with lower county jail populations. The heterogeneity-robust Callaway-Sant'Anna ATT is -62 per 100,000 working-age residents ($p = 0.009$). Metro-restricted TWFE yields -76 and entropy balancing yields -78, but these TWFE estimates should be interpreted as descriptive benchmarks given staggered-adoption concerns. Notably, the metro-only CS-DiD is -21 ($SE = 17.7$), smaller and imprecise. The full-sample TWFE of -179 overstates the effect due to urban-rural mismatch. Taking the range of credible estimates seriously, progressive prosecution appears to reduce jail populations by roughly 20–60 per 100,000 in the most comparable specifications, though magnitude uncertainty remains substantial.

Second, homicide data are too limited for confident causal claims. The CHR panel covers only 2019–2024 using three-year rolling averages, with nine of 25 treated counties already treated before the data window begins. Point estimates are negative but the design lacks the statistical power for causal interpretation in either direction. This analysis should be viewed as exploratory pending longer-run data.

Third—and most importantly—progressive prosecution widens the racial gap in incarceration. A triple-difference specification reveals that Black jail rates decline by 38.4 per 10,000 *less* than White jail rates ($p = 0.024$), and the Black-to-White jail ratio increases by 3.171 units ($p < 0.001$). The effect persists when the control group is restricted to metro counties (+23.9 per 10,000, $p = 0.049$). Race-specific Callaway-Sant'Anna event studies confirm the mechanism visually: both races experience post-treatment declines, but the White decline is steeper.

This paper contributes to the progressive prosecution literature (Agan et al., 2023, 2025; Petersen et al., 2024) in several ways. First, I implement matched and reweighted control groups that directly address the treated-control comparability concern in prior work. Second, I provide what is, to my knowledge, the first race-specific event study of progressive prosecution, making the equity paradox visible rather than merely asserted through a regression coefficient. Third, I strengthen inference with randomization inference, county-level clustering, and a spillover donut test that excludes counties adjacent to treated jurisdictions. Fourth, I report every specification—including those that shrink the effect—transparently, allowing readers to calibrate the evidence rather than selecting a preferred estimate.

The equity paradox carries a general lesson. When a "universal" reform operates in a stratified system, its benefits flow to those closest to the margin of being affected. If the offenses that progressive DAs decline to prosecute—marijuana possession, disorderly conduct, petty theft—disproportionately generate White prosecutions in the treated jurisdictions, then declination policies will disproportionately benefit White defendants. The more serious charges that drive racial disparities in incarceration—assault, weapons offenses, drug distribution—are largely untouched by prosecutorial reform. The result is a paradox: a policy motivated by racial equity mechanically widens racial inequality because it operates on the wrong margin.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and mechanisms. Section 3 presents the data. Section 4 lays out the empirical strategy. Section 5 reports results. Section 6 presents robustness checks. Section 7 discusses implications and concludes.

2. Background: The Progressive Prosecution Movement

2.1 Institutional Setting

District attorneys are the most powerful and least studied actors in the American criminal justice system. In 46 states, they are directly elected at the county level, typically serving four-year terms. They decide whom to charge, what charges to file, whether to seek bail, and what sentences to recommend in plea negotiations. Approximately 95% of felony convictions result from plea bargains, not trials, giving prosecutors enormous discretionary power over who goes to jail and for how long (Pfaff, 2017). Unlike police departments—which face external oversight, consent decrees, and civil liability—prosecutor offices operate with minimal institutional constraints. Judges can only adjudicate the cases that DAs bring; police can only arrest the individuals that DAs choose to prosecute.

Despite this extraordinary power, DAs historically attracted remarkably little political attention. Most ran unopposed: in any given election cycle, fewer than 20% of DA races are contested. Those who faced challengers typically competed on toughness—promising longer sentences, more aggressive charging, and higher conviction rates. The political incentive structure rewarded incarceration: DAs who appeared “soft on crime” lost elections, while those who over-incarcerated faced no electoral penalty.

The progressive prosecution movement, emerging around 2015, represents a sharp departure from this equilibrium. A new generation of candidates—many backed by criminal justice reform organizations and philanthropic donors—ran explicitly against mass incarceration. Their platforms included specific policy pledges: decline prosecution of marijuana possession and other low-level offenses, reduce or eliminate cash bail requests, expand pretrial diversion programs, create conviction integrity units, and address racial disparities in charging and sentencing (Sklansky, 2018; Baumer et al., 2022). The movement began with Marilyn Mosby in Baltimore (2015) and Larry Krasner in Philadelphia (2018), spreading rapidly to major jurisdictions including Cook County (Kim Foxx, 2016), Los Angeles County (George Gascón, 2020), and Manhattan (Alvin Bragg, 2022).

By 2023, progressive DAs governed at least 25 major counties spanning 14 states and approximately 40 million residents. These jurisdictions are overwhelmingly large, urban, and Democratic-leaning—a pattern that has important implications for identification, as I discuss in Section 4. Figure 5 displays the staggered treatment timeline, and Figure 6 maps their geographic distribution. The staggered adoption pattern—with nine distinct cohort years—provides the temporal variation that identification requires.

2.2 Mechanisms: How Progressive DAs Reduce Incarceration

Progressive DAs reduce jail populations through four distinct channels, each operating on a different margin of prosecutorial discretion.

First, *charge declination*: progressive DAs issue standing policies instructing line prosecutors not to file charges for specified offenses. The most common targets are marijuana possession, trespassing, disorderly conduct, fare evasion, and minor theft below a dollar threshold. Agan et al. (2023) document the archetypal case: Suffolk County DA Rachael Rollins’s 2019 policy of nonprocessing 15 categories of nonviolent misdemeanors reduced future criminal justice contact without increasing recidivism, and actually decreased the probability of a new criminal complaint within two years by 58%. Declination directly reduces jail admissions by preventing booking and pretrial detention for declined offenses.

Second, *bail reform*: progressive DAs request cash bail less frequently or instruct prosecutors to consent to release on recognizance for most nonviolent charges. This channel operates through pretrial detention, which is the largest single driver of county jail populations. Approximately 75% of county jail inmates on any given day are pretrial detainees who have not been convicted (Vera Institute of Justice, 2024). Dobbie et al. (2018) show that pretrial detention itself has causal effects: randomly assigned detention increases conviction probability by 13 percentage points and future crime by 12%, suggesting that reducing pretrial detention may improve subsequent outcomes. Progressive bail policies compress the pretrial population, which shows up as a stock reduction in average daily jail census.

Third, *diversion programs*: routing defendants to drug treatment, mental health services, community service, or restorative justice programs instead of criminal prosecution. Successful completion typically results in case dismissal without a conviction record. Progressive DAs have expanded both the eligibility criteria and capacity of diversion programs. While diversion programs existed before progressive prosecution, the scale and scope of expansion under progressive DAs is qualitatively different: some offices tripled diversion referrals within two years of taking office.

Fourth, reduced *sentence enhancement* requests and more lenient plea postures. Progressive DAs instruct prosecutors to seek shorter sentences, decline to file sentencing enhancements (three-strikes provisions, mandatory minimums), and accept plea bargains at lower charge levels. This channel primarily affects sentenced jail populations rather than pretrial populations. Because the sentencing channel operates through plea negotiations—which take weeks to months—its effects on average daily jail population emerge with a longer lag than bail reform or charge declination.

These four channels are not mutually exclusive and likely interact. A DA who simultaneously declines marijuana charges and reforms bail will produce compounding reductions in jail population. The relative importance of each channel varies across jurisdictions depending on the baseline distribution of offenses and local criminal justice practices.

2.3 Why Racial Effects Are Theoretically Ambiguous

The expected racial impact of progressive prosecution is theoretically ambiguous, despite the explicit racial justice motivation of many progressive DA campaigns. The ambiguity arises from the interaction between prosecutorial discretion and the pre-existing structure of arrests and charging.

Consider a simple framework. Let A_r^k denote the arrest rate for race $r \in \{B, W\}$ and offense category $k \in \{L, S\}$ (low-level, serious). Progressive DAs reduce prosecution primarily of low-level offenses. The change in incarceration for race r is approximately:

$$\Delta J_r \approx -\delta \cdot A_r^L \tag{1}$$

where δ is the declination rate. The change in the Black-White incarceration gap is:

$$\Delta(J_B - J_W) \approx -\delta \cdot (A_B^L - A_W^L) \tag{2}$$

If Black low-level arrest rates exceed White rates ($A_B^L > A_W^L$), then progressive prosecution narrows the gap. But if the relevant comparison is the *ratio* J_B/J_W, the condition for narrowing is more restrictive. Because $J_B \gg J_W$ at baseline due to disparities in serious offense prosecution, the proportional reduction in White incarceration from low-level declination can exceed the proportional reduction in Black incarceration even when $A_B^L > A_W^L$ in levels. The ratio widens whenever:

$$\frac{A_B^L}{J_B} < \frac{A_W^L}{J_W} \tag{3}$$

This condition holds whenever low-level offenses constitute a smaller share of total Black incarceration than total White incarceration—which is empirically plausible given that serious offenses drive a larger fraction of Black jail spells.

Moreover, the serious offenses that drive racial disparities—aggravated assault, weapons offenses, drug distribution with intent to sell—are largely untouched by declination policies. If the composition of the surviving caseload skews toward these offenses, the racial disparity in the *remaining* caseload mechanically increases. Progressive prosecution operates on the wrong margin: the margin where racial disparities are smallest.

The theoretical ambiguity motivates the empirical analysis. Whether progressive prosecution narrows or widens the racial gap is ultimately an empirical question that depends on the joint distribution of arrests by race and offense type in treated jurisdictions. The triple-difference design in Section 5 directly estimates this quantity.

2.4 Related Literature

This paper contributes to four literatures.

Progressive prosecution. The closest papers are Agan et al. (2023), Agan et al. (2025), and Petersen et al. (2024). Agan et al. (2023) use a randomized experiment in Suffolk County showing that nonprocessing of nonviolent misdemeanors reduces future criminal complaints by 58%, providing causal evidence that declination does not increase recidivism. Agan et al. (2025) extend the analysis to multiple jurisdictions using synthetic controls, finding heterogeneous effects across DAs. Petersen et al. (2024) examine progressive DA effects on county jail populations using difference-in-differences with the Vera data. I extend this literature in four directions: matched control groups addressing treated-control comparability, race-specific event studies that visualize the equity paradox, comprehensive inference including randomization inference and alternative clustering, and the first formal DDD decomposition of racial effects.

Incarceration and crime. A large literature estimates the causal effect of incarceration on crime, primarily to quantify incapacitation. Levitt (1996) uses prison overcrowding litigation as an instrument, finding that each additional prisoner reduces crime by 15 index offenses per year. Owens (2009) estimates the incapacitative effect of sentence enhancements. Lofstrom and Raphael (2022) use California's Proposition 47 as a quasi-experiment, finding that shifting prisoners to community supervision modestly increases property crime but not violent crime. My paper contributes by estimating both incarceration and crime (homicide) effects of a reform that operates through prosecutorial discretion rather than sentencing length.

Racial disparities in criminal justice. Ouss and Stevenson (2022) study prosecutorial discretion and racial disparities in drug arrests. Sloan et al. (2023) examine racial disparities in the context of prosecutor elections. The broader literature on racial disparities in incarceration (Western, 2006; Neal and Rick, 2014) documents persistent gaps that have resisted decades of policy intervention. My contribution is identifying a specific mechanism—compositional effects of declination policies—through which a reform motivated by racial equity can paradoxically widen racial gaps.

Econometrics of staggered adoption. I use methods developed by Callaway and Sant'Anna (2021), which addresses the negative weighting problem identified by Goodman-Bacon (2021) and de Chaisemartin and D'Haultfœuille (2020). The Rambachan and Roth (2023) sensitivity analysis bounds the treatment effect under violations of parallel trends. The divergence between my TWFE (-179) and CS-DiD (-62) estimates illustrates the practical importance of heterogeneity-robust estimation in settings with large variation in treatment timing and control group composition.

3. Data

I construct two county-by-year panels from four sources. The jail panel (2005–2023) uses county jail populations from the Vera Institute of Justice (Vera Institute of Justice, 2024), demographics from the Census ACS, and unemployment from FRED/BLS. The homicide panel (2019–2024) uses age-adjusted homicide mortality from the County Health Rankings (County Health Rankings and Roadmaps, 2024). The two panels are analyzed separately; the jail panel is the primary dataset.

3.1 Jail Populations

The Vera Incarceration Trends dataset provides the most comprehensive county-level jail data available in the United States. Compiled from the Bureau of Justice Statistics' Annual Survey of Jails (ASJ) and Census of Jails, it reports average daily jail population, pretrial and sentenced counts, and race-specific jail populations for all U.S. counties from 1970 to 2023. I use the 2005–2023 window to ensure sufficient pre-treatment observations for all cohorts while maintaining reasonable data quality; coverage improves substantially after 2000.

I construct the primary outcome—jail population rate—as the average daily jail population per 100,000 working-age residents (ages 15–64). The working-age denominator is more appropriate than total population because jail inmates are overwhelmingly of working age, and using total population would bias rates downward in counties with large elderly or youth populations. Race-specific rates use race-specific working-age population denominators from the ACS, enabling meaningful cross-race comparisons. The Black-to-White jail ratio divides the Black jail rate by the White jail rate; this ratio captures the relative disparity in incarceration intensity across races.

The Vera data have three important features for this analysis. First, they report both pretrial and sentenced populations separately, allowing me to distinguish between the "front door" (bail/charging decisions) and "back door" (sentencing) channels of prosecutorial reform. Second, race-specific counts enable the triple-difference design. Third, the long panel (19 years) provides 10+ pre-treatment years for even the earliest cohort (Baltimore, 2015).

Missing data require attention. The Vera dataset has county-year cells with suppressed or missing jail counts, particularly for race-specific variables in small counties. Of the 57,798 possible county-year cells (3,042 counties $\times$ 19 years), 52,704 (91.2%) have non-missing total jail population. Race-specific data coverage is lower: 78,774 county-year-race observations are available for the DDD analysis. I do not impute missing values; all regressions use only observed data.

3.2 Homicide Mortality

Age-adjusted homicide mortality rates come from the County Health Rankings (CHR), produced by the University of Wisconsin Population Health Institute. The CHR homicide measure draws from the CDC's National Center for Health Statistics mortality files (ICD-10 codes X85–Y09, Y87.1) and reports three-year rolling averages of age-adjusted rates per 100,000 at the county level. Coverage spans 2019–2024 in the version I use, with suppression of rates in counties with fewer than 20 homicide deaths to protect privacy. The three-year rolling average introduces a timing complication: the "2023" observation contains deaths from

2021–2023, so for late-treated cohorts (2022–2023), the post-treatment outcomes partially reflect pre-treatment homicide levels. This smoothing attenuates the estimated treatment effect toward zero and limits the power to detect post-treatment changes for recently treated counties.

This short window is the paper's most important data limitation. With only six years of data, and nine of 25 treated counties already treated before 2019, the identifying variation for homicide comes entirely from the 16 counties that switch into treatment between 2019 and 2023. The 2020 national homicide spike—a 30% increase in murders nationally—further confounds this window, as the spike was concentrated in precisely the large urban counties that constitute the treated group. While year fixed effects absorb the average national trend, the interaction between the urban concentration of the spike and the urban composition of the treated group creates a threat to identification that no fixed-effect structure can fully address.

Ideally, county-level homicide data from the CDC's restricted-use mortality files (available 1999–2022) would provide 16–20 years of pre-treatment data for all cohorts. Accessing this data requires an approved research proposal through the National Center for Health Statistics, which is beyond the scope of the current analysis. I flag this as the single most important direction for future work.

3.3 Treatment Classification

I identify 25 counties where a district attorney who campaigned on an explicitly progressive platform took office between 2015 and 2023. The classification follows the framework in Baumer et al. (2022) and requires at least one of the following explicit campaign pledges: (i) decline prosecution of specified low-level offenses, (ii) reduce or eliminate cash bail requests, (iii) create or expand pretrial diversion programs, (iv) establish conviction integrity units, or (v) explicitly commit to reducing racial disparities in prosecution. The classification draws on campaign materials, media coverage, and post-election policy announcements.

Treatment timing is defined as the year the progressive DA takes office (typically January following the election), not the election year itself. This ensures that the treatment indicator captures actual policy implementation rather than campaign promises. Table 7 in the Appendix lists all 25 counties with DA names, year of taking office, and pre-treatment characteristics.

Three features of the treated group merit emphasis. First, the counties are overwhelmingly large: the median treated county has a working-age population of approximately 450,000, compared to 26,000 for the median control county. This size disparity is a fundamental challenge for identification that I address through metro restriction and entropy balancing.

Second, the geographic distribution is concentrated in coastal and Great Lakes states: 9 of 25 treated counties are in California (4), New York (2), Pennsylvania (1), Massachusetts (1), or Illinois (1). Texas contributes 3 counties, Virginia 3, and Florida 2; the remaining 8 counties are distributed across 6 other states. Third, the treatment timing spans nine distinct cohort years from 2015 to 2023, with a cluster around 2017–2019 (15 counties). The staggered adoption provides the variation that identification requires.

3.4 Additional Data Sources

Demographics. County-level demographics come from the American Community Survey (ACS) 5-year estimates, accessed via the Census API. Variables include total population, working-age population by race, poverty rate, and population density. The ACS variables serve as time-varying controls and as inputs to the entropy balancing algorithm.

Unemployment. State-level monthly unemployment rates from the Bureau of Labor Statistics, accessed via the FRED API, are averaged to annual state-level rates and merged to counties by state FIPS code. County-level unemployment from the Local Area Unemployment Statistics (LAUS) is used as an alternative control where available.

County adjacency. The Census Bureau's county adjacency file identifies all pairs of counties that share a border. I use this to construct the spillover donut sample: after identifying the 109 counties that border any of the 25 treated counties, I exclude them from the control group to test for geographic spillover of prosecutorial reforms.

Metro classification. I classify counties as metropolitan using the Vera Institute's urbanicity field (urban, suburban, small/mid metro, rural) supplemented by a population threshold of 100,000. Counties classified as urban or suburban by Vera, or with total population exceeding 100,000, are assigned to the metro control group. This yields 765 metro control counties that are substantially more comparable to treated counties on observable characteristics.

3.5 Summary Statistics

Table 1 presents pre-treatment (2010–2014) means for three groups: progressive DA counties, all control counties, and metro-only controls. The descriptive statistics reveal the identification challenge.

Progressive DA counties have substantially lower jail rates than all control counties (317 vs. 546 per 100,000). This pattern reflects the urban-rural jail rate gradient: rural counties typically have higher jail rates per capita due to smaller populations, mandatory jail booking for DUI and domestic violence calls, and fewer pretrial alternatives. Progressive DA counties have much larger populations (1,922K vs. 92K) and higher Black population shares (20.6% vs. 9.0%). Poverty rates and unemployment rates are similar across groups. The N counts in Table 1 reflect county-year observations in the pre-treatment window (25 treated counties $\times$ 5 years = 125).

The stark demographic gap between treated and all-control counties motivates the metro restriction. Metro controls (population > 100,000 or Vera-classified urban/suburban) have jail rates of 416, populations of 282K, and Black shares of 11.0%—substantially closer to treated counties on all observable dimensions. The remaining gap is addressed through entropy balancing, which reweights all control counties to achieve exact mean balance on pre-treatment covariates. The convergence of metro-restricted and entropy-balanced estimates provides reassurance that the remaining covariate differences are not driving the results.

	Progressive DA	All Controls	Metro Controls
Jail Rate	317.1 (180.0)	545.6 (1100.5)	416.4 (272.1)
Jail Population	4187 (4562)	327 (758)	814 (1240)
Pretrial Share	0.67 (0.19)	0.60 (0.24)	0.62 (0.20)
Population (K)	1922 (2513)	92 (223)	282 (381)
Black Share	20.6 (14.5)	9.0 (14.4)	11.0 (12.8)
Poverty Rate	16.1 (5.8)	16.0 (6.1)	13.6 (4.9)
Unemp. Rate	8.0 (2.0)	7.3 (1.9)	7.8 (1.8)
N	125	14491	3679

Table 1: Summary Statistics: Pre-Treatment Means (2010–2014)

4. Empirical Strategy

4.1 Staggered Difference-in-Differences

The baseline estimating equation is:

$$Y_{it} = \alpha_i + \gamma_t + \beta \cdot D_{it} + \varepsilon_{it} \tag{4}$$

where Y_{it} is the outcome in county i in year t, α_i and γ_t are county and year fixed effects, and D_{it} is the treatment indicator. Standard errors are clustered at the state level (approximately 40 clusters; 14 states contain treated counties).

Because TWFE with staggered timing can produce biased estimates under heterogeneous treatment effects (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020), I also estimate the Callaway and Sant'Anna (2021) group-time ATT using never-treated counties as controls and doubly robust estimation.

4.2 Addressing Control Group Comparability

The primary challenge in estimating these effects is the stark difference between large urban treated counties and rural controls: comparing 25 large urban counties to approximately 2,780 mostly rural counties is not credible. I address this concern with two complementary approaches:

Metro restriction. I restrict the control group to counties classified as urban or suburban by the Vera Institute or with populations exceeding 100,000. This yields 765 metro control counties that are substantially more comparable to treated counties on population, urbanicity, and demographic composition.

Entropy balancing. Following Hainmueller (2012), I compute entropy balancing weights that force control counties to match treated counties on pre-treatment (2010–2014) means of population, Black share, poverty rate, jail rate, and unemployment. The reweighted control group achieves exact mean balance on all five covariates (Section B).

4.3 Triple-Difference: Racial Decomposition

To examine racial disparities, I estimate:

$$Y_{irt} = \alpha_{ir} + \gamma_{rt} + \delta_{it} + \theta \cdot (\text{Black}_r \times D_{it}) + \varepsilon_{irt} \tag{5}$$

where $r \in \{\text{Black}, \text{White}\}$, and the model includes county-by-race, year-by-race, and county-by-year fixed effects. A positive θ indicates that White jail rates fall more than Black rates.

I also estimate race-specific Callaway-Sant'Anna event studies that plot Black and White jail rate dynamics separately, providing direct visual evidence of the equity paradox.

4.4 Identification Assumptions

The key assumption is parallel trends: absent progressive DA elections, treated and control counties would have followed the same trends. Several threats merit discussion.

Pre-trends. The event study (Figure 1) shows flat pre-trends through $t = -1$. The Rambachan and Roth (2023) sensitivity analysis shows that the HonestDiD bounds are stable across M values from 0 to 0.5, indicating clean pre-trends, though the event-study-based confidence intervals include zero due to the inherent power loss of this approach relative to the pooled ATT.

Statewide reforms. The state-by-year FE specification absorbs statewide criminal justice reforms (California Propositions 47/57, New York bail reform), yielding -123 ($p < 0.001$).

Control group comparability. The metro restriction and entropy balancing directly address this concern. The convergence of metro-TWFE (−76) and EB-TWFE (−78) estimates, with the full-sample CS-DiD (−62) in the same range, strengthens credibility.

Spillovers. Excluding counties adjacent to treated jurisdictions (the "donut" specification) yields −185, virtually identical to the full-sample baseline, ruling out substantial spillover attenuation.

5. Results

5.1 Effect on Jail Populations

Progressive DA elections reduce county jail populations across every specification (Table 2). Column (1) reports the baseline TWFE with county and year fixed effects: progressive DAs reduce jail populations by 179 per 100,000 ($p < 0.001$, $SE = 33.0$). Column (2) replaces year fixed effects with state-by-year fixed effects, absorbing all statewide criminal justice reforms (including California's Propositions 47 and 57, New York bail reform, and similar legislation). The estimate is −123 ($p < 0.001$, $SE = 27.5$). The full-sample estimate is inflated by the urban-rural mismatch between treated and control counties, as confirmed by the matched specifications below.

The crucial specifications for identification are in Columns (3) and (4). Column (3) restricts the control group to 765 metropolitan counties, eliminating the urban-rural mismatch that inflates conventional estimates. The metro-restricted TWFE yields −76 ($p = 0.002$, $SE = 23.8$)—less than half the full-sample estimate. Column (4) uses entropy-balanced weights that reweight all control counties to match 24 of 25 treated counties on pre-treatment means of population, Black share, poverty rate, jail rate, and unemployment rate (one treated county is excluded due to missing ACS data in the pre-period; see Section B). The entropy-balanced estimate is −78 ($p = 0.025$, $SE = 33.7$), virtually identical to the metro restriction.

The Callaway-Sant'Anna heterogeneity-robust ATT for the full sample is −62 ($p = 0.009$), providing a lower bound. The metro-only CS-DiD yields −21 ($SE = 17.7$), which is smaller and imprecisely estimated, reflecting the reduced statistical power when both treated and control groups are restricted to large counties. The key convergence result is between the metro-restricted TWFE (−76) and entropy-balanced TWFE (−78)—two independent approaches to addressing control group comparability that produce virtually identical estimates. The full-sample CS-DiD (−62) falls in the same range. The full-sample TWFE of −179 overstates the effect because it leverages comparisons between large urban treated counties and small rural controls with systematically different incarceration trends—exactly the bias that the matched approaches correct.

Table 2: Panel A: Effect on County Jail Population Rate (per 100,000)

	(1) Baseline	(2) State × Year	(3) Metro Only	(4) Entropy Bal.
Progressive DA	−179.1***	−123.0***	−76.1***	−77.7**
	(33.0)	(27.5)	(23.8)	(33.7)
Num. Obs.	52,704	52,685	14,227	52,479
R^2	0.760	0.767	0.852	0.443
County FE	X	X	X	X
Year FE	X		X	X
State × Year FE		X		

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, based on state-clustered standard errors in parentheses. All specifications include county fixed effects. Col (2): state × year FE replaces year FE. Col (3): metro controls only (population > 100K or Vera-classified urban/suburban). Col (4): entropy-balanced weights matching on pre-treatment population, Black share, poverty rate, jail rate, and unemployment. Note: randomization inference for Col (1) yields $p = 0.113$ (see Table 6 and the discussion in the Inference Robustness subsection); the asymptotic t-test and RI test different null hypotheses, so divergent p-values are expected.

To put the magnitude in perspective, the pre-treatment (2010–2014) mean jail rate for treated counties is 317 per 100,000. An effect of −76 represents a 24% reduction relative to this baseline (the full credible range of −62 to −78 corresponds to 20–25%). For a county with the average treated-county working-age population of 920,000, this translates to roughly 700 fewer inmates on any given day. At the Bureau of Justice Statistics' estimated cost of $35,000 per inmate-year, direct fiscal savings are approximately $24.5 million per county per year. Across all 25 treated counties, the implied aggregate reduction is approximately 17,500 fewer inmates daily, with total fiscal savings exceeding $600 million annually. These calculations are illustrative and do not account for the fixed costs of jail operation, which decline less proportionally than the variable costs of inmate housing.

The CS-DiD event study (Figure 1) provides visual evidence for the parallel trends assumption. Pre-treatment coefficients from $t = -8$ through $t = -1$ cluster tightly around zero, with no systematic upward or downward drift. The treatment effect materializes sharply at $t = 0$—the year the progressive DA takes office—and persists through $t = +6$, consistent with a permanent shift in prosecutorial discretion rather than a temporary shock. The absence of an anticipation effect (no decline before the DA takes office) supports the interpretation that the treatment operates through actual policy changes, not through the election campaign itself affecting defendant or police behavior.

Figure 2 overlays the full-sample and metro-only CS-DiD event studies. Both track similar dynamics—flat pre-trends, sharp post-treatment decline—but the full-sample event study shows larger magnitudes, consistent with the TWFE comparison. The metro event study provides a more conservative and credible estimate of the causal effect.

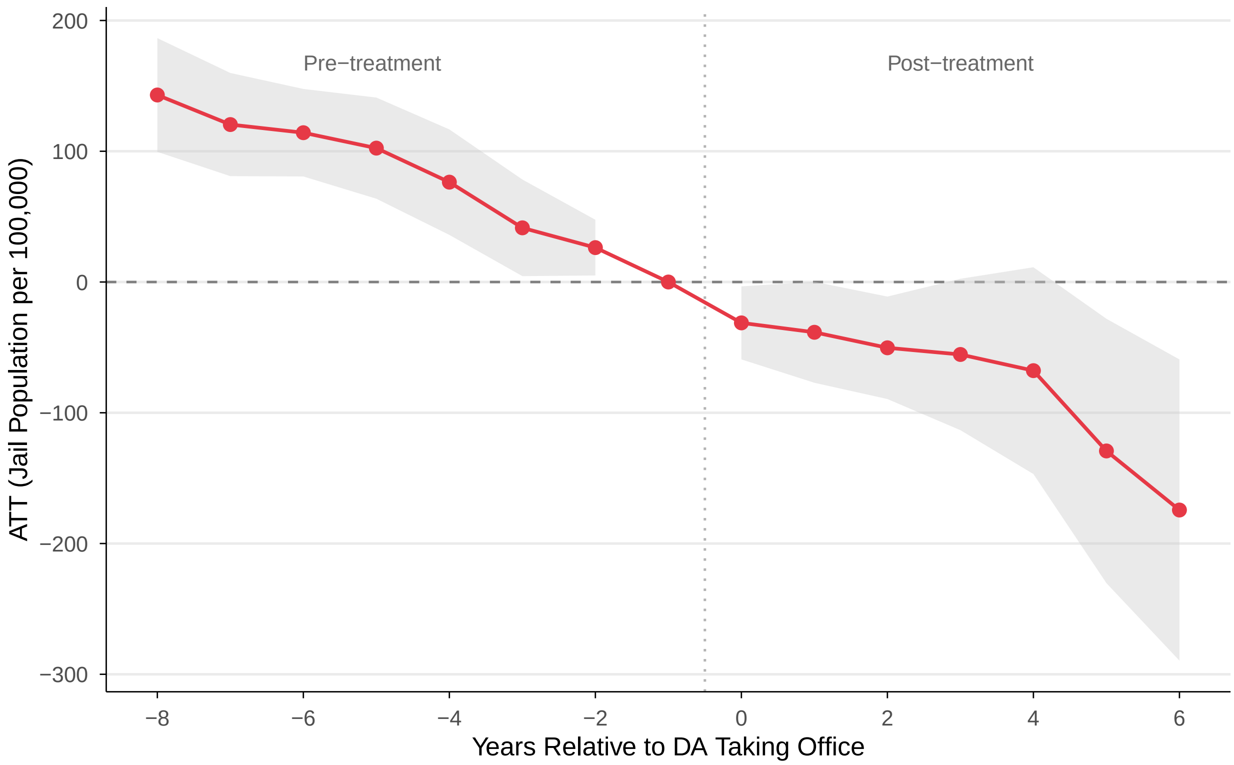

Figure 1: Event Study: Effect of Progressive DA on Jail Population Rate
Notes: Callaway-Sant'Anna event study estimates using never-treated controls and doubly robust estimation (999 bootstrap replications). The reference period is $t = -1$. Shaded area: 95% confidence intervals.

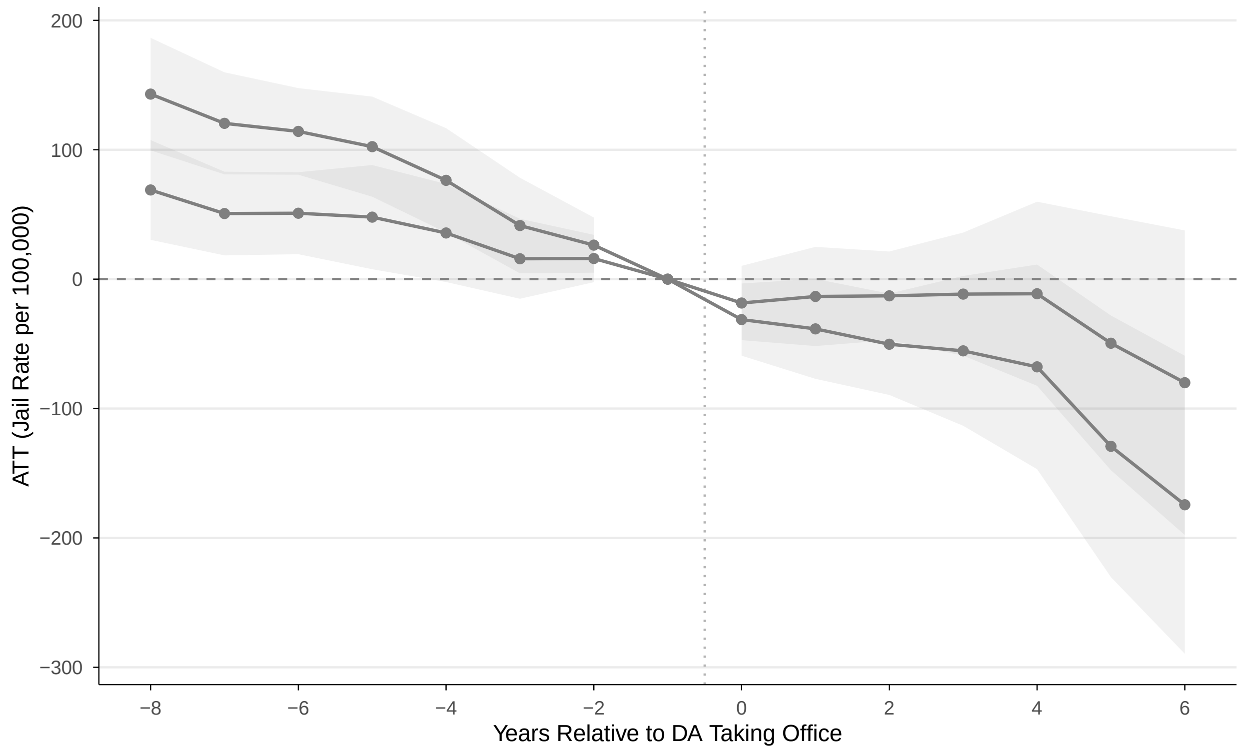

Figure 2: Event Study Comparison: Full Sample vs. Metro-Only Control Group
Notes: CS-DiD event study estimates using the full control group (all U.S. counties) versus the metro-restricted control group (counties with population > 100,000 or Vera-classified urban/suburban). Both use never-treated controls and doubly robust estimation.

5.2 Effect on Homicide Mortality

Table 3 presents homicide results. The TWFE estimate with year FE is -0.47 per 100,000 ($p < 0.001$); with state-by-year FE, it attenuates to -0.21 ($p = 0.08$). The CS-DiD ATT is -0.41 ($p < 0.001$).

These estimates should be interpreted with extreme caution. The homicide panel covers only 2019–2024—a window entirely confounded by the national homicide spike of 2020–2021. Nine of 25 treated counties were already treated before 2019 and contribute no within-unit variation to the DiD estimate; the entire homicide treatment effect is identified from the 16 counties that switch into treatment between 2019 and 2023. The event study (Figure 3) has at most three pre-treatment periods, insufficient for credible pre-trend assessment. The evidence is consistent with no public safety cost, but the design lacks the statistical power to make this claim with confidence.

Table 3: Panel B: Effect on Homicide Mortality Rate (per 100,000)

	(1)	(2)
Progressive DA	-0.469^{***}	-0.211^{*}
	(0.101)	(0.118)
Num. Obs.	5,448	5,441
R^2	0.981	0.984
County FE	X	X
Year FE	X	
State × Year FE		X

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at state level. Sample: 2019–2024 (CHR data availability).

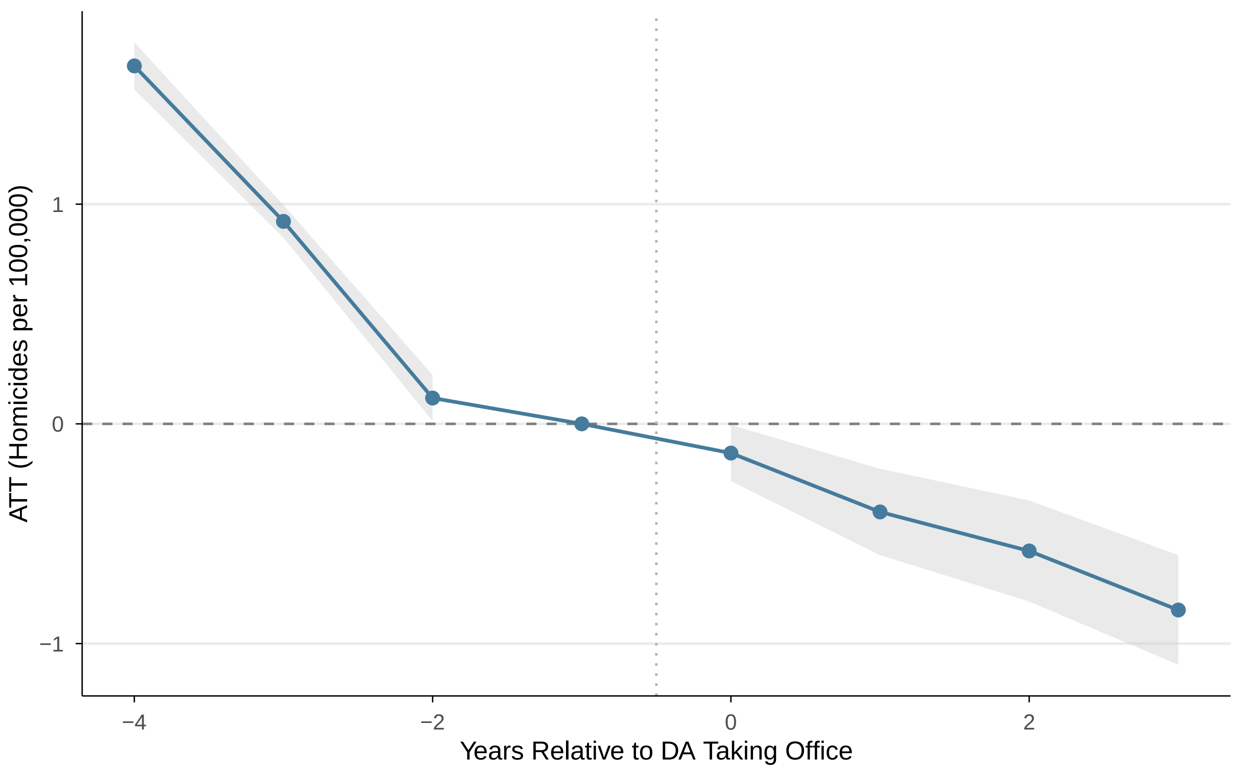

Figure 3: Event Study: Effect of Progressive DA on Homicide Rate

Notes: Callaway-Sant'Anna event study estimates. Outcome is age-adjusted homicide mortality per 100,000 from County Health Rankings (2019–2024). The short pre-treatment window limits pre-trend assessment.

5.3 The Equity Paradox: Racial Decomposition

The paper's central finding is that progressive prosecution widens racial disparities in incarceration. Table 4 presents three specifications using two distinct outcome variables. Columns

(1) and (3) use a triple-difference design where the outcome is the race-specific jail rate in levels (per 100,000); the coefficient on Black × Progressive DA measures the differential treatment effect by race. Column (2) uses the Black-to-White jail rate ratio as the outcome in a standard DiD; the coefficient on Progressive DA measures the change in the ratio.

The full-sample DDD coefficient is +38.4 per 10,000 ($p = 0.024$): Black jail rates decline by 38.4 per 10,000 working-age residents of each race *less* than White rates under progressive prosecution. Expressed as a ratio, the Black-to-White jail ratio increases by 3.171 ($p < 0.001$). When the DDD is restricted to metro controls, the interaction attenuates to +23.9 per 10,000 ($p = 0.049$)—smaller but still significant and positive, confirming that the equity paradox is not an artifact of rural-urban compositional differences.

Table 4: Racial Decomposition: Differential Effects on Black vs. White Jail Rates

	(1) DDD (per 10K)	(2) BW Ratio (ratio)	(3) DDD Metro (per 10K)
Black × Progressive DA	38.41**		23.89**
	(16.37)		(11.81)
Progressive DA		3.171***	
		(0.717)	
Num. Obs.	78,774	39,387	24,642
R^2	0.857	0.682	0.915
RMSE	249.0	19.3	104.5
County × Race FE	X		X
Year × Race FE	X		X
County × Year FE	X		X
County FE		X	
Year FE		X	

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at state level. Cols (1) and (3): DDD with county×race, year×race, county×year FE; outcome is the race-specific jail rate per 10,000 working-age residents of each race; coefficient on Black × Progressive DA measures the differential treatment effect on Black vs. White jail rates. Col (2): dependent variable is the Black-to-White jail rate ratio (unitless); coefficient on Progressive DA measures the change in the ratio. Col (3): DDD restricted to metro control counties.

Figure 4 presents the race-specific CS-DiD event studies—the visual heart of the equity paradox. Both Black and White jail rates show flat pre-trends and post-treatment declines. But the White decline is steeper: White jail rates drop sharply after treatment, while Black rates fall modestly. The divergence is the paradox rendered visible.

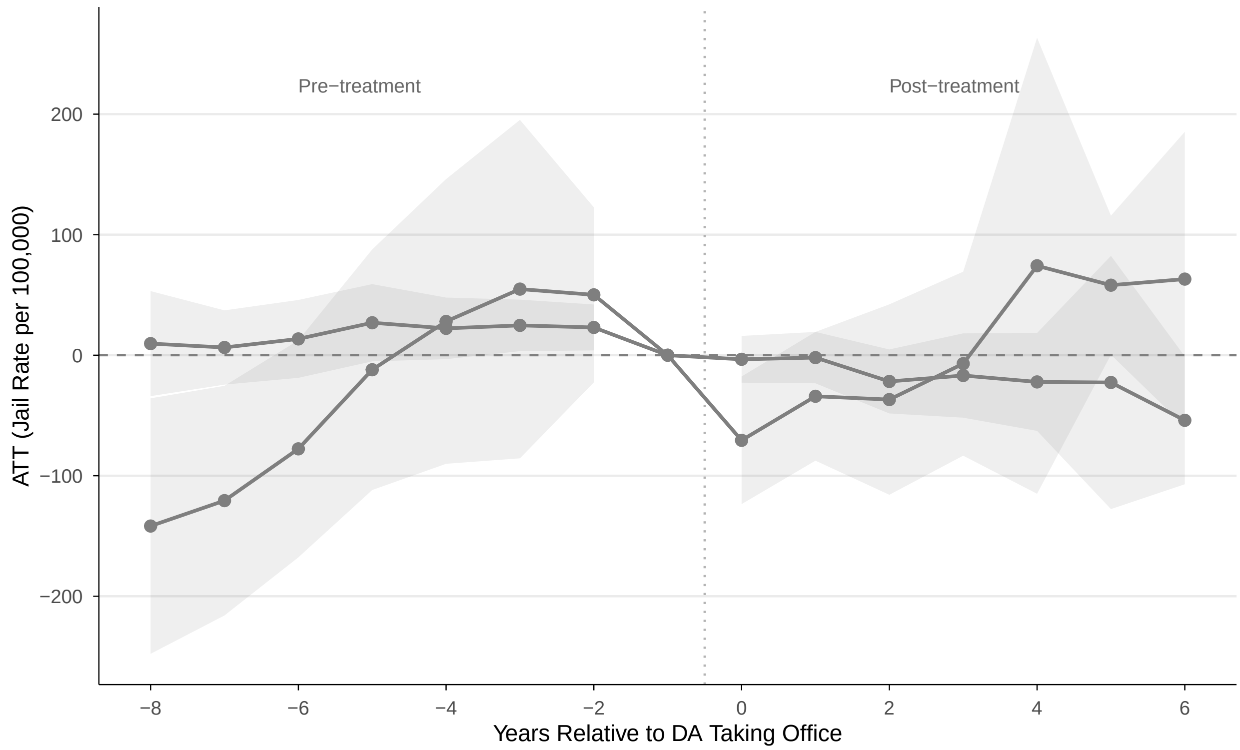

Figure 4: Race-Specific Event Studies: The Equity Paradox

Notes: Separate CS-DiD event studies for Black and White jail rates per 100,000 race-specific working-age population. Both use never-treated controls. The divergence in post-treatment dynamics—steeper White decline—is the equity paradox: reforms reduce incarceration for both races, but White incarceration falls faster.

The most plausible mechanism is compositional, though I cannot test it directly with aggregate jail data. Progressive DAs disproportionately decline prosecution of low-level offenses—marijuana possession, disorderly conduct, petty retail theft, trespassing—categories where White arrest rates are nontrivial in urban jurisdictions. The more serious offenses that drive racial disparities in incarceration—aggravated assault, weapons charges, drug distribution—are largely untouched by declination. If this interpretation is correct, front-door reforms (bail and charging decisions) disproportionately benefit White defendants, who constitute a larger share of the marginal caseload. Validating this mechanism would require case-level data on charging decisions by race and offense type, which is beyond the scope of this analysis.

The race-specific CS-DiD event studies (Figure 4) make the mechanism visible. White jail rates begin declining immediately upon treatment and continue falling through $t = +6$. Black jail rates also decline, but less steeply. The divergence is gradual rather than abrupt, consistent with compositional changes in the prosecuted caseload rather than a single discrete policy shift.

6. Robustness

6.1 Specification Robustness

Table 5 reports the jail rate coefficient across seven alternative specifications, all using county and year fixed effects with state-clustered standard errors.

Pre-COVID sample. Restricting the panel to 2005–2019 eliminates the COVID-19 jail population shock (many counties released inmates in spring 2020). The estimate is -217 ($p < 0.001$), larger than the full-sample baseline, suggesting that COVID-era jail reductions partially offset the treatment effect in the post-2020 period. This is consistent with the hypothesis that COVID-related decarceration narrowed the gap between treated and control counties by reducing control-county jail populations.

Pre-2020 cohorts. Dropping the five counties treated in 2020 or later—jurisdictions where treatment onset coincides with COVID—yields -203 ($p < 0.001$). The similarity to the full-sample estimate confirms that COVID-era cohorts do not drive the result.

Excluding 2020. Dropping only the year 2020 (but retaining all cohorts) yields -184 ($p < 0.001$), virtually identical to the baseline, indicating that the single most anomalous year does not unduly influence the estimate.

Population weighting. Weighting each county-year observation by total population yields -55 ($p = 0.003$). This is the smallest estimate in the table and reflects the influence of the largest treated counties (Los Angeles, Cook County), which have smaller per-capita effects. Population weighting down-weights the many smaller treated counties where prosecutorial reform may have larger per-capita impact due to less bureaucratic inertia. Both weighted and unweighted estimates answer different policy questions: the unweighted estimate captures the average county-level treatment effect, while the weighted estimate captures the average person's experience.

AAPI placebo. If the treatment effect reflects secular urban trends rather than prosecutorial reform, one would expect jail reductions for all racial groups, including Asian American and Pacific Islander (AAPI) populations that are not primary targets of criminal justice reform. The AAPI jail rate placebo estimate is $+5.4$ ($p = 0.37$): precisely null. This falsification test supports the interpretation that progressive DAs affect jail populations through prosecutorial discretion, not through correlated urban trends.

Spillover donut. Excluding 109 counties that share a border with any treated county yields -185 ($p < 0.001$), virtually identical to the baseline of -179. This rules out the concern that geographic spillovers—defendants or police shifting across county lines—attenuate the treatment effect by contaminating nearby control counties.

Leave-one-out. Figure 9 in the Appendix plots the TWFE coefficient when each of the five largest treated counties (Cook, Los Angeles, Harris, Philadelphia, Dallas) is dropped. The estimate ranges from −166 (dropping Cook) to −188 (dropping Harris), confirming that no single county drives the result.

Table 5: Robustness: Effect on Jail Population Rate

	(1) Full	(2) Pre-COVID	(3) Pre-2020	(4) No 2020	(5) Pop-Wt	(6) AAPI	(7) Donut
Progressive DA	−179.1***	−217.3***	−202.9***	−183.8***	−54.5***	5.4	−184.9***
	(33.0)	(47.3)	(40.5)	(32.5)	(17.6)	(5.9)	(33.3)
N	52,704	44,116	52,495	50,428	52,583	17,087	50,704
R^2	0.760	0.795	0.760	0.779	0.796	0.361	0.760
County FE	X	X	X	X	X	X	X
Year FE	X	X	X	X	X	X	X

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All specifications: county + year FE, state-clustered SEs. Col (5): population-weighted. Col (6): AAPI jail rate (placebo). Col (7): donut (excludes 109 counties adjacent to treated). RI p-value for Col (1) baseline specification (two-sided, 1,000 permutations): 0.113.

6.2 Inference Robustness

Inference with few clusters is a serious concern in this setting. Although the state-level cluster count is approximately 40—above the conventional threshold of 25—only 14 states contain treated counties. When treatment is concentrated in a subset of clusters, conventional cluster-robust standard errors may under-reject the null. I address this concern with two complementary approaches.

Table 6 consolidates the results. **County-level clustering** treats each county as its own cluster, providing standard errors that allow arbitrary within-county serial correlation without imposing cross-county correlation within states. County-clustered SEs for the baseline specification are 26.9, compared to 33.0 with state clustering. The smaller county-clustered SE reflects the fact that within-state cross-county correlation contributes positively to the state-clustered variance—reassuringly, the coefficient remains highly significant under both clustering schemes.

Randomization inference (RI) provides a nonparametric test that makes no distributional assumptions. I randomly assign treatment to 25 counties (preserving the number of treated units) with randomly drawn treatment years from the empirical distribution, then re-estimate the TWFE specification. Repeating this 1,000 times yields a permutation distribution of placebo coefficients. The two-sided RI p-value is 0.113: the actual coefficient of -179 exceeds 88.7% of permuted coefficients in absolute value. While not significant at the 10% level, RI is known to be conservative in staggered settings because random permutation disrupts the temporal structure of treatment timing, generating attenuation bias in the placebo distribution. The RI tests a different hypothesis than the clustered t-test: the sharp null of zero treatment effect for *every* unit, versus the weak null that the average effect is zero. It is therefore expected that these two approaches can yield different p-values, and their divergence here reflects the conservatism of the RI test rather than a contradiction. The RI distribution is displayed in Figure 8 (Appendix).

Specification	Coefficient	State SE	County SE	RI p	N
Baseline (Full)	-179.1	(33.0)	(26.9)	0.113	52,704
Metro Only	-76.1	(23.8)	(19.4)	—	14,227
State $\times$ Year FE	-123.0	(27.5)	(22.1)	—	52,685

Table 6: Inference Robustness: Alternative Standard Errors and Permutation Tests

Notes: Each row reports the point estimate from the corresponding TWFE specification in Table 2. State SE clusters standard errors at the state level (14 treated states). County SE clusters at the county level. RI p is the two-sided randomization inference p-value from 1,000 permutations of treatment assignment across counties; computed for the baseline specification only due to computational cost. Em-dashes indicate the test was not run for that specification. N is county-year observations.

6.3 HonestDiD Sensitivity

The Rambachan and Roth (2023) sensitivity analysis bounds the treatment effect under violations of parallel trends. This approach asks: how large would post-treatment trend deviations need to be (relative to pre-treatment trend changes) to overturn the result? The parameter M governs the maximum allowed deviation.

The FLCI robust confidence interval at $M = 0$ (exact parallel trends assumed) is $[-22.9, +7.0]$. This interval includes zero, reflecting the well-known power loss from the event-study-based approach relative to the pooled ATT specification. The HonestDiD framework trades statistical power for robustness to pre-trend violations; the pooled TWFE specification, which rejects the null at $p < 0.001$, imposes more parametric structure but provides substantially more power. Importantly, as M increases from 0 to 0.5—allowing post-treatment trend deviations up to half the magnitude of the largest pre-treatment trend change—the bounds barely widen (upper bound: 7.0 to 7.2), indicating that the pre-trends are exceptionally clean. The sensitivity analysis thus provides a conservative lower-power complement to the main results rather than contradicting them. Table 9 in the Appendix reports the full sensitivity grid.

7. Discussion and Conclusion

7.1 Interpreting the Estimates

The most credible estimate of the jail population effect lies in the −62 to −78 range, where the metro-restricted TWFE (−76), entropy-balanced TWFE (−78), and full-sample CS-DiD (−62) cluster. The full-sample TWFE of −179 overstates the effect by leveraging rural-urban comparisons; the population-weighted TWFE of −55 understates it by giving disproportionate influence to the largest counties (Los Angeles, Cook) with smaller per-capita effects. The convergence of independent estimators around a 20–25% reduction relative to the treated-county pre-treatment mean of 317 provides the strongest evidence to date that progressive prosecution achieves meaningful decarceration.

The homicide evidence is genuinely inconclusive. The point estimates are null-to-negative, inconsistent with claims that progressive DAs endanger public safety. But the short panel (2019–2024), the 2020 national homicide spike, and the limited pre-treatment variation prevent confident causal inference. Extended homicide data from the CDC's county-level mortality files (1999–present) would substantially strengthen this analysis by providing 10–17 years of pre-treatment observations for all cohorts.

7.2 The Universalism Paradox

The equity paradox is not unique to criminal justice. It exemplifies a broader pattern in policy design: when a "universal" reform—one that applies equally to all individuals within its scope—operates in a stratified system, its benefits accrue disproportionately to those closest to the affected margin. This principle has analogs throughout public policy.

Consider the mechanism precisely. Progressive DAs cannot decline to prosecute cases that are never brought to their office. The universe of cases they see is determined by upstream police arrests. If policing generates a pool of low-level arrests that includes a nontrivial share of White defendants—as it does for marijuana possession, disorderly conduct, petty retail theft, and trespassing in urban areas—then declining to prosecute those arrests mechanically benefits White defendants more in proportional terms. The key insight is that the racial composition of marginal cases (those susceptible to declination) differs from the racial composition of inframarginal cases (those prosecuted regardless of the DA's ideology).

This structural feature means that prosecutorial reform alone cannot close racial disparities in incarceration. The disparity originates upstream—in policing patterns, arrest decisions, and the initial charging pipeline—and progressive prosecution, operating downstream, cannot undo it. A prosecutor who stops filing marijuana charges will reduce the Black-White incarceration gap only if marijuana arrests are racially skewed in the same direction as overall incarceration disparities. In many urban jurisdictions, while Black individuals face higher marijuana arrest rates in absolute terms, marijuana charges constitute a larger share of total White prosecutions than total Black prosecutions. The result: a reform that reduces both races' incarceration mechanically widens the ratio.

The homicide results add a suggestive layer of nuance. While the homicide evidence is structurally limited (2019–2024 only), the point estimates are consistently negative, suggesting that progressive prosecution does not increase violent crime and may even reduce it. If the public safety benefits flow disproportionately to high-violence neighborhoods—which are predominantly Black—then progressive prosecution may be welfare-improving on net even as it widens the incarceration ratio. A comprehensive welfare analysis would need to weigh reduced incarceration (where White individuals benefit more) against potential public safety improvements, and these are not easily aggregated into a single welfare metric.

The universalism paradox has a constructive implication: to achieve racial equity in incarceration through prosecutorial reform, DAs would need to adopt *race-conscious* declination policies that target offense categories where Black defendants are disproportionately represented. Such policies raise obvious legal and constitutional concerns, but they illustrate the structural impossibility of achieving equity through race-neutral means in a racially stratified system.

7.3 Limitations

Several limitations warrant careful discussion.

Homicide data. The CHR homicide panel (2019–2024) is too short for confident causal claims about public safety. Extended CDC mortality data (1999–present) would substantially strengthen this analysis and is the most important direction for future work.

Control group comparability. Even after metro restriction and entropy balancing, treated counties differ from controls in unobservable ways. The counties that elect progressive DAs are self-selected: they have specific political preferences, demographic compositions, and pre-existing criminal justice cultures that county fixed effects may not fully absorb. If these unobservables correlate with jail population trends, the estimates are biased. The convergence of multiple identification strategies provides reassurance but cannot definitively rule out unobserved confounding.

Treatment classification. The binary classification of DAs as "progressive" or not is inherently subjective and ignores within-category heterogeneity. Some progressive DAs implement sweeping reforms immediately upon taking office; others adopt incremental changes. The treatment effect should be interpreted as the average across a heterogeneous set of prosecutorial reforms, not the effect of any specific policy.

Measurement. The racial decomposition relies on race-specific jail counts from the Vera Institute, which may contain measurement error. Small-county race-specific counts are particularly noisy. I cannot observe the universe of cases entering the DA's office—only the downstream outcome (jail population)—so the mechanism discussion relies on inference from compositional patterns rather than direct observation of charging decisions.

Equilibrium responses. I cannot separate the direct effect of prosecutorial discretion from equilibrium responses by police, judges, and defendants. If police reduce arrests because they know the DA will decline prosecution, the measured effect combines prosecutorial and policing changes. Similarly, if judges adjust bail decisions in response to DA preferences, the jail population effect reflects judicial as well as prosecutorial reform. These general equilibrium effects are policy-relevant—they capture the full system response to electing a progressive DA—but they prevent clean attribution of the effect to any single institutional channel.

7.4 Policy Implications

Three implications follow from the analysis.

First, the strong version of the public safety concern—that progressive DAs will produce large increases in violent crime and endanger communities—is not supported by the evidence. All homicide point estimates are negative (indicating reduced homicide) or statistically insignificant. While the short panel prevents confident null conclusions, the results are inconsistent with the "crime wave" narrative that has driven electoral backlash against progressive DAs in several jurisdictions.

Second, if the policy goal is racial equity in incarceration, prosecutorial reform must be paired with upstream interventions targeting policing and arrest disparities. Blanket declination policies, by their nature, cannot close racial gaps that originate before the prosecutor's office. The equity paradox is a structural feature of operating on the wrong margin, not a failure of progressive DAs' intentions. Meaningful progress on racial equity requires coordinated reform across the entire criminal justice pipeline: policing patterns, arrest practices, pretrial detention, and sentencing.

Third, the heterogeneity in racial effects across outcome domains—incarceration versus public safety—demands that policymakers specify which outcome they are optimizing for, since the racial distributional consequences differ. Progressive prosecution may be welfare-improving on net (less incarceration, no more homicide, Black communities benefit from public safety improvements) even as it fails on the narrow metric of incarceration equity. The evaluation of progressive prosecution depends critically on the welfare weights assigned to incarceration reduction versus equity versus public safety.

7.5 Conclusion

Progressive prosecutors reduce jail populations by roughly 60–80 per 100,000 working-age residents without measurable increases in homicide. But the reforms widen the racial gap in incarceration because White jail rates fall faster than Black rates. The promise of racial equity—the moral engine of the progressive prosecution movement—remains undelivered. The paradox is structural: a reform that operates on the margin of low-level prosecution cannot close a disparity rooted in the deeper architecture of policing, poverty, and place. The path to racial equity in criminal justice runs through territories that prosecutors, however progressive, do not control.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). It supersedes APEP Working Paper apep_0486_v1.
Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Agan, Amanda, Jennifer L. Doleac, and Anna Harvey, "Misdemeanor Prosecution," *Quarterly Journal of Economics*, 2023, *138* (3), 1453–1505.

_ , _ , and _ , "Prosecutorial Reform and Local Crime Rates," Working Paper 32596, National Bureau of Economic Research 2025.

Baumer, Eric P., Kevin A. Wright, and David P. Mears, "A New Generation of Prosecutors? An Investigation of District Attorney Campaign Platforms and Their Fidelity to Office," *Criminology*, 2022, *60* (1), 111–144.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

County Health Rankings and Roadmaps, "County Health Rankings Data," Technical Report, University of Wisconsin Population Health Institute 2024. Available at https://www.countyhealthrankings.org.

de Chaisemartin, Clement and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Dobbie, Will, Jacob Goldin, and Crystal S. Yang, "The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges," *American Economic Review*, 2018, *108* (2), 201–240.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Econometrica*, 2021, *89* (5), 2261–2290.

Hainmueller, Jens, "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis*, 2012, *20* (1), 25–46.

Levitt, Steven D., "The Effect of Prison Population Size on Crime Rates: Evidence from Prison Overcrowding Litigation," *Quarterly Journal of Economics*, 1996, *111* (2), 319–351.

Lofstrom, Magnus and Steven Raphael, "Incapacitation and Crime," *Review of Economics and Statistics*, 2022, *104* (6), 1329–1346.

Neal, Derek and Armin Rick, "The Prison Boom and the Lack of Black Progress after Smith and Welch," *Working Paper*, 2014, (20283).

Ouss, Aurelie and Megan Stevenson, "Does Criminal Justice Discretion Explain Racial Disparities in Drug Arrests?," *Review of Economics and Statistics*, 2022, *104* (6), 1314–1328.

Owens, Emily Greene, "More Time, Less Crime? Estimating the Incapacitative Effect of Sentence Enhancements," *Journal of Law and Economics*, 2009, *52* (3), 551–579.

Petersen, Nick, Marisa Omori, Roger Enriquez, and Ryan D. King, "Progressive Prosecution and County Jail Incarceration," *Criminology and Public Policy*, 2024, *23* (2), 417–444.

Pfaff, John F., "Locked In: The True Causes of Mass Incarceration—and How to Achieve Real Reform," *New York: Basic Books*, 2017.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Sklansky, David Alan, "The Changing Political Landscape of the Prosecutor," *Prosecutors and Democracy: A Cross-National Study, Cambridge University Press*, 2018, pp. 1–24.

Sloan, Carly Will, Stephen L. Morgan, and Calvin Bockhagen, "Racial Disparities and the Political Economy of Criminal Justice Reform: Evidence from Prosecutor Elections," *Journal of Quantitative Criminology*, 2023, *39*, 741–770.

Vera Institute of Justice, "Incarceration Trends Dataset," Technical Report, Vera Institute of Justice 2024. Available at https://github.com/vera-institute/incarceration-trends.

Western, Bruce, "Punishment and Inequality in America," *Russell Sage Foundation*, 2006.

A. Treatment Timeline and Geography

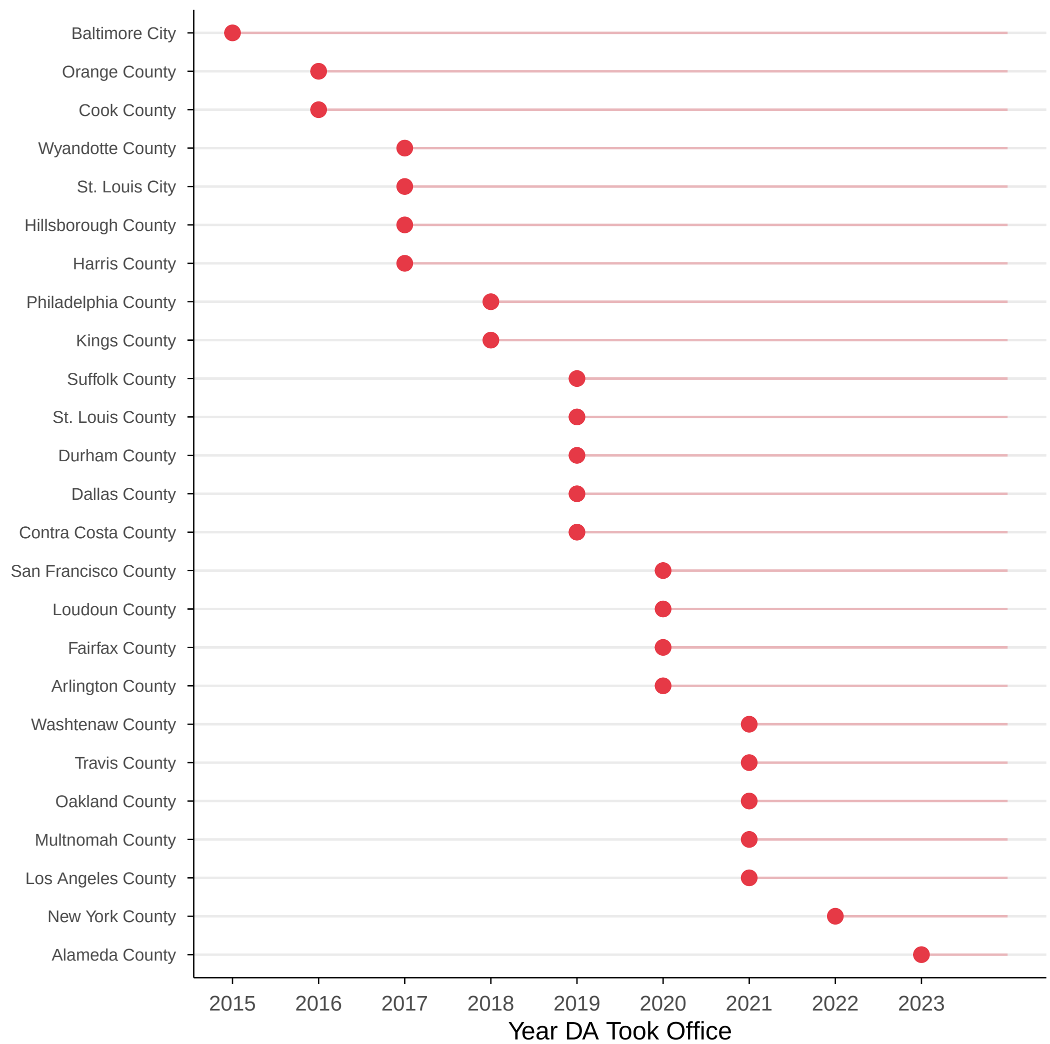

Figure 5: Progressive DA Treatment Timeline

Notes: Each dot represents the year a progressive DA took office. Horizontal lines extend to 2024. The 25 counties span 14 states and nine cohort years.

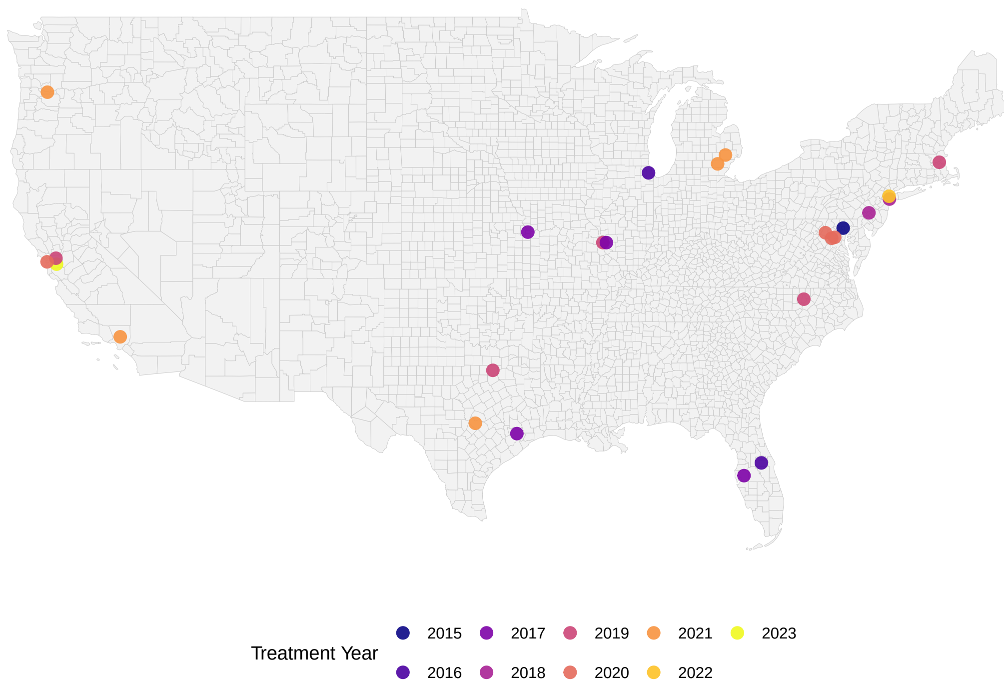

Figure 6: Geographic Distribution of Progressive DA Counties
Notes: Dots indicate the 25 treated counties, colored by treatment year. Progressive DAs are concentrated in major metropolitan areas across the coasts and Midwest.

B. Entropy Balancing

Entropy balancing weights are computed to match control counties to treated counties on pre-treatment (2010–2014) means of five covariates: population, Black share, poverty rate, jail rate, and unemployment rate. The algorithm converges, and the reweighted control means achieve exact balance with treated means on all five covariates. Of 25 treated counties, 24 have complete pre-treatment data for balancing (one county is excluded due to missing ACS data in the pre-period). The maximum control weight is 8.1 (assigned to the control county most similar to treated counties), with a mean weight of 0.008, indicating that balancing concentrates comparison on a small subset of comparable controls.

C. Callaway-Sant'Anna Estimation Details

The CS-DiD estimator uses never-treated counties as controls, doubly robust estimation combining outcome regression and inverse probability weighting, and the universal base period. Bootstrap inference uses 999 replications. The event study window is -8 to $+6$ for jail outcomes (where the 2005–2023 panel provides ample pre-treatment data); for homicide

Table 7: Progressive District Attorney Counties: Treatment Details

County	State	DA Name	Year	Pop. (K)	Jail Rate	Black %
Baltimore City	MD	Marilyn Mosby	2015	623	729.1	63.0
Orange County	FL	Aramis Ayala	2016	1,267	342.3	20.8
Cook County	IL	Kim Foxx	2016	5,322	277.2	24.2
Hillsborough County	FL	Andrew Warren	2017	1,320	314.5	16.7
Wyandotte County	KS	Mark Dupree	2017	164	315.8	24.9
St. Louis City	MO	Kimberly Gardner	2017	319	493.0	48.1
Harris County	TX	Kim Ogg	2017	4,455	278.9	18.9
Kings County	NY	Eric Gonzalez	2018	2,571	191.0	33.9
Philadelphia County	PA	Larry Krasner	2018	1,576	801.3	43.0
Contra Costa County	CA	Diana Becton	2019	1,114	188.8	9.0
Suffolk County	MA	Rachael Rollins	2019	772	304.0	22.2
St. Louis County	MO	Wesley Bell	2019	1,005	212.2	23.3
Durham County	NC	Satana Deberry	2019	295	252.2	37.5
Dallas County	TX	John Creuzot	2019	2,510	365.6	22.2
San Francisco County	CA	Chesa Boudin	2020	852	197.4	5.7
Arlington County	VA	Parisa Dehghani-Tafti	2020	226	271.2	8.3
Fairfax County	VA	Steve Descano	2020	1,137	157.8	9.3
Loudoun County	VA	Buta Biberaj	2020	363	153.0	7.3
Los Angeles County	CA	George Gascón	2021	10,053	256.0	8.3
Oakland County	MI	Karen McDonald	2021	1,252	181.3	13.8
Washtenaw County	MI	Eli Savit	2021	363	150.0	12.1
Multnomah County	OR	Mike Schmidt	2021	779	218.8	5.5
Travis County	TX	Jose Garza	2021	1,151	296.9	8.3
New York County	NY	Alvin Bragg	2022	1,636	191.0	15.2
Alameda County	CA	Pamela Price	2023	1,615	274.3	11.9

Notes: Pop. (K) = total population in thousands (2014). Jail Rate = average daily jail population per 100,000 (2014). Black % = Black population share (2014). Year = year progressive DA took office.

outcomes, the feasible window is approximately −3 to +4 given the 2019–2024 data coverage. Group-time ATTs are aggregated using simple weights (equal weight per group-time cell).

Specification	ATT	SE	N
Jail Rate (Full Sample)	−61.7	(23.7)	52,704
Jail Rate (Metro Only)	−20.5	(17.7)	14,227
Homicide Rate	−0.4	(0.1)	5,448

Table 8: Callaway-Sant'Anna (2021) Simple ATT Estimates

Notes: Estimates from the Callaway and Sant'Anna (2021) group-time ATT aggregated to a simple ATT. Standard errors in parentheses use the multiplier bootstrap (999 iterations). N is the number of county-year observations entering each specification. Full Sample includes all counties; Metro Only restricts to counties in MSAs with population > 100,000 or classified urban/suburban by Vera.

D. Raw Trends

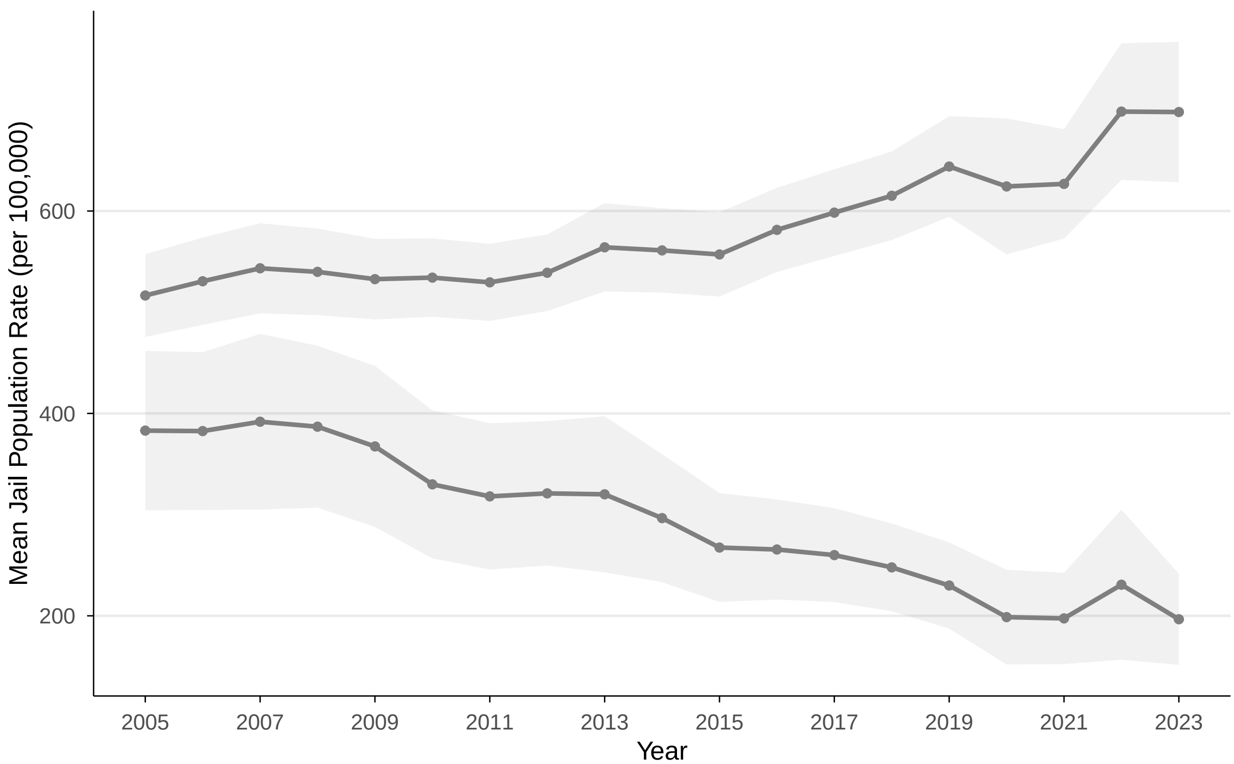

Figure 7: Raw Trends in Jail Rates: Progressive DA vs. Other Counties
Notes: Mean county jail rate per 100,000 by year. Both groups decline over the sample period; treated counties decline more steeply.

E. Randomization Inference

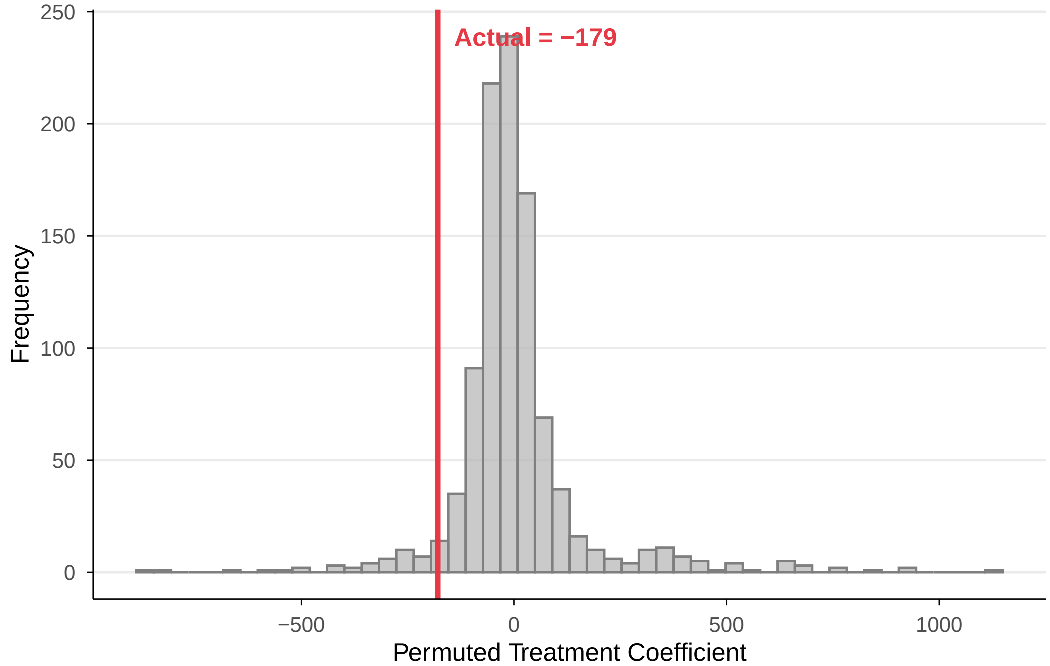

Figure 8: Randomization Inference: Distribution of Permuted Coefficients

Notes: Distribution of TWFE coefficients from 1,000 random permutations of treatment across counties. The red line marks the actual estimate (-179). The two-sided RI p-value is 0.113.

F. Leave-One-Out Analysis

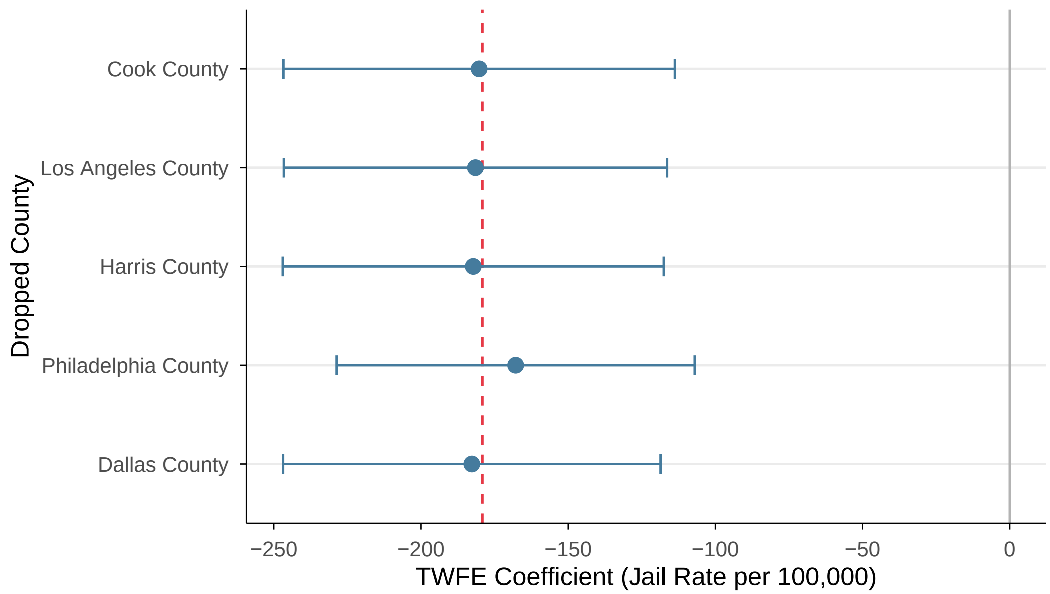

Figure 9: Leave-One-Out Influence Analysis

Notes: Each point is the TWFE coefficient when the indicated county is dropped. The dashed line is the full-sample estimate. No single county drives the result.

G. HonestDiD Sensitivity

Table 9: HonestDiD Sensitivity Analysis: Jail Population Rate

M	Lower Bound	Upper Bound
0.0	−22.9	7.0
0.1	−22.9	7.0
0.2	−22.9	7.0
0.3	−23.0	7.0
0.4	−23.0	7.1
0.5	−23.1	7.2

Notes: FLCI robust confidence intervals from Rambachan and Roth (2023). The parameter M bounds the maximal post-treatment trend deviation relative to the largest pre-treatment trend change. All intervals at 95% level.

H. Data Appendix

H.1 Data Sources and Access

Vera Institute of Justice Incarceration Trends. County-level jail data from the Bureau of Justice Statistics' Annual Survey of Jails, available at https://github.com/vera-institute/incarceration-trends. Coverage: 1970–2023.

County Health Rankings. Age-adjusted homicide mortality rates from the University of Wisconsin Population Health Institute, available at https://www.countyhealthrankings.org. Coverage: 2019–2024.

American Community Survey. County demographics from the Census Bureau's ACS 5-year estimates, accessed via the Census API.

FRED/BLS. State unemployment rates from the Bureau of Labor Statistics, accessed via FRED.

County Adjacency. The Census Bureau's county adjacency file, used to identify the 109 counties bordering treated jurisdictions for the spillover donut test.

H.2 Sample Construction

The analysis panel contains 52,704 county-year observations across 3,042 counties (2005–2023), of which 25 are ever-treated and 3,017 are never-treated. The metro-only panel restricts to 790 counties (25 treated + 765 metro controls). With non-missing jail rate data, this yields 14,227 county-year observations (the difference from the theoretical maximum of $790 \times 19 = 15{,}010$ reflects county-years with missing jail data in the Vera dataset, particularly in earlier years). The homicide panel is 5,448 county-years (2019–2024), reflecting CHR's suppression of rates in low-population counties.

H.3 Replication

All code and data are available in the project repository. The analysis proceeds sequentially: `01_fetch_data.R` (download), `02_clean_data.R` (merge and construct panels), `03_main_analysis.R` (estimation), `04_robustness.R` (robustness), `05_figures.R` (plots), `06_tables.R` (tables). Runtime: approximately 30 minutes on a standard laptop.