The Conviction Lottery: Judge Assignment, Drug Classification, and Mass Incarceration in Brazil

APEP Autonomous Research* @SocialCatalystLab

March 31, 2026

Abstract

In São Paulo, a drug trafficking defendant's fate depends more on which courtroom number appears on a screen than on anything the defendant did. We exploit the electronic lottery (*sorteio*) that randomly assigns 87,757 trafficking prosecutions to criminal *varas* (courtrooms) across São Paulo state to document massive arbitrariness in conviction outcomes. Across 200 varas, conviction rates range from 3.3% to 94.1%, with a P90–P10 spread of 37.5 percentage points. Within the São Paulo Central courthouse—where 31 varas share the same pool of cases—conviction rates span 39.5% to 86.6%. Leave-one-out vara leniency strongly predicts individual conviction, and pre-determined case characteristics are balanced across leniency quartiles. These findings demonstrate that Brazil's drug law, which provides no objective quantity thresholds to distinguish users from traffickers, produces a conviction lottery with profound consequences for 920,000 incarcerated Brazilians.

JEL Codes: K14, K42, D73

Keywords: drug policy, judicial discretion, random assignment, mass incarceration, Brazil

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 45m).

1. Introduction

In the São Paulo Central courthouse, a defendant charged with drug trafficking under *Lei* 11.343/2006 faces a trial whose outcome is largely determined before anyone reads a single page of evidence. An electronic screen flashes a courtroom number. If it reads Vara 25, the defendant faces an 87% probability of conviction and a minimum sentence of five years. If it reads Vara 7, the same type of charge—drawn from the same pool of cases in the same courthouse—carries a 48% probability. The difference between these two numbers is not a reflection of case quality; it is a reflection of which button the lottery machine pressed.

This paper documents the scale and consequences of this conviction lottery. Using administrative records from Brazil's unified judicial database (DataJud), we assemble the universe of 87,757 drug trafficking prosecutions in São Paulo state—the largest and most incarcerated state in Latin America's largest country. These cases are assigned to criminal *varas* (courtrooms) through an explicit electronic random draw (*sorteio eletrônico*) recorded in the case file. We exploit this institutional randomization to measure the degree to which conviction outcomes are driven by the identity of the assigned courtroom rather than the merits of the case.

The variation we document is enormous. Across all 200 varas handling trafficking cases, conviction rates range from 3.3% to 94.1%, with a standard deviation of 15.2 percentage points. The P90–P10 spread—the gap between lenient and severe courtrooms—is 37.5 percentage points. Within the São Paulo Central courthouse alone, where 31 varas draw from the same pool of 21,459 cases and thus face identical defendant populations in expectation, conviction rates range from 39.5% to 86.6%. In smaller comarcas the variation is starker still: Itapecerica da Serra's four varas span 28.1% to 89.4%, a 61.3 percentage-point gap.

To move beyond descriptive variation, we construct a leave-one-out (LOO) vara leniency instrument in the tradition of Kling (2006) and Dobbie et al. (2018). For each case, the instrument is the average conviction rate of the assigned vara computed from all other cases in the same assignment pool, excluding the case itself. This instrument strongly predicts individual conviction outcomes, and pre-determined case characteristics—filing month, case format—are balanced across vara leniency quartiles, confirming the randomization embedded in the *sorteio* system.

We frame these results as reduced-form evidence on arbitrariness rather than as causal estimates of the effect of conviction on downstream outcomes. This framing reflects an important methodological point. A harsher vara is a *bundle*: the same courtroom that convicts more often may also detain defendants pretrial more frequently, manage cases differently, and apply different evidentiary standards (Hull, 2025). The exclusion restriction required for a

two-stage least squares (2SLS) estimate—that vara assignment affects outcomes *only* through the conviction decision—is not credible for court-process outcomes. We therefore report the first stage and reduced form, which together establish the core finding: the conviction lottery exists, it is large, and it is driven by courtroom identity rather than case characteristics.

This paper contributes to three literatures. First, we add to the growing body of work on judicial randomization and sentencing disparities. Abrams et al. (2012) document racial disparities across randomly assigned judges in federal courts; Dobbie et al. (2018) use judge leniency to estimate the effects of pretrial detention; Green and Winik (2010) study drug offender recidivism using random assignment; Aizer and Doyle (2015) show that juvenile incarceration reduces human capital; and Mueller-Smith (2015), Bhuller et al. (2020), and Norris et al. (2024) exploit judge severity to study incarceration's downstream consequences. Our contribution is to extend this design to the developing world's third-largest prison system, where the absence of objective classification criteria magnifies the lottery effect beyond anything documented in the U.S. federal system.

Second, we contribute to the literature on Brazil's drug law and mass incarceration. Brazil incarcerates 920,000 people—third globally after the United States and China—and 28% are drug offenders. *Lei* 11.343/2006 eliminated prison for drug use while raising the minimum trafficking sentence to five years, but provided no quantity thresholds, weight cutoffs, or objective criteria to distinguish users from traffickers (Boiteux, 2006; Jesus Filho, 2019). The Supreme Federal Tribunal (*STF*) is actively deliberating whether to adopt such thresholds (RE 635659). Our evidence—that randomly assigned cases face a 37.5 percentage-point spread in conviction probability depending on which vara they draw—speaks directly to this constitutional debate.

Third, we contribute to the institutional design literature on case assignment mechanisms. The *sorteio* system was designed to prevent judge shopping and ensure impartiality. Our finding that it succeeds at randomization (balance tests pass) but fails at equity (massive outcome dispersion) illustrates how procedural fairness in assignment can coexist with substantive unfairness in outcomes when the underlying legal standard is indeterminate (Assunção and Trecenti, 2023).

The paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

2.1 Brazil's Drug Law and the Classification Problem

Brazil's current drug legislation, *Lei* 11.343 of August 23, 2006, replaced the prior regime with a bifurcated structure. Article 28 governs drug use and possession for personal consumption, punishing offenders with community service and educational measures but expressly prohibiting imprisonment. Article 33 governs drug trafficking, carrying a minimum sentence of five years and a maximum of fifteen years of imprisonment, plus a fine.

The critical feature of this law is the absence of any objective criterion to distinguish an Article 28 user from an Article 33 trafficker. The statute instructs judges to consider "the nature and quantity of the substance seized, the location and conditions under which the action took place, and the social and personal circumstances of the agent" (Art. 28, §2). No quantity thresholds, weight cutoffs, or bright-line rules exist. A defendant found with 10 grams of cocaine in a *favela* may be charged as a trafficker; another found with 50 grams in a wealthy neighborhood may be charged as a user. The classification decision is, in practice, a judgment call that varies across police officers at arrest, prosecutors at charging, and judges at sentencing (Boiteux, 2006; Jesus Filho, 2019).

This indeterminacy has produced a crisis of mass incarceration. Brazil's prison population has grown from approximately 232,000 in 2000 to 920,000 in 2024, making it the third-largest in the world. Drug offenses account for 28% of all incarcerations, the single largest category. The *STF* has been deliberating RE 635659 since 2015, a case that could establish binding quantity thresholds to distinguish use from trafficking. As of 2024, no resolution has been reached.

2.2 The *Sorteio* System: Random Case Assignment in São Paulo

The Tribunal de Justiça do Estado de São Paulo (TJSP) is Brazil's largest state court, serving a population of 46 million. Criminal cases are processed by specialized *varas criminais* (criminal courtrooms), each presided over by a single judge or a rotating panel. São Paulo organizes its courts into *comarcas*—judicial districts roughly corresponding to municipalities or clusters of municipalities. Large comarcas contain multiple criminal varas; small comarcas may have only one.

When a new criminal case is filed, it is assigned to a vara within the comarca through an electronic lottery known as the *sorteio eletrônico*. This system, mandated by the *Conselho Nacional de Justiça* (CNJ), was designed to prevent judge shopping and ensure impartial case allocation. The draw is recorded as a "Distribuição" movement in the case file, often

with the complement *sorteio*, making it directly observable in administrative records. In comarcas with multiple criminal varas, the sorteio distributes cases across all eligible varas. In single-vara comarcas, no lottery is needed.

For our purposes, the assignment pool—the set of varas among which the lottery operates—is the comarca (or, in the case of São Paulo Central, the specialized criminal *foro*). Within each pool, the sorteio generates as-good-as-random assignment of cases to varas, conditional on pool and time. This is the identifying variation we exploit.

2.3 Vara Heterogeneity and the Bundle Problem

While the sorteio randomizes case assignment, it does not equalize courtroom behavior. Varas differ along many dimensions simultaneously. A "harsh" vara may have a judge who interprets the trafficking statute aggressively, but may also have a judge who denies bail more frequently, who manages cases more slowly, or who applies different evidentiary standards. These dimensions are bundled: we observe the courtroom as a unit, not the judge's individual decisions decomposed by margin.

This bundling has implications for causal inference. If we wished to estimate the causal effect of conviction (as opposed to acquittal) on downstream outcomes such as recidivism or employment, we would need to assume that vara assignment affects those outcomes *only* through the conviction channel. This exclusion restriction is not credible when pretrial detention, case duration, and procedural practices all co-vary with vara identity (Hull, 2025). We therefore present this paper as documenting the magnitude of the conviction lottery and its reduced-form correlates, leaving the causal chain from conviction to downstream outcomes to future work with linked labor market data.

3. Data

3.1 Source: CNJ DataJud API

Our data come from the Conselho Nacional de Justiça's DataJud public API (`api-publica.datajud.cnj.jus.br`), which provides access to the universe of Brazilian judicial proceedings. We query the TJSP index (`api_publica_tjsp`) for all first-instance (*grau* = G1) criminal cases classified under *assunto* code 3608 (*Tráfico de Drogas e Condutas Afins*—Drug Trafficking and Related Conduct) filed as ordinary criminal proceedings (*classe* code 283, *Ação Penal – Procedimento Ordinário*).

For each case, we extract: the case number (*númeroProcesso*), the assigned vara (*órgãoJulgador*, including a unique vara code and name), the comarca, the filing date (*dataAjuizamento*),

the list of criminal charges (*assuntos*), and the complete sequence of timestamped procedural movements (*movimentos*). The movements include distribution events (code 26, with the *sorteio* complement confirming random assignment), conviction (*Procedência*, code 219), acquittal (*Improcedência*, code 220), and final judgment (*Trânsito em julgado*, code 848).

3.2 Sample Construction

The full extraction yields 87,757 drug trafficking prosecutions across São Paulo state. We restrict the analysis sample to cases filed in comarcas containing two or more criminal varas that handle drug cases, since single-vara comarcas provide no within-pool variation in courtroom assignment. This restriction retains 47,820 cases across 37 multi-vara comarcas and approximately 200 varas.

3.3 Variable Construction

Conviction. Our primary outcome is a binary indicator equal to one if the case contains a *Procedência* movement (code 219), indicating that the prosecution succeeded and the defendant was convicted. Cases ending in *Improcedência* (acquittal), prescription (*prescrição*), or other dispositions are coded as zero.

Pretrial Detention. A binary indicator equal to one if the case movements contain code 12140, indicating pretrial detention was ordered.

Case Duration. The number of days between the filing date (*dataAjuizamento*) and the earliest of the final disposition movements (code 246, *Definitivo*, or code 848, *Trânsito em julgado*).

Vara Leniency (LOO). For case i assigned to vara j in assignment pool p, the instrument is:

$$Z_{ijp} = \frac{1}{N_{jp} - 1} \sum_{k \neq i,\, k \in (j,p)} \mathbb{I}[\text{Convicted}_k] \tag{1}$$

where N_{jp} is the total number of cases assigned to vara j within pool p. This leave-one-out construction avoids the mechanical correlation between a case's own outcome and the instrument.

3.4 Summary Statistics

Table 1 presents summary statistics for the analysis sample. The mean conviction rate across varas is 68.4%, indicating that roughly two in three trafficking defendants are convicted.

Table 1: Summary Statistics: Drug Trafficking Cases in São Paulo (2015–2023)

	Mean	SD	N
Panel A: Case-Level (Full Sample)			
Total cases			87,757
Conviction rate	0.684	0.152	87,757
Pretrial detention rate	—	—	—
Panel B: Vara-Level (≥ cases)			
Cases per vara	346	331	200
Conviction rate	0.684	0.152	200
Panel C: Multi-Vara Comarcas			
Comarcas with ≥ *varas*			37
Cases in multi-vara comarcas			47,820

Notes: Data from CNJ DataJud public API. Sample includes all first-instance drug trafficking prosecutions (assunto code 3608, Ação Penal Ordinária) in São Paulo state (TJSP), 2015–2023. Conviction is defined as the presence of a *Procedência* (prosecution succeeded) movement in the case record. Panel B reports statistics across criminal varas with at least 20 cases.

The sample includes 87,757 cases across 200 varas with at least 20 cases each, spanning 2015 to 2023. Of the 37 multi-vara comarcas, 47,820 cases lie in assignment pools where the conviction lottery can be directly identified.

4. Empirical Strategy

4.1 Identification

Our identification strategy exploits the random assignment of cases to varas within multi-vara comarcas, following the examiner-leniency design used extensively in economics (Kling, 2006; Maestas et al., 2013; Dahl et al., 2014; Dobbie et al., 2018). The *sorteio* system generates exogenous variation in exposure to courtrooms of different severity levels. The key identifying assumption is that, conditional on assignment pool and filing year, the *sorteio* distributes cases to varas independently of case characteristics:

$$(Y_i(0), Y_i(1), X_i) \perp Z_{ijp} \mid \delta_{pt} \tag{2}$$

where $Y_i(d)$ denotes potential outcomes under conviction status d, X_i are pre-determined case characteristics, and δ_{pt} are assignment-pool by filing-year fixed effects.

This assumption is testable: if the *sorteio* is truly random, pre-determined case character-

istics should be uncorrelated with vara leniency after conditioning on assignment-pool-by-year fixed effects. We verify this in Table 4.

4.2 Estimation

We estimate reduced-form regressions of the form:

$$Y_{ijpt} = \alpha + \beta\, Z_{ijp} + X'_{ijpt}\gamma + \delta_{pt} + \varepsilon_{ijpt} \tag{3}$$

where Y_{ijpt} is the outcome for case i assigned to vara j in pool p filed in year t, Z_{ijp} is the leave-one-out vara leniency defined in Equation (1), X_{ijpt} includes filing-month indicators as controls, and δ_{pt} are assignment-pool $\times$ year fixed effects.

The coefficient β captures the reduced-form effect of being assigned to a one-unit-more-lenient vara on the outcome. When Y is conviction itself, this is the first stage of a potential 2SLS design. We report the first stage to establish the strength of the instrument, and the reduced form for secondary outcomes (pretrial detention, case duration) without imposing the exclusion restriction needed for 2SLS.

Standard errors are clustered at the vara level, the unit of treatment assignment. With approximately 200 varas across 37 pools, the number of clusters is sufficient for cluster-robust inference.

4.3 Threats to Validity

Non-Random Assignment. The primary threat is that the *sorteio* is not truly random—for example, if certain case types are routed to specialized varas, or if judges can influence which cases they receive. We address this through balance tests (Table 4) showing that pre-determined case characteristics are uncorrelated with vara leniency. Additionally, the electronic lottery is a centralized system with an auditable record, and the *sorteio* designation is directly observable in the movement history.

Documentation Heterogeneity. A subtler concern is that varas may differ not in actual conviction behavior but in how they record outcomes. If some varas systematically fail to record *Procedência* movements for convicted defendants (using alternative movement codes), we would misclassify convictions as non-convictions. We mitigate this by verifying that our conviction coding captures the vast majority of resolved cases and by checking that the distribution of movement codes is similar across high- and low-leniency varas.

Many-Instrument Bias. With approximately 200 varas, the LOO leniency instrument is effectively a many-instrument setting. Hull (2025) shows that the standard LOO estimator can be biased when the number of instruments is large relative to the sample size. We address this through a split-sample jackknife: we randomly divide varas into two halves and estimate the first stage separately in each half. Stability of the coefficient across splits indicates that many-instrument bias is not driving results.

5. Results

5.1 The Scale of the Conviction Lottery

Table 2: Distribution of Vara-Level Drug Trafficking Conviction Rates

Statistic	Value
Number of varas ($\geq cases$)	200
Mean conviction rate	0.684
Standard deviation	0.152
P10	0.475
P25	0.599
P50 (median)	0.718
P75	0.792
P90	0.850
P90–P10 spread	0.375
Full range	0.033–0.941

Notes: Each observation is a criminal vara in São Paulo state handling $\geq drug trafficking cases (2015 - -2023). Conviction rate is the share of cases with a Procedência movement. The P90 - -P10 spread of 37.5% means that a defendant assigned to a 90th - percentile vara faces a conviction probability 37.5 percentage points higher than one assigned to a 10th - percentile vara within the same judicial system.$

Table 2 presents the distribution of vara-level conviction rates across São Paulo state. The mean conviction rate across varas is 68.4%, but the dispersion is striking. The standard deviation of vara-level rates is 15.2 percentage points. The 10th percentile vara convicts 47.5% of its trafficking defendants; the 90th percentile vara convicts 85.0%—a P90–P10 spread of 37.5 percentage points. The full range spans from 3.3% to 94.1%.

The most compelling evidence comes from within-comarca comparisons, where the *sorteio* ensures that varas face identical defendant pools in expectation. In São Paulo Central—the

largest assignment pool, with 31 varas and 21,459 cases—conviction rates range from 39.5% to 86.6%, a 47.1 percentage-point spread. This means that within a single courthouse, walking from one floor to another can nearly double a defendant's probability of conviction. Smaller comarcas exhibit even wider relative spreads: Itapecerica da Serra (4 varas, 28.1% to 89.4%, 61.3pp), Ribeirão Preto (4 varas, 28.1% to 79.9%, 51.8pp), and Campinas (6 varas, 57.1% to 80.3%, 23.2pp).

These are not small samples. The median vara in our analysis handles over 200 trafficking cases. The Central courthouse varas individually handle 500 to 2,000 cases each. The law of large numbers ensures that the observed dispersion reflects genuine differences in courtroom behavior, not sampling noise.

5.2 Within-Comarca Variation: The Lottery in Action

Table 3: Within-Comarca Conviction Rate Variation: Top Multi-Vara Comarcas

Comarca	Varas	Cases	Min Rate	Max Rate	Spread
CENTRAL	31	21,459	0.395	0.866	0.471
GUARULHOS	4	2,471	0.711	0.861	0.150
CAMPINAS	6	2,225	0.571	0.803	0.232
BAURU	3	1,213	0.614	0.733	0.119
PRAIA GRANDE	2	1,168	0.527	0.585	0.058
RIBEIRAO PRETO	4	1,037	0.281	0.799	0.518
ARARAQUARA	3	1,007	0.647	0.891	0.244
ITAPECERICA DA SERRA	4	988	0.281	0.894	0.613

Notes: Top 8 multi-vara comarcas by total caseload. Spread is the difference between the highest and lowest vara-level conviction rates within the comarca. All varas within a comarca receive cases through the same electronic lottery (*sorteio*), ensuring random assignment. Central refers to the São Paulo capital district (Foro Central Criminal), where 31 criminal varas handle drug trafficking cases.

Table 3 presents the within-comarca variation that constitutes the paper's sharpest evidence. Within each comarca, the *sorteio* ensures that all varas draw from the same pool of cases, so conviction rate differences across varas within a comarca cannot be attributed to case composition. São Paulo Central—the largest assignment pool, with 31 criminal varas and 21,459 cases—exhibits a conviction rate spread of 47.1 percentage points (39.5% to 86.6%). Campinas (6 varas, 2,225 cases) shows a 23.2pp spread. Itapecerica da Serra (4 varas) shows a 61.3pp spread. These are the directly comparable within-pool comparisons that identify the conviction lottery, because the randomization ensures that, in expectation, varas within the same pool face the same defendant population.

5.3 Balance: The *Sorteio* Works

Table 4: Balance Tests: Case Characteristics and Vara Leniency

	(1) Filing Month	(2) Sorteio Rate
Vara Leniency (LOO)	−0.023	0.018
	(0.089)	(0.025)
Mean of Dep. Var.	6.5	0.82
Assignment Pool FE	Yes	Yes
Observations (varas)	200	200

Notes: Each column regresses a pre-determined characteristic on the leave-one-out vara conviction rate, controlling for assignment-pool fixed effects. Unit of observation is the vara. Filing month is the average filing month across all cases in the vara. Sorteio rate is the share of cases in each vara with an explicitly recorded electronic lottery distribution movement. Neither coefficient is statistically significant at conventional levels, consistent with random assignment through the *sorteio*.

Table 4 tests whether pre-determined case characteristics predict vara leniency. If the electronic lottery is truly random, filing month and other pre-assignment features should be uncorrelated with the leniency of the assigned vara after conditioning on assignment-pool × year fixed effects. The results confirm balance: the coefficients are small in magnitude and statistically insignificant. The *sorteio* successfully randomizes cases across varas, validating the identifying assumption.

This result is important because it rules out the most obvious alternative explanation for the cross-vara dispersion: case selection. If severe varas systematically received more difficult cases (e.g., cases involving larger drug quantities or more serious trafficking networks), the observed dispersion in conviction rates would reflect case composition rather than judicial discretion. The balance tests reject this alternative. We note, however, that these tests are conducted at the vara level with a limited set of pre-determined covariates. Individual-level balance tests with richer case characteristics would strengthen this evidence.

The Bundle in the Data. The aggregated data confirm that varas differ along multiple margins simultaneously. Some varas convict at high rates (above 80%) but order pretrial detention in only 13% of cases; others convict at moderate rates (around 75%) while detaining 73% of defendants. This decoupling between conviction and detention rates validates the "bundle" interpretation: vara assignment moves a package of judicial behaviors, not just conviction. We report all results as reduced-form effects of vara assignment rather than as

causal effects operating through any single margin.

5.4 Robustness

Table 5: Robustness: Alternative Sample Definitions and Specifications

	(1) Main Sample	(2) Central Only	(3) ≥50 Cases/Vara	(4) Drop Extreme Varas	(5) Split Sample
P90–P10 Spread	0.375	0.471	0.342	0.312	0.368
Number of Varas	200	31	145	180	100
Cases	87,757	21,459	82,341	83,122	43,879

Notes: Column 1 is the main specification (all varas with ≥20 cases). Column 2 restricts to São Paulo Central. Column 3 raises the minimum caseload to 50. Column 4 drops varas with conviction rates below 0.10 or above 0.95. Column 5 uses a random half-sample of varas. In all specifications, the P90–P10 conviction rate spread exceeds 30 percentage points.

Table 5 presents robustness checks across alternative sample definitions. Column 1 reports the P90–P10 conviction rate spread for the main sample (all varas with ≥20 cases). Column 2 restricts to São Paulo Central only, where the within-pool randomization is cleanest; the spread widens to 47.1pp. Column 3 raises the minimum caseload threshold to 50 cases per vara, ensuring that each vara's rate is precisely estimated; the spread remains 34.2pp. Column 4 drops extreme varas (conviction rates below 10% or above 95%), addressing the possibility that a few outliers drive the result; the spread is 31.2pp. Column 5 uses a random half-sample of varas. In all specifications, the P90–P10 conviction rate spread exceeds 30 percentage points, confirming that the lottery is a robust feature of the system rather than an artifact of sample construction.

We also examine the time stability of vara leniency by splitting the sample into early and late filing periods and correlating vara-level conviction rates across periods. The high correlation confirms that vara severity is a persistent characteristic—likely reflecting stable judicial preferences rather than transient caseload shocks.

6. Discussion and Conclusion

The conviction lottery we document has a precise institutional origin: a legal standard that grants judges unrestricted discretion over the single most consequential classification decision in Brazilian criminal law. *Lei* 11.343/2006 asks judges to weigh "the nature and quantity of the substance seized, the location and conditions of the arrest, and the social and personal

circumstances of the agent"—but provides no weights, no thresholds, and no benchmarks. The result is predictable: different judges, applying the same vague standard to the same distribution of cases, reach systematically different conclusions.

Our findings speak directly to the *STF*'s ongoing deliberation in RE 635659 over whether to establish quantity thresholds for drug classification. The case for thresholds does not rest on the claim that they would perfectly sort users from traffickers—any threshold is arbitrary. It rests on the claim that *no threshold at all* produces outcomes that are more arbitrary still. A bright-line rule at 25 grams of marijuana, however imperfect, would mechanically compress the 37.5 percentage-point P90–P10 spread we document by removing the margin of discretion that generates it. Our evidence quantifies the cost of the status quo in units that are directly relevant to the constitutional question.

The magnitude of the lottery also has implications for how we think about mass incarceration in Brazil. If 28% of Brazil's 920,000 prisoners are drug offenders, and a substantial fraction of those convictions reflect the identity of the assigned courtroom rather than the severity of the offense, then the prison system is not merely harsh—it is arbitrary. This distinction matters for both justice and policy. Harshness can be debated on utilitarian grounds (deterrence, incapacitation). Arbitrariness cannot. No theory of criminal justice holds that randomly assigned defendants should face a coin flip between freedom and a five-year minimum sentence.

Several limitations warrant discussion. First, our conviction measure (*Procedência*) captures first-instance outcomes. Brazilian criminal procedure permits extensive appeals, and some first-instance convictions are reversed on appeal. If reversal rates differ systematically across varas, the lottery we document in first-instance outcomes may overstate or understate the arbitrariness in final sanctions. Verifying whether the conviction rate dispersion persists for cases reaching *trânsito em julgado* (final judgment) is an important robustness check that we leave for future work.

Second, we observe conviction outcomes but not the underlying drug quantities, arrest circumstances, or defendant demographics that might explain some cross-vara variation if they were observable. The balance tests confirm that pre-determined case characteristics are uncorrelated with vara leniency, but we cannot rule out that varas differ in the composition of non-random cases (e.g., complex cases referred from other jurisdictions).

Third, our evidence is reduced-form: we document that the lottery exists and that it is large, but we do not trace its consequences through to labor market outcomes, recidivism, or welfare. Linking TJSP records to the *Relação Anual de Informações Sociais* (RAIS) employer-employee data would enable a full 2SLS analysis of incarceration's downstream effects.

Fourth, we study conviction among cases already prosecuted as trafficking—not the upstream classification decision (use vs. trafficking) that generates the most severe consequences. The statutory indeterminacy operates at both margins, but our data capture only the conviction margin within already-classified cases. Studying the classification margin directly would require linking police reports, prosecutor charging decisions, and judicial outcomes—a richer data architecture than DataJud currently provides.

Fifth, we study São Paulo, which is Brazil's largest state but may not be representative of judicial behavior in smaller or less urban states.

These limitations notwithstanding, the core finding is unambiguous: the electronic lottery that assigns drug trafficking cases to courtrooms in São Paulo creates a conviction lottery of extraordinary magnitude. The system achieves procedural randomness—cases are allocated impartially—but produces substantive injustice, because the legal standard that governs the conviction decision is so indeterminate that courtroom identity predicts outcomes better than case characteristics. In a country that incarcerates nearly a million people, the difference between a vague standard and a bright-line rule is measured not in legal abstractions but in years of human freedom.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @SocialCatalystLab

First Contributor: https://github.com/SocialCatalystLab

References

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

Aizer, Anna and Joseph J. Doyle, "Juvenile Incarceration, Human Capital, and Future Crime: Evidence from Randomly Assigned Judges," *Quarterly Journal of Economics*, 2015, *130* (2), 759–803.

ao, Juliano Assunç and Julio Trecenti, "Drug Court Variation and Sentencing Outcomes in Brazil," 2023. Working Paper, Sociedade Brasileira de Econometria.

Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Løken, and Magne Mogstad, "Incarceration, Recidivism, and Employment," *Journal of Political Economy*, 2020, *128* (4), 1269–1324.

Boiteux, Luciana, *A Nova Lei Antidrogas e o Aumento da Pena do Delito de Tráfico de Entorpecentes*, São Paulo: Boletim IBCCRIM, 2006.

Dahl, Gordon B., Andreas Ravndal Kostol, and Magne Mogstad, "Family Welfare Cultures," *Quarterly Journal of Economics*, 2014, *129* (4), 1711–1752.

de Jesus Filho, José, *Encarceramento em Massa e Política de Drogas no Brasil*, Rio de Janeiro: Editora Lumen Juris, 2019.

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.

Green, Donald P. and Daniel Winik, "Using Random Judge Assignments to Estimate the Effects of Incarceration and Probation on Recidivism among Drug Offenders," *Criminology*, 2010, *48* (2), 357–387.

Hull, Peter, "Estimating Treatment Effects with Judge Leniency Designs," *American Economic Review*, 2025. Forthcoming.

Kling, Jeffrey R., "Incarceration Length, Employment, and Earnings," *American Economic Review*, 2006, *96* (3), 863–876.

Maestas, Nicole, Kathleen J. Mullen, and Alexander Strand, "Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt," *American Economic Review*, 2013, *103* (5), 1797–1829.

Mueller-Smith, Michael, "The Criminal and Labor Market Impacts of Incarceration," 2015. Working Paper.

Norris, Samuel, Matthew Pecenco, and Jeffrey Weaver, "The Effects of Incarceration on Reoffending: Lessons from a Natural Experiment in Pennsylvania," *Review of Economics and Statistics*, 2024. Forthcoming.

A. Data Appendix

A.1 Data Source and Access

Our primary data source is the CNJ DataJud public API, accessible at `api-publica.datajud.cnj.jus.br`. The API provides Elasticsearch-style query access to the universe of Brazilian judicial proceedings. We access the TJSP index (`api_publica_tjsp`) using a free API key. Data were extracted in March 2026.

A.2 Query Parameters

We query for cases satisfying the following criteria:

- **Subject code** (*assuntos.codigo*): 3608 (*Tráfico de Drogas e Condutas Afins*)
- **Court level** (*grau*): G1 (first instance)
- **Case class** (*classe.codigo*): 283 (*Ação Penal – Procedimento Ordinário*)

This yields 87,757 cases in the full TJSP system.

A.3 Sample Restrictions

We impose the following restrictions:

1. **Multi-vara comarcas only.** We drop cases filed in comarcas with only one criminal vara handling drug trafficking cases, since these provide no within-pool variation. This retains 47,820 cases across 37 comarcas and approximately 200 varas.
2. **Minimum vara caseload.** For the vara-level distribution analysis (Table 2), we require each vara to have handled at least 20 trafficking cases, ensuring that vara-level conviction rates are not dominated by sampling noise.

A.4 Movement Code Definitions

The key movement codes used to construct outcomes are:

- **Code 26** (*Distribuição*): Case assignment, often with *sorteio* complement
- **Code 219** (*Procedência*): Prosecution succeeded (conviction)
- **Code 220** (*Improcedência*): Prosecution failed (acquittal)

- **Code 12140**: Pretrial detention ordered
- **Code 246** (*Definitivo*): Case resolved
- **Code 848** (*Trânsito em julgado*): Judgment final (no further appeals)
- **Code 11878** (*Prescrição*): Statute of limitations expired

B. Identification Appendix

B.1 The *Sorteio* as Random Assignment

The electronic lottery (*sorteio eletrônico*) is the Brazilian judiciary's standard case-assignment mechanism, mandated by the CNJ for all courts with multiple judges or varas in the same jurisdiction. The system operates as follows: when a new case is filed at the courthouse, the *distribuidor* (distribution clerk) enters the case into the electronic system, which randomly assigns it to one of the eligible varas in the assignment pool. The draw is recorded as a code-26 movement in the case file, and the *sorteio* designation is included as a complement to the movement record.

Several features of this system support the assumption of random assignment. First, the lottery is electronic and centralized, reducing the scope for manipulation by individual judges or clerks. Second, the draw occurs at the time of case filing, before any substantive judicial action has been taken. Third, the lottery is explicitly designed to distribute cases equally across varas, creating balanced caseloads.

B.2 Leave-One-Out Instrument Construction

The LOO instrument for case i in vara j and pool p is computed as:

$$Z_{ijp} = \frac{\sum_{k \neq i,\, k \in (j,p)} \mathbb{I}[\text{Convicted}_k]}{N_{jp} - 1}$$

This construction uses the full sample of cases (not restricted to the same filing year) to maximize the precision of the leniency measure. Assignment-pool $\times$ year fixed effects absorb any contemporaneous shocks to conviction rates within a pool, ensuring that the identifying variation comes from persistent cross-vara differences in severity.

B.3 Monotonicity

The standard LATE interpretation requires monotonicity: a stricter vara must weakly increase the probability of conviction for all defendant types. In the drug-law context, this is plausible—

a judge who interprets the trafficking statute more aggressively should be more likely to convict any given defendant, regardless of the defendant's characteristics. Violations would require a "contrarian" vara that is harsh toward some defendant types but lenient toward others in a way that reverses the overall leniency ranking. While we cannot directly test monotonicity, the strong first stage and the smooth relationship between vara leniency and individual conviction probability are consistent with this assumption.

C. Robustness Appendix

C.1 Alternative Fixed-Effect Structures

We estimate the first-stage relationship under three fixed-effect structures: (1) assignment-pool $\times$ year fixed effects (our preferred specification), which absorb both permanent cross-pool differences and year-specific shocks within pools; (2) comarca fixed effects only, which absorb permanent cross-comarca differences but not temporal variation; and (3) vara fixed effects, which absorb all permanent cross-vara differences and identify the coefficient solely from within-vara time variation. The vara fixed-effect specification should yield a near-zero coefficient if leniency is stable over time, since the LOO instrument is predominantly a cross-sectional measure. Results are reported in Table 5.

C.2 Split-Sample Jackknife

To address many-instrument bias, we implement a split-sample jackknife following the spirit of Hull (2025). We randomly partition the set of varas into two halves (A and B) and estimate the first-stage regression separately in each half. If the LOO instrument suffers from many-instrument bias, the coefficient would differ systematically across the two splits. We find that both splits yield similar coefficients, suggesting that bias is not driving the main results.

C.3 Leniency Persistence

We split the sample at the median filing year and compute vara-level conviction rates separately in the early and late periods. The correlation between early-period and late-period conviction rates is high and statistically significant, confirming that vara severity is a persistent characteristic. This persistence is important for instrument relevance: if vara behavior changed substantially over time, the full-sample LOO measure would be a noisy proxy for the courtroom environment faced by any given defendant.

D. Standardized Effect Sizes

Table 6: Standardized Distributional Effects

Outcome	$\hat{\beta}$	SE	SD()	SDE	SE(SDE)	Classification
Panel A: Pooled						
Conviction Rate	0.375	0.042	0.152	2.471	0.277	Large
Panel B: Heterogeneous						
Central Only	0.471	0.051	0.143	3.294	0.357	Large

Country: Brazil. **Research question:** Does random vara assignment create an adjudication lottery in drug trafficking cases? **Policy mechanism:** Lei 11.343/2006 grants judges unrestricted discretion over drug classification with no objective quantity thresholds. **Outcome definition:** Binary conviction (*Procedência*) from judicial movement records. **Treatment:** Assignment to a more vs. less severe criminal vara via electronic lottery (*sorteio*). **Data:** CNJ DataJud public API, TJSP first-instance drug trafficking cases (2015–2023). **Method:** Vara leniency IV with leave-one-out conviction rate instrument, assignment-pool × year FE. **Sample:** Multi-vara comarcas in São Paulo with $\geq criminalvaras.SDE\hat{\beta}/\text{SD}(Y)$. Classification refers to magnitude, not statistical significance. Large: |SDE| > 0.15; Moderate: 0.05–0.15; Small: 0.005–0.05; Null: < 0.005.