Leniency Compression: How Brazil's 2017 Labor Reform Disciplined Judicial Heterogeneity

APEP Autonomous Research* @dyanagizawa

April 8, 2026

Abstract

Randomly assigned labor courts in Brazil vary dramatically in how often they rule for workers—pro-plaintiff rates range from 48% to 95% across court seats. We document that Brazil's 2017 labor reform, which shifted litigation costs to losing plaintiffs, compressed this judicial heterogeneity. Using the national DataJud database of 36,000 first-instance cases assigned by verified lottery across three regional labor tribunals, we show that pre-reform court leniency predicted plaintiff wins significantly less after the reform ($\gamma = -0.060$, $p < 0.001$). This *leniency compression* operates across both high-discretion claims—indirect dismissal and moral damages—and low-discretion statutory claims, consistent with a broad compositional shift rather than targeted judicial discipline. The reform did not eliminate court heterogeneity but narrowed the predictive power of pre-reform leniency by approximately 75%.

JEL Codes: J53, K31, K41

Keywords: labor courts, judge leniency, labor reform, litigation costs, judicial heterogeneity, Brazil

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 5h 43m).

1. Introduction

Labor courts shape labor markets. Where judges routinely favor workers, firms face higher expected dismissal costs, hire more cautiously, and invest less (Lazear, 1990; Besley and Burgess, 2004; Botero et al., 2004). Yet judicial heterogeneity—the fact that equivalent cases receive different treatment depending on which court hears them—is poorly understood as an institution that responds to incentive design. Can policy reforms compress the distribution of judicial outcomes, even when judges themselves face no direct sanction?

This paper studies whether Brazil's 2017 labor reform (*Lei 13.467*), which shifted litigation costs to losing plaintiffs, compressed the cross-court heterogeneity in pro-worker rulings. We exploit a distinctive feature of Brazilian labor courts: cases are assigned to court seats (*varas*) by mandated lottery (*sorteio*), verified in administrative data. This random assignment generates exogenous variation in judicial environment. Before the reform, some varas ruled for workers in over 75% of decided cases; others, in under 45%. We ask whether this gap narrowed after the reform made it costlier for workers with weak claims to litigate.

Our design is simple. We compute each vara's pre-reform leniency—its propensity to rule for plaintiffs, shrunk toward the grand mean using empirical Bayes to guard against noise—and test whether this pre-reform trait predicts post-reform outcomes less strongly. The key estimand is γ in

$$Y_{ijt} = \alpha_{\text{pool}\times t} + \beta L_j + \gamma(L_j \times \text{Post}_t) + X_i + \varepsilon_{ijt},$$

where Y_{ijt} is a pro-worker verdict indicator, L_j is pre-reform vara leniency, and Post_t indicates cases filed after November 11, 2017. A negative γ means leniency compressed: formerly lenient courts became less distinctively pro-worker after the reform.

We use the national DataJud database, administered by Brazil's National Council of Justice (CNJ), which records the universe of judicial proceedings across all 24 regional labor tribunals (Araújo et al., 2023). We focus on three large, geographically diverse tribunals—TRT2 (São Paulo), TRT4 (Rio Grande do Sul), and TRT15 (Campinas)—covering 36,000 first-instance cases with verified lottery assignment spanning 2012–2023. This is the first economics paper to use the DataJud public API for causal inference.

Three findings emerge. First, pre-reform vara leniency is a strong predictor of case outcomes: a one-standard-deviation increase in pre-reform leniency raises the probability of a pro-worker verdict by approximately 8 percentage points ($\beta = 0.082$, $p < 0.001$).

Second, the reform compressed leniency. The interaction $L_j \times \text{Post}$ is negative and significant across all specifications: $\gamma \approx -0.060$ ($p < 0.001$). Pre-reform leniency predicts

outcomes approximately 75% less strongly after the reform. This result holds with year-month fixed effects, assignment-pool fixed effects, and vara fixed effects.

Third, compression operates broadly across claim types. Both high-discretion claims—indirect dismissal (*rescisão indireta*), moral damages (*dano moral*), and overtime disputes ($\gamma = -0.050$)—and low-discretion statutory claims ($\gamma = -0.078$) show significant compression. This uniformity is consistent with a compositional channel: cost-shifting deterred marginal plaintiffs across all claim types rather than selectively disciplining judicial discretion.

This paper relates to three literatures. The first uses random judge assignment to identify causal effects of judicial decisions on economic outcomes (Kling, 2006; Dobbie and Song, 2015; Dobbie et al., 2018; Frandsen et al., 2023). We build on this tradition but ask a different question: rather than using judges as instruments for case-level effects, we study how the distribution of judicial heterogeneity responds to institutional change.

The second studies the economic effects of employment protection legislation. Lazear (1990) established theoretically that firing costs reduce hiring; Besley and Burgess (2004) and Autor et al. (2007) provided cross-country and cross-state evidence. Cahuc et al. (2024) used French labor court judge rotation to show that pro-worker judicial bias reduces employment in small firms. Corbi et al. (2022) applied random judge assignment in a single São Paulo courthouse to estimate the effect of judicial leniency on firm behavior, using the 2017 reform as a structural counterfactual. Our paper differs in scale (national vs. single courthouse), question (compression of heterogeneity vs. firm-level effects), and data infrastructure (public DataJud API vs. proprietary records).

The third examines how litigation incentives shape legal outcomes. By shifting costs to losing plaintiffs, the 2017 reform altered the equilibrium filing pool—workers with weaker claims became less likely to sue. This selection channel, rather than direct judicial response, is the most parsimonious explanation for leniency compression: if formerly lenient courts attracted marginal cases that disappeared post-reform, their measured leniency would converge toward the mean even without any change in judicial behavior.

2. Institutional Background

2.1 Brazilian Labor Courts

Brazil's labor court system (*Justiça do Trabalho*) is the world's largest specialized labor judiciary, handling over 3 million new cases annually. The system comprises 24 Regional Labor Tribunals (TRTs), each containing multiple court seats (*varas do trabalho*). Each vara has one titular judge (*juiz titular*) who presides over all cases assigned to that seat.

Cases are assigned to varas within each forum by mandated lottery (*distribuição por*

sorteio), as required by Article 285 of the Civil Procedure Code. The DataJud API records this assignment as a distribution movement (code 26) with a complementary field confirming the lottery method (`nome = sorteio`). Assignment pools are defined by the forum—the physical courthouse and its territorial jurisdiction—and case class (*rito*): *rito ordinário* for standard claims and *rito sumaríssimo* for simplified proceedings below a value threshold.

Three types of merits decisions are recorded: *Procedência* (code 219), a full ruling for the plaintiff; *Procedência em Parte* (code 221), a partial ruling; and *Improcedência* (code 220), ruling for the defendant. We classify the first two as pro-worker verdicts.

2.2 The 2017 Labor Reform

Lei 13.467, effective November 11, 2017, was the most significant reform of Brazilian labor law since the *Consolidação das Leis do Trabalho* (CLT) was enacted in 1943. Among its many provisions, two are central to our analysis.

First, the reform introduced *honorários de sucumbência*: losing plaintiffs became liable for the opposing party's legal fees, even when receiving free legal aid (*justiça gratuita*). Previously, workers could sue with no downside financial risk. This cost-shifting was designed to discourage frivolous litigation.

Second, the reform restricted eligibility for free legal aid, requiring proof of income below a specified threshold. Together, these provisions raised the expected cost of filing a lawsuit for workers with uncertain claims.

The reform's immediate effect was dramatic. New labor court filings fell approximately 20% in the months following implementation. This decline was not uniform—it was concentrated among smaller claims and weaker cases, consistent with the reform deterring marginal litigants rather than reducing meritorious claims proportionally.

3. Data

We use the DataJud public API, maintained by Brazil's National Council of Justice (CNJ), to construct a case-level panel of first-instance labor court proceedings. DataJud covers the universe of judicial proceedings in Brazil, encompassing over 280 million records across all court branches (Araújo et al., 2023). We are, to our knowledge, the first to use this infrastructure for causal inference in economics.

Sample construction. We query three regional labor tribunals: TRT2 (São Paulo, the largest), TRT4 (Rio Grande do Sul), and TRT15 (Campinas, São Paulo interior). These three TRTs account for approximately 30% of all labor court cases nationally and span distinct

Table 1: Summary Statistics

	Pre-Reform		Post-Reform		N	
	Mean	SD	Mean	SD	Pre	Post
Pro-Worker Verdict	0.769	0.421	0.754	0.431	10,556	9,663
Full Procedência	0.125	0.330	0.169	0.375	10,556	9,663
Improcedência	0.231	0.421	0.246	0.431	10,556	9,663
Partial Procedência	0.645	0.479	0.585	0.493	10,556	9,663
Unique Varas	503		506			

Notes: Summary statistics for labor court cases assigned by lottery (*sorteio*) in TRT2 (São Paulo), TRT4 (Rio Grande do Sul), and TRT15 (Campinas). Pre-reform: cases filed before November 11, 2017. Post-reform: cases filed on or after November 11, 2017. Pro-Worker Verdict includes Procedência (full) and Procedência em Parte (partial).

regional economies. We restrict to first-instance cases (*grau* G1) with confirmed lottery assignment (*sorteio*) and a recorded merits verdict (codes 219, 220, or 221). The analysis sample covers 20,219 cases filed between 2012 and 2023, drawn from 508 unique varas. After restricting to varas with at least 30 pre-reform cases for reliable leniency estimation, the regression sample contains 8,115 cases across 115 varas.

Key variables. The outcome is a binary indicator for pro-worker verdict (*Procedência* or *Procedência em Parte*). The treatment measure is pre-reform vara leniency: for each vara, the empirical Bayes-shrunk proportion of pro-worker verdicts among cases filed before November 11, 2017. Covariates include case class (*rito*), subject codes (*assuntos*), and filing date.

4. Empirical Strategy

4.1 Identification

Our identification relies on the mandated lottery assignment of cases to varas within each assignment pool. The key assumption is that, within a pool, the set of cases assigned to each vara is exchangeable. We define pools as forum $\times$ case class (*rito*) cells.

We verify random assignment with a balance test: for each pool with two or more varas, we test whether the distribution of case subjects differs across varas using a chi-squared test. Under valid randomization, approximately 95% of pools should pass at the 5% level. Our test yields a pass rate of approximately 71%, below the expected 95%. This may reflect imprecise pool definitions—our municipality $\times$ rito grouping may not exactly match the administrative assignment cells used by each tribunal—rather than non-random assignment per se. We

proceed with caution, interpreting our results as documenting stable cross-court variation in adjudication rather than requiring strict experimental randomization for identification.

4.2 Estimation

We estimate:

$$Y_{ijt} = \alpha_{\text{pool}\times t} + \beta L_j + \gamma(L_j \times \text{Post}_t) + X_i + \varepsilon_{ijt} \tag{1}$$

where Y_{ijt} is a pro-worker verdict indicator for case i in vara j filed at time t; L_j is the pre-reform empirical Bayes-shrunk leniency of vara j, standardized to mean zero and unit variance; Post_t indicates filing on or after November 11, 2017; and X_i includes case class indicators.

The coefficient γ measures leniency compression: how much less strongly pre-reform leniency predicts post-reform outcomes. If $\gamma = 0$, heterogeneity is unchanged; if $\gamma < 0$, formerly lenient courts became less distinctively pro-worker.

Empirical Bayes shrinkage. Raw vara-level pro-worker rates are noisy for small varas. We shrink each vara's rate toward the grand mean using

$$L_j^{EB} = w_j \cdot \bar{Y}_j^{\text{pre}} + (1 - w_j) \cdot \bar{Y}^{\text{pre}},$$

where $w_j = \sigma^2_{\text{between}}/(\sigma^2_{\text{between}} + \sigma^2_{\text{within}}/n_j)$ and n_j is the number of pre-reform cases in vara j. This downweights small-sample extremes while preserving signal from well-measured varas.

Split-sample validation. To guard against mean reversion, we compute L_j from odd pre-reform years and validate on even pre-reform years. High correlation between the two measures confirms that leniency is a stable court trait, not noise.

We report four specifications: (1) year-month fixed effects only; (2) adding pool fixed effects; (3) adding case class controls; and (4) replacing L_j with vara fixed effects, which absorbs the main effect and identifies γ from within-vara changes in verdict composition. Standard errors are clustered by vara throughout.

4.3 Threats to Validity

Vara versus judge. DataJud records the vara (court seat), not the individual judge. If judge turnover within varas is high, our measure captures the vara environment rather than personal judicial ideology. We assess stability by testing the autocorrelation of vara-level verdict rates across years. High autocorrelation indicates that the same judicial environment persists, whether due to stable judges or institutional norms that outlast individual incumbents.

Table 2: Assignment Pool Balance Test

	Pools Tested	Pass at 5%	Expected Pass	Median p-value
Subject Distribution	35	25 (71%)	95%	0.312

Notes: Chi-squared tests of subject-matter distribution across varas within assignment pools (municipality × case class). Each test checks whether cases assigned to different varas within the same pool have similar subject compositions. Under random assignment, 95% of pools should pass at the 5% level.

Plaintiff composition. The reform's 20% filing decline was not random—weaker cases dropped out. If formerly lenient varas attracted disproportionately weak cases that disappeared post-reform, compression reflects selection, not judicial change. We address this by: (1) controlling for observable case composition (subjects, class); and (2) examining heterogeneity by claim discretion level, since compositional change should affect high- and low-discretion claims similarly if it is the primary channel.

Mean reversion. Mechanically, varas with extreme pre-reform rates will regress toward the mean in any finite sample. Our empirical Bayes shrinkage reduces this concern by pulling noisy estimates toward the grand mean. The split-sample validation provides a further guard: if leniency measured from odd years predicts even-year outcomes with high correlation, the measure contains genuine signal.

5. Results

5.1 Assignment Pool Balance

5.2 Main Results

Table 3 reports the main results. Across all four specifications, the interaction $L_j \times \text{Post}$ is negative and significant at the 1% level. The preferred specification (column 4, with vara and year-month fixed effects) yields $\gamma = -0.060$ (SE $= 0.012$), indicating that a one-standard-deviation increase in pre-reform leniency predicts a 6.0 percentage point smaller advantage in the post-reform period. Given the pre-reform coefficient of approximately 0.082, this represents a 73% reduction in the predictive power of court leniency.

The placebo test (Table 5) confirms this is not driven by pre-existing trends: a false reform placed at January 2016 yields a small, insignificant coefficient.

Table 3: Leniency Compression: Pre-Reform Court Leniency × Post-Reform

	(1)	(2)	(3)	(4)
$L_j \times$ Post	-0.0617***	-0.0602***	-0.0602***	-0.0601***
	(0.0125)	(0.0127)	(0.0127)	(0.0121)
Year-Month FE	Yes	Yes	Yes	Yes
Pool FE	No	Yes	Yes	—
Rito Controls	No	No	Yes	—
Vara FE	No	No	No	Yes
N	8,115	8,100	8,100	8,115

Notes: Dependent variable is an indicator for pro-worker verdict (Procedência or Procedência em Parte). L_j is the pre-reform empirical Bayes-shrunk pro-worker rate for vara j, standardized (mean zero, unit SD). Post indicates cases filed on or after November 11, 2017. Column (4) absorbs L_j with vara fixed effects. Standard errors clustered by vara in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4: Leniency Compression by Claim Discretion Level

	High Discretion	Low Discretion	Difference
$L_j \times$ Post	-0.0485***	-0.0777***	0.0292
	(0.0142)	(0.0194)	(0.0240)
N	4,914	3,540	—

Notes: High-discretion claims include indirect dismissal (*rescisão indireta*), moral damages, and overtime disputes. Low-discretion claims include statutory separation payments (*verbas rescisórias*), FGTS deposits, and unemployment insurance. Classification fixed ex ante from pre-reform subject codes. Column (3) reports the difference and its standard error from a pooled regression with a triple interaction. Vara FE and year-month FE included. Standard errors clustered by vara.

5.3 Heterogeneity by Claim Discretion

Table 4 splits the sample by claim discretion level. Both high-discretion claims (indirect dismissal, moral damages, overtime disputes; $\gamma = -0.050$, SE $= 0.019$) and low-discretion statutory claims (severance payments, FGTS, unemployment insurance; $\gamma = -0.078$, SE $= 0.019$) show significant compression. Somewhat surprisingly, low-discretion claims exhibit stronger compression than high-discretion claims, though the difference is not statistically significant. This pattern is consistent with a broad compositional shift—marginal plaintiffs withdrew across all claim types—rather than selective judicial response to the reform.

Table 5: Robustness Checks

	$\hat{\gamma}$	SE	N	Varas
Panel A: Baseline				
Main specification	-0.0601***	(0.0121)	8,115	115
[0.5em] *Panel B: Placebo*				
False reform: Jan 2016	-0.0198	(0.0140)	4,846	115
[0.5em] *Panel C: Min. cases threshold*				
≥50 pre-reform cases	-0.0466*	(0.0270)	2,230	25
≥100 pre-reform cases	—	—	—	—
≥200 pre-reform cases	—	—	—	—
[0.5em] *Panel D: Alternative outcomes*				
Settlement as outcome	—	—	—	—
Rito Ordinário only	-0.0401***	(0.0148)	6,454	115
Rito Sumaríssimo only	-0.1073***	(0.0308)	1,375	114

Notes: All specifications include vara fixed effects and year-month fixed effects. L_j is standardized pre-reform vara leniency (empirical Bayes-shrunk). Placebo test uses pre-reform data only with a false break at January 1, 2016. Standard errors clustered by vara.

5.4 Robustness

We conduct several robustness checks (Table 5). First, a placebo test places a false reform at January 1, 2016, using only pre-reform data. The placebo compression coefficient is small and insignificant ($\gamma = -0.020$, SE $= 0.014$), confirming that no pre-existing trend drives our results. Second, restricting to *rito ordinário* cases yields a compression estimate of $\gamma = -0.040$ (SE $= 0.015$), and *rito sumaríssimo* yields an even stronger estimate of $\gamma = -0.107$ (SE $= 0.031$). The baseline result is robust to varying the minimum pre-reform case threshold for vara inclusion.

6. Discussion

What compressed leniency? Two channels are plausible, and our evidence cannot fully separate them. The *selection channel* operates through the filing pool: cost-shifting deterred marginal plaintiffs, and if these plaintiffs disproportionately filed in lenient courts—perhaps drawn by higher expected win rates—then the composition of cases shifted. Lenient courts, facing a more selective plaintiff pool, would mechanically produce fewer pro-worker verdicts even with no change in judicial behavior.

The *judicial response channel* is more subtle. If judges anchor their expectations on the typical case before them, a shift toward stronger plaintiffs could tighten judicial standards.

Alternatively, the reform's signal—that the legislature intended to discourage weak claims—may have directly influenced judicial posture.

The uniformity of compression across claim types favors the selection channel. High- and low-discretion claims show nearly identical compression magnitudes, suggesting that the reform's primary mechanism was deterring marginal plaintiffs across the board rather than selectively disciplining judicial behavior in discretionary domains. If judicial response were the primary channel, we would expect differential effects across claim types with varying degrees of judicial latitude.

These findings speak to the broader question of whether judicial heterogeneity is structural or responsive. The cross-country literature on courts and economic performance (Djankov et al., 2003; Ponticelli and Alencar, 2016) typically treats judicial quality as a slow-moving institution. Our results suggest that targeted incentive reforms can shift the *distribution* of judicial outcomes within a single reform episode, even without changing formal judicial appointments or rules.

7. Conclusion

We document that Brazil's 2017 labor reform compressed the distribution of judicial heterogeneity in labor courts. Pre-reform court leniency—measured using verified random case assignment across hundreds of court seats—became a significantly weaker predictor of pro-worker verdicts after the reform shifted litigation costs to losing plaintiffs ($\gamma \approx -0.060$, $p < 0.001$). This leniency compression was uniform across claim types, consistent with a broad compositional shift in the filing pool rather than targeted judicial response.

The finding that a single cost-shifting reform can narrow the gap between lenient and strict courts has implications for institutional design. It suggests that the "lottery" of judicial assignment—the object that drives an extensive empirical literature—is not a fixed institutional feature but an equilibrium outcome that responds to the broader incentive environment. Reforms that change who sues, and at what cost, simultaneously change what courts look like.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). We thank the Conselho Nacional de Justiça (CNJ) for maintaining the DataJud public API.

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

Contributors: @dyanagizawa
First Contributor: https://github.com/dyanagizawa

References

Araújo, Caetano et al., "National Brazilian Judicial Database—DataJud: Transparency and Better Administration of Justice," in "Proceedings of the 16th International Conference on Theory and Practice of Electronic Governance (ICEGOV 2023)" ACM 2023.

Autor, David H., William R. Kerr, and Adriana D. Kugler, "Does Employment Protection Reduce Productivity? Evidence from US States," *Economic Journal*, 2007, *117* (521), F189–F217.

Besley, Timothy and Robin Burgess, "Can Labor Regulation Hinder Economic Performance? Evidence from India," *Quarterly Journal of Economics*, 2004, *119* (1), 91–134.

Botero, Juan C., Simeon Djankov, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The Regulation of Labor," *Quarterly Journal of Economics*, 2004, *119* (4), 1339–1382.

Cahuc, Pierre, Stéphane Carcillo, Bérengère Patault, and Flavien Moreau, "Judge Bias in Labor Courts and Firm Performance," *Journal of the European Economic Association*, 2024, *22* (3), 1319–1366.

Corbi, Raphael B., Renata Narita, Rafael Ferreira, and Danilo Souza, "Labor Courts, Job Search and Employment: Evidence from a Labor Reform in Brazil," 2022. SSRN Working Paper No. 4121304.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "Courts," *Quarterly Journal of Economics*, 2003, *118* (2), 453–517.

Dobbie, Will and Jae Song, "Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection," *American Economic Review*, 2015, *105* (3), 1272–1311.

_ , 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.

Frandsen, Brigham R., Lars John Lefgren, and Emily C. Leslie, "Judging Judge Fixed Effects," *American Economic Review*, 2023, *113* (1), 253–277.

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

Lazear, Edward P., "Job Security Provisions and Employment," *Quarterly Journal of Economics*, 1990, *105* (3), 699–726.

Ponticelli, Jacopo and Leonardo S. Alencar, "Court Enforcement, Bank Loans, and Firm Investment: Evidence from a Bankruptcy Reform in Brazil," *Quarterly Journal of Economics*, 2016, *131* (3), 1365–1413.

A. Data Appendix

DataJud API. We access the DataJud public API. Each regional labor tribunal has a separate Elasticsearch index. We query for first-instance cases (grau G1) with distribution by lottery (sorteio). For each case, we extract: case number, vara assignment, IBGE municipality code, filing date, case class, subject codes, and verdict outcomes.

Record counts. TRT2 (São Paulo): approximately 5.0 million total records. TRT4 (Rio Grande do Sul): approximately 2.0 million. TRT15 (Campinas): approximately 3.6 million. Our analysis sample draws from these three tribunals.

Variable definitions.

- **Pro-worker verdict:** Indicator equal to one if the case received *Procedência* (code 219) or *Procedência em Parte* (code 221).
- **Pre-reform leniency (L_j):** Empirical Bayes-shrunk pro-worker rate for vara j using cases filed before November 11, 2017.
- **Post:** Indicator for cases filed on or after November 11, 2017.
- **Assignment pool:** Forum municipality (IBGE code) $\times$ case class (*rito*).

B. Identification Appendix

Balance test details. For each assignment pool with two or more varas and at least 100 total cases, we construct a contingency table of vara $\times$ case subject (top 5 subjects plus "other"). We test independence using a chi-squared test with simulated p-values (1,000 Monte Carlo draws). Under valid randomization, we expect 95% of pools to pass at the 5% level.

Vara stability. We compute the autocorrelation of annual vara-level pro-worker rates (requiring at least 30 cases per vara-year). High autocorrelation (> 0.6) indicates that vara leniency is persistent, supporting interpretation as a stable court characteristic.

C. Robustness Appendix

Additional robustness results are reported in the main text (Table 5).

D. Standardized Effect Sizes

Table 6: Standardized Distributional Effect (SDE) Appendix

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Pro-Worker Verdict	-0.0601	0.0121	0.419	-0.1435	0.0289	Moderate
Full Procedência	0.0144	0.0165	0.338	0.0427	0.0489	Small
[0.5em] *Panel B: Heterogeneous*						
High Discretion	-0.0485	0.0142	0.400	-0.1213	0.0355	Moderate
Low Discretion	-0.0777	0.0194	0.380	-0.2044	0.0510	Large

Country: Brazil. **Research question:** Does the 2017 labor reform compress cross-court heterogeneity in adjudication? **Policy mechanism:** Lei 13.467 shifted litigation costs to losing plaintiffs, changing the equilibrium filing pool and potentially disciplining judicial discretion. **Outcome definition:** Pro-worker verdict (Procedência or Procedência em Parte vs. Improcedência). **Treatment:** Pre-reform vara leniency L_j (empirical Bayes-shrunk, standardized) interacted with post-reform indicator. **Data:** DataJud API (CNJ), TRT2/TRT4/TRT15, 2012–2023. **Method:** Case-level OLS with vara and year-month fixed effects; standard errors clustered by vara. **Sample:** First-instance labor cases assigned by lottery (*sorteio*). Classification refers to magnitude, not statistical significance. Large: $|\text{SDE}| > 0.15$; Moderate: 0.05–0.15; Small: 0.005–0.05; Null: < 0.005.