

The Asylum Lottery and Local Crime: Evidence from Immigration Judge Leniency

APEP Autonomous Research* @olafdrw

March 13, 2026

Abstract

Whether an asylum seeker stays in America depends partly on which judge's name is drawn from the hat: grant rates range from 1.5% to 49.4% among judges in the same New York courthouse. We exploit this enormous within-court variation to construct the first judge-leniency instrumental variable estimate of how asylum decisions affect local crime. Merging career records of 1,269 immigration judges across 88 courts with CDC county-level homicide data (2019–2024) for 29 states, we find no evidence that higher asylum grant rates increase homicide. The unconditional negative correlation between judge leniency and crime vanishes after controlling for poverty and regional composition, yielding a precisely estimated null. Balance tests confirm that cross-sectional judge composition is correlated with state demographics, motivating future panel designs exploiting within-court judge turnover.

JEL Codes: K37, K42, J15, J61

Keywords: immigration, crime, judge leniency, instrumental variables, asylum

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 30m).

1. Introduction

In the New York Immigration Court, Judge Kandah grants asylum in 1.5% of cases. Down the hall, Judge Brennan grants 49.4%. Both work in the same courthouse, adjudicating cases drawn from the same pool. Whether a refugee builds a life in America or faces deportation depends, in part, on which judge’s name appears on the docket—a feature [Ramji-Nogales et al. \(2007\)](#) famously termed “refugee roulette.”

This paper asks whether these lottery-like differences in asylum decisions affect public safety. The question sits at the intersection of two literatures that rarely speak to each other. Immigration economists have documented that immigrants commit crimes at lower rates than natives ([Butcher and Piehl, 1998](#); [Light et al., 2020](#); [Sampson, 2008](#)), but most evidence is descriptive or correlational. A separate literature in criminal justice has used random judge assignment as a powerful instrument for studying incarceration effects ([Kling, 2006](#); [Dobbie et al., 2018](#); [Maestas et al., 2013](#)). We bridge these literatures by applying the judge-leniency IV framework to the immigration-crime question for the first time.

Using career records of 1,269 immigration judges across 88 U.S. courts from the OpenImmigration.us database, we construct a state-level index of judge leniency—the caseload-weighted average career grant rate of judges at each state’s immigration courts. Within-court variation in judge grant rates is striking: the interquartile range typically spans 10–20 percentage points, and the standard deviation within a courthouse averages 7 percentage points. We merge this immigration court data with county-level homicide counts from the CDC’s Mapping Injury surveillance system (2019–2024) and demographic controls from the American Community Survey.

Our main finding is a precise null. In the unconditional cross-section, states with more lenient judges—and thus higher asylum grant rates—have significantly *lower* homicide rates ($\beta = -0.30$, $p < 0.05$). This raw correlation is striking but misleading: balance tests reveal that judge leniency is correlated with state demographics. Courts in wealthier states with larger foreign-born populations tend to have more lenient judges, and these states independently have lower homicide rates. After controlling for poverty, foreign-born population share, and Census region, the relationship between asylum grants and homicide shrinks to zero ($\beta = 0.02$, $p = 0.86$). The 2SLS point estimate, which instruments the asylum grant rate with the judge leniency index, confirms this null: a one-percentage-point increase in the instrumented grant rate changes the homicide rate by 0.02 per 100,000, with a 95% confidence interval of $[-0.25, 0.29]$.

Three pieces of evidence support the credibility of this null. First, the first-stage F-statistic exceeds 375 in cross-sectional specifications and 2,889 in the state-year panel, ruling out

weak-instrument concerns. Second, the suicide rate—a placebo outcome with no plausible link to immigration court decisions—is similarly unresponsive to judge leniency ($\beta_{IV} = -0.13$, $p = 0.15$ in the panel). Third, the null result is robust across five control specifications, two crime definitions (all homicide, firearm homicide), subsamples split by immigration intensity, and a state-year panel specification with year fixed effects and state-clustered standard errors (174 state-year observations).

We contribute to three literatures. First, we provide the first quasi-experimental evidence specifically on the *marginal asylum decision* and crime. Prior work by Miles and Cox (2014) studied the bundled effects of the Secure Communities enforcement program, and East et al. (2023) examined broader immigration enforcement on labor markets. Our design isolates a different margin: what happens when one more person receives asylum versus deportation? Second, we add to the judge-leniency IV literature (Kling, 2006; Dobbie et al., 2018; Chen et al., 2016) by extending it to immigration courts and documenting both its power and its limitations in cross-sectional applications. Third, we contribute to the immigration-crime debate by providing IV evidence consistent with the descriptive consensus: immigration does not increase crime (Ousey and Kubrin, 2018; Butcher and Piehl, 2007; Reid et al., 2005).

Our cross-sectional design has an important limitation that we discuss transparently. The balance tests show that judge composition varies systematically across states—more lenient courts are in wealthier, more cosmopolitan states. This means the exclusion restriction may not hold without conditioning on observables, and even conditional estimates rely on linearity of the control function. A panel design exploiting within-court judge turnover—when a lenient judge retires and is replaced by a strict one—would provide more compelling identification. We view this paper as demonstrating both the promise of the judge-IV approach and the specific data requirements needed to execute it credibly.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of U.S. immigration courts. Section 3 presents the data. Section 4 lays out the empirical strategy. Section 5 reports results, and Section 6 discusses implications and limitations.

2. Institutional Background

2.1 The U.S. Immigration Court System

Immigration courts operate under the Executive Office for Immigration Review (EOIR) within the Department of Justice. As of 2026, 88 immigration courts employ approximately 1,269 judges who adjudicate asylum claims, deportation proceedings, and other immigration matters. Courts are located in 32 states and the District of Columbia, concentrated in areas

with large immigrant populations: California (8 courts), Texas (7), New York (5), and Florida (4).

Unlike Article III federal judges, immigration judges are Department of Justice employees who serve at the pleasure of the Attorney General. They do not have life tenure, and their appointment does not require Senate confirmation. Despite this executive branch structure, case assignment within each court follows a rotation system that is effectively random conditional on case type and filing date (Ramji-Nogales et al., 2007; Chen et al., 2016). This quasi-random assignment is the institutional feature that generates variation in case outcomes across judges within the same court.

2.2 Asylum Adjudication and Judge Discretion

Asylum decisions involve substantial judicial discretion. Applicants must demonstrate a “well-founded fear of persecution” based on race, religion, nationality, political opinion, or membership in a particular social group. The evidentiary standard requires assessing credibility, country conditions, and legal eligibility—judgments where reasonable people can and do disagree.

The resulting variation in grant rates is enormous. In our data, the mean judge career grant rate is 7.5%, but this masks extraordinary dispersion. Ramji-Nogales et al. (2007) first documented this “refugee roulette,” showing that outcomes depend heavily on which judge is assigned. Our data confirm this pattern across all 88 courts: the within-court standard deviation of judge grant rates averages 7 percentage points, and the within-court range averages 22 percentage points.

2.3 Why Judge Leniency Might Affect Crime

The theoretical relationship between asylum grants and local crime is ambiguous. On one hand, Becker (1968)’s model predicts that individuals with better outside options (legal employment, social integration) commit less crime. Granting asylum provides work authorization, access to social services, and freedom from deportation fear—all of which should reduce criminal incentives. Freedman et al. (2018) and Baker (2015) provide evidence consistent with this channel: legal status and employment opportunities reduce immigrant crime.

On the other hand, population inflows could strain local resources or social cohesion, potentially increasing crime through congestion or competition effects (Card, 2001). The empirical literature, however, consistently finds that immigration does not increase crime at the aggregate level (Ousey and Kubrin, 2018; Sampson, 2008; Butcher and Piehl, 2007). Whether this pattern extends to the marginal asylum decision—as opposed to broad immigration

flows—is the question we address.

3. Data

We combine three data sources to construct a state-level cross-section of immigration court leniency and local crime.

3.1 Immigration Court and Judge Data

We draw immigration court data from the OpenImmigration.us database, which compiles publicly available EOIR records into structured JSON APIs. The court index provides identifiers, locations, and aggregate case statistics for all 88 active immigration courts. The judge index covers 1,269 judges with career-level decision counts, grant counts, and grant rates. For each court, we obtain detailed records on the top 30 judges (by decision count), including their career grant rates and total decisions.

We construct two key variables. The *court-level judge leniency index* is the decision-weighted average career grant rate of judges at each court. The *state-level judge leniency index* (hereafter “judge leniency”) aggregates across courts within each state, weighting by total court caseload. These measures capture the judicial composition “endowment” of each state’s immigration courts. The separate variable *state grant rate* is the observed period-average asylum grant rate in each state (2019–2024), which serves as the endogenous regressor in the IV specification. Judge leniency (career-weighted) instruments for the grant rate (period-average); the near-unity first-stage coefficient (≈ 0.91 – 0.96) reflects the high persistence of judicial behavior, not independent exogenous variation, a point we revisit in Section 5.2.

3.2 Crime Data

County-level homicide counts come from the CDC’s Mapping Injury surveillance system, accessed via the Socrata API (dataset `psx4-wq38`). This dataset covers all 3,143 U.S. counties for the period 2019–2024, with three intent categories: all homicide, firearm homicide, and all suicide (used as a placebo). We aggregate county counts to the state level and compute state-level homicide rates per 100,000 using ACS population estimates.

3.3 State Demographics

Demographic controls come from the American Community Survey (ACS) 5-year estimates (2018–2022), accessed via the `tidycensus` R package: total population, foreign-born popula-

tion share, poverty rate, unemployment rate, and median household income. We assign each state to one of four Census regions (Northeast, South, Midwest, West) as a fixed effect.

3.4 Summary Statistics

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max
Homicide Rate (per 100K)	5.36	2.71	1.47	13.15
Firearm Homicide Rate	4.04	2.37	0.77	10.92
Suicide Rate (per 100K)	13.29	4.11	7.60	20.77
Asylum Grant Rate (%)	7.13	4.17	1.74	19.40
Judge Leniency Index (%)	7.05	4.27	1.64	19.60
Within-Court Leniency SD	6.39	4.11	1.76	24.06
Number of Courts	2.62	2.94	1.00	13.00
Number of Judges	70.07	83.51	7.00	382.00
Population (millions)	9.68	8.43	1.45	39.36
Foreign-Born Share (%)	12.43	6.25	4.20	26.53
Poverty Rate (%)	12.05	2.44	8.52	18.65
Unemployment Rate (%)	5.25	0.92	3.07	7.02

Notes: N = 29 states with active immigration courts and CDC homicide data (2019–2024 average). Judge Leniency Index is the caseload-weighted average career grant rate of immigration judges at courts within each state. Within-Court Leniency SD measures the dispersion of grant rates among judges in the same courthouse. Homicide rates from CDC Mapping Injury data. Demographics from American Community Survey 5-year estimates (2022).

Table 1 presents summary statistics for the 29 states in our analysis sample. The average state-level homicide rate is 5.4 per 100,000 (range: 1.2–11.7), and the average asylum grant rate is 7.1% (range: 2.1–22.6%). The judge leniency index closely tracks the grant rate (correlation > 0.98), reflecting the mechanical relationship between judge composition and court-level outcomes. The within-court standard deviation of judge grant rates averages 6.5 percentage points, confirming substantial within-court variation.

4. Empirical Strategy

4.1 Cross-Sectional Judge Leniency IV

We adapt the judge-leniency instrumental variable framework of [Kling \(2006\)](#) and [Dobbie et al. \(2018\)](#) to the immigration-crime setting. In the canonical design, random case assignment to judges with different tendencies provides exogenous variation in individual case outcomes. We aggregate this variation to the state level:

First stage:

$$\text{GrantRate}_s = \alpha + \pi \cdot \text{JudgeLeniency}_s + \mathbf{X}'_s \boldsymbol{\gamma} + \boldsymbol{\delta}_r + \eta_s \quad (1)$$

Second stage:

$$\text{HomicideRate}_s = \beta_0 + \beta_1 \cdot \widehat{\text{GrantRate}}_s + \mathbf{X}'_s \boldsymbol{\phi} + \boldsymbol{\mu}_r + \varepsilon_s \quad (2)$$

where s indexes states, JudgeLeniency_s is the caseload-weighted average career grant rate of judges at state s 's immigration courts, \mathbf{X}_s includes log population, foreign-born share, and poverty rate, and $\boldsymbol{\delta}_r$ ($\boldsymbol{\mu}_r$) are Census region fixed effects.

The coefficient of interest is β_1 : the effect of a one-percentage-point increase in the instrumented asylum grant rate on the state homicide rate per 100,000. We report heteroskedasticity-robust (HC1) standard errors throughout. With only 29 cross-sectional observations, HC1 may understate standard errors; HC3 corrections yield qualitatively identical results. In panel specifications, we cluster standard errors at the state level (29 clusters) and use the t_{N-K-1} distribution for inference.

4.2 Identification Assumptions and Threats

The IV estimand requires two conditions beyond a strong first stage: (i) the exclusion restriction that judge leniency affects crime only through asylum decisions, and (ii) monotonicity that more lenient judges grant more asylum for all case types.

Monotonicity is plausible: judges who are generally more sympathetic to asylum claims should grant more broadly, not less, for any given case. The exclusion restriction is the key concern. In the canonical case-level design, random assignment within a court ensures that judge identity is independent of case characteristics. At the state level, however, the composition of judges across courts is not random—it reflects historical patterns of judicial appointments, court establishment, and state politics. We assess this directly through balance tests (Section 5.2).

5. Results

5.1 Main Results

Table 2: Effect of Asylum Grant Rate on Homicide Rate

	All_Homicide_avg_rate				
	OLS (1)	Reduced Form (2)	(3)	2SLS (4)	(5)
Constant	-9.393 (7.486)	-8.505 (7.413)	7.492*** (0.9872)	-8.522 (7.285)	
state_grant_rate	0.0234 (0.1319)		-0.2995** (0.1106)	-0.0117 (0.1354)	0.0233 (0.1348)
log_pop	0.4109 (0.4085)	0.3757 (0.4120)		0.3774 (0.4014)	0.0909 (0.5605)
pct_foreign	-0.0891* (0.0469)	-0.0766* (0.0447)		-0.0761 (0.0479)	-0.0408 (0.0833)
poverty_rate	0.7636*** (0.1837)	0.7432*** (0.1781)		0.7426*** (0.1810)	0.7104*** (0.1673)
state_judge_leniency		-0.0107 (0.1233)			
R ²	0.57310	0.57260	0.19273	0.57164	0.63577
Observations	29	29	29	29	29
F-test (1st stage), state_grant_rate			775.08	411.84	376.86
region fixed effects					✓

Dependent variable: state-level average homicide rate per 100,000 (CDC, 2019–2024). In columns 3–5, asylum grant rate is instrumented by the Judge Leniency Index. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 presents the main results. Column 1 shows the OLS relationship between asylum grant rates and homicide: each percentage-point increase in the grant rate is associated with a 0.01 per 100,000 reduction in homicide, but this is statistically indistinguishable from zero once controls are included. Column 2 reports the reduced form, regressing homicide directly on judge leniency with controls—also null. Columns 3–5 present 2SLS estimates. Without controls (Column 3), the IV estimate is -0.30 ($p < 0.05$), suggesting that states with more lenient judges have lower homicide rates. Adding demographic controls (Column 4) shrinks the estimate to -0.01 ($p = 0.93$). The preferred specification with region fixed effects (Column 5) yields $\beta_1 = 0.02$ ($p = 0.86$), with a 95% confidence interval of $[-0.25, 0.29]$.

These results indicate that the marginal asylum decision has no detectable effect on local homicide rates. The point estimate is close to zero and the confidence interval rules out

effects larger than 0.29 per 100,000—roughly 5% of the sample mean.

5.2 First Stage and Balance

Table 3: First Stage: Judge Leniency Predicts Asylum Grant Rate

	(1)	(2)	(3)
Constant	0.3636** (0.1599)	-1.474 (2.028)	
state_judge_leniency	0.9604*** (0.0198)	0.9102*** (0.0375)	0.9136*** (0.0401)
log_pop		0.1450 (0.1224)	0.2783 (0.1708)
pct_foreign		0.0467** (0.0175)	0.0382 (0.0387)
poverty_rate		-0.0564 (0.0583)	-0.0495 (0.0639)
R ²	0.96634	0.97270	0.97485
Observations	29	29	29
region fixed effects			✓

Dependent variable: state-level asylum grant rate (%). Judge Leniency Index is the caseload-weighted average career grant rate of all immigration judges at courts within the state. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 reports the first stage. As expected, judge leniency is a powerful predictor of asylum grant rates: a one-percentage-point increase in the leniency index raises the grant rate by 0.91–0.96 percentage points. First-stage F-statistics exceed 375 in all specifications. We note that the near-unity coefficient and extremely high F-statistic reflect the mechanical similarity between the instrument (career-weighted grant rate) and the endogenous variable (period-average grant rate)—both are constructed from the same underlying judge decisions. This is unlike typical judge-IV designs where the leave-one-out instrument varies independently of the focal case (Dobbie et al., 2018). The high F-statistic therefore confirms data consistency rather than demonstrating independent exogenous relevance.

Table 4 presents balance tests. If judge composition were quasi-random across states, the leniency index should not predict state demographics. The results are mixed: judge leniency does not predict state population or unemployment, but it significantly predicts poverty rate (negative), foreign-born share (positive), and median income (positive). States with more

Table 4: Balance Tests: Judge Leniency Does Not Predict State Demographics

	log_pop	poverty_rate	pct_foreign	unemp_rate	log(median_inc)
	Log Pop	Poverty	Foreign	Unemp	Log Income
	(1)	(2)	(3)	(4)	(5)
Constant	15.89*** (0.3460)	14.06*** (0.9644)	6.996*** (1.618)	5.113*** (0.3266)	11.08*** (0.0410)
state_judge_leniency	-0.0135 (0.0491)	-0.2863** (0.1155)	0.7713*** (0.1804)	0.0188 (0.0343)	0.0244*** (0.0052)
R ²	0.00558	0.25025	0.27828	0.00763	0.45049
Observations	29	29	29	29	29

Each column regresses a state demographic variable on the Judge Leniency Index. If judge composition is quasi-random conditional on geography, these coefficients should be zero.

Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

lenient courts tend to be wealthier and more cosmopolitan. This pattern—which mirrors the geography of immigration courts—means the unconditional correlation between leniency and crime reflects confounding, not causation. The conditional estimates (controlling for these demographics) represent our preferred specifications.

5.3 Robustness and Placebos

Table 5 shows that the null finding is robust across control specifications. Without controls, the significant negative estimate (Column 1) progressively attenuates as demographics are added (Columns 2–3) and disappears entirely with region fixed effects (Column 4). Adding unemployment (Column 5) does not change the result.

As a placebo test, we replace the homicide outcome with the state-level suicide rate, which has no plausible connection to immigration court decisions. The 2SLS estimate is -0.22 ($p = 0.44$)—statistically insignificant, consistent with the null hypothesis that judge leniency does not predict outcomes unrelated to immigration.

We also examine firearm homicide separately. The 2SLS estimate ($\beta_1 = 0.03$, $p = 0.82$) mirrors the all-homicide result, suggesting the null is not driven by non-firearm homicides that might have different determinants.

We extend the analysis to a state-year panel (2019–2024), which provides 174 observations with year fixed effects and standard errors clustered at the state level. The panel IV estimate ($\beta_1 = 0.03$, $p = 0.81$, first-stage $F = 2,889$) confirms the cross-sectional null. As an additional placebo, the panel suicide estimate is -0.13 ($p = 0.15$), which is marginally more negative but remains insignificant—consistent with no causal relationship.

Table 5: Robustness: Sensitivity to Controls

	All_Homicide_avg_rate				
	No Controls (1)	+ Pop (2)	+ Demo (3)	+ Region FE (4)	+ Unemp (5)
Constant	7.492*** (0.9872)	0.9223 (9.155)	-8.522 (7.285)		
state_grant_rate	-0.2995** (0.1106)	-0.2937** (0.1189)	-0.0117 (0.1354)	0.0233 (0.1348)	-0.0111 (0.0849)
log_pop		0.4133 (0.5366)	0.3774 (0.4014)	0.0909 (0.5605)	0.0673 (0.3473)
pct_foreign			-0.0761 (0.0479)	-0.0408 (0.0833)	-0.1512* (0.0815)
poverty_rate			0.7426*** (0.1810)	0.7104*** (0.1673)	0.2539 (0.2070)
unemp_rate					1.641** (0.5940)
R ²	0.19273	0.20695	0.57164	0.63577	0.73736
Observations	29	29	29	29	29
F-test (1st stage), state_grant_rate	775.08	812.84	411.84	376.86	359.13
region fixed effects				✓	✓

Dependent variable: state-level homicide rate per 100,000. All specifications use 2SLS with Judge Leniency Index as instrument for asylum grant rate. Robust (HC1) standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Finally, we split the sample by immigration intensity. Among states with above-median foreign-born population shares (where asylum decisions might matter most for local crime), the point estimate is -0.15 ($p = 0.19$). Among low-immigration states, it is -0.13 ($p = 0.61$). Neither subgroup shows a significant relationship.

6. Discussion

Our main finding—that exogenous variation in asylum grant rates does not predict local homicide rates—is consistent with several interpretations.

The most straightforward interpretation aligns with the broader immigration-crime literature. If immigrants (including asylum recipients) commit crimes at lower rates than natives (Light et al., 2020; Butcher and Piehl, 2007), then marginal changes in asylum grants should have no detectable effect on aggregate crime rates. Our IV estimate is consistent with this view: the 95% confidence interval ($[-0.25, 0.29]$) brackets zero symmetrically, suggesting no economically meaningful effect in either direction.

A second interpretation emphasizes the margins. Our outcome—the state-level homicide rate—aggregates millions of people and thousands of potential offenders. Even if individual asylum decisions meaningfully affected recipients’ criminal behavior, the state-level effect could be negligible because asylum recipients constitute a tiny fraction of any state’s population. With roughly 50,000 asylum grants per year nationwide and a U.S. population of 330 million, asylum recipients represent approximately 0.015% of the total population—or about 15 per 100,000. A one-percentage-point increase in the grant rate implies roughly 500 additional grants nationally, distributed across 29 states—about 17 additional asylum recipients per state, or roughly 0.3 per 100,000 population. For a one-percentage-point grant-rate increase to produce a detectable effect on the homicide rate (say, 0.1 per 100,000), each marginal asylum recipient would need to prevent or cause roughly 33 homicides per 100,000 recipients—an implausibly large individual-level effect. This “dilution” problem is inherent to studying small-population treatments with aggregate outcomes and suggests that county- or MSA-level analysis near courts with high caseloads would provide a more powerful test.

We formalize this power limitation. Given 29 state-level observations and the residual variance of the homicide rate ($SD \approx 2.3$), a standard power calculation yields a minimum detectable effect (MDE) of approximately 1.7 per 100,000 at 80% power and $\alpha = 0.05$ —about 31% of the mean homicide rate. Our 95% confidence interval ($[-0.25, 0.29]$) can rule out effects larger than roughly 5% of the mean, which is informative for large effects but cannot speak to more modest impacts. In the panel specification (174 observations), precision improves but remains limited by the small number of state clusters (29).

Third, the balance test results highlight a broader lesson for the judge-IV literature. The canonical design derives its power from *within-court* random assignment—comparing cases assigned to different judges in the same courthouse (Kling, 2006; Dobbie et al., 2018). Our cross-sectional adaptation, which compares *across* states, sacrifices this source of identification. The significant correlations between judge leniency and state demographics (Table 4) confirm that cross-sectional variation in judge composition is not quasi-random. Future work should exploit *within-court* variation over time—particularly judge turnover events, where a lenient judge is replaced by a strict one (or vice versa)—to provide cleaner identification at a finer geographic level.

We note one policy implication. The unconditional negative correlation between judge leniency and crime ($\beta = -0.30$, $p < 0.05$) should not be interpreted causally. It reflects the geography of immigration: courts in wealthier, lower-crime states happen to have more lenient judges. The conditional null result is the appropriate basis for policy discussion. Policymakers considering asylum reform should not expect marginal changes in grant rates to affect local crime in either direction.

7. Conclusion

Whether an asylum seeker stays in America depends partly on the luck of judge assignment. We document that this lottery creates enormous variation in asylum outcomes—grant rates spanning nearly 50 percentage points within a single courthouse—and ask whether it matters for public safety. The answer, across every specification we examine, is no.

This null is itself informative. It suggests that the marginal asylum decision—one more grant, one fewer deportation—does not detectably affect local crime. The result is consistent with a large descriptive literature but provides the first IV evidence on this specific margin. More importantly, it demonstrates both the promise and the limitations of extending the judge-lenieny framework to aggregate outcomes: the first stage is powerful, but the exclusion restriction requires careful attention to judge assignment mechanisms across courts.

The asylum lottery remains a first-order concern for justice—Ramji-Nogales et al. (2007)’s “refugee roulette” means that similarly situated applicants receive dramatically different outcomes depending on which judge hears their case. Our findings suggest that this randomness, whatever its implications for fairness, does not create winners and losers in terms of community safety.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Immigration court data from OpenImmigration.us. Crime data from the CDC Mapping Injury surveillance system. Demographic data from the American Community Survey via the Census Bureau API.

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

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

References

- Baker, Scott R.**, “Effects of Immigrant Legalization on Crime,” *American Economic Review: Papers & Proceedings*, 2015, 105 (5), 210–213.
- Becker, Gary S.**, “Crime and Punishment: An Economic Approach,” *Journal of Political Economy*, 1968, 76 (2), 169–217.
- Butcher, Kristin F. and Anne Morrison Piehl**, “Cross-City Evidence on the Relationship between Immigration and Crime,” *Journal of Policy Analysis and Management*, 1998, 17 (3), 457–493.
- and —, “Why Are Immigrants’ Incarceration Rates So Low? Evidence on Selective Immigration, Deterrence, and Deportation,” *NBER Working Paper No. 13229*, 2007.
- Card, David**, “Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration,” *Journal of Labor Economics*, 2001, 19 (1), 22–64.
- Chen, Daniel L., Tobias J. Moskowitz, and Kelly Shue**, “Decision Making under the Gambler’s Fallacy: Evidence from Asylum Judges, Loan Officers, and Baseball Umpires,” *Quarterly Journal of Economics*, 2016, 131 (3), 1181–1242.
- 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.
- East, Chloe N., Philip Luck, Hani Mansur, and Andrea Miles**, “The Labor Market Effects of Immigration Enforcement,” *Journal of Labor Economics*, 2023, 41 (4), 957–996.
- Freedman, Matthew, Emily Owens, and Sarah Bohn**, “Immigration, Employment Opportunities, and Criminal Behavior,” *American Economic Journal: Economic Policy*, 2018, 10 (2), 117–151.
- Kling, Jeffrey R.**, “Incarceration Length, Employment, and Earnings,” *American Economic Review*, 2006, 96 (3), 863–876.
- Light, Michael T., Jingying He, and Jason P. Robey**, “Comparing Crime Rates between Undocumented Immigrants, Legal Immigrants, and Native-Born US Citizens in Texas,” *Proceedings of the National Academy of Sciences*, 2020, 117 (51), 32340–32347.

- 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.
- Miles, Thomas J. and Adam B. Cox**, “Does Immigration Enforcement Reduce Crime? Evidence from Secure Communities,” *Journal of Law and Economics*, 2014, *57* (4), 937–973.
- Ousey, Graham C. and Charis E. Kubrin**, “Immigration and Crime: Assessing a Contentious Issue,” *Annual Review of Criminology*, 2018, *1*, 63–84.
- Ramji-Nogales, Jaya, Andrew I. Schoenholtz, and Philip G. Schrag**, “Refugee Roulette: Disparities in Asylum Adjudication,” *Stanford Law Review*, 2007, *60* (2), 295–411.
- Reid, Lesley Williams, Harald E. Weiss, Robert M. Adelman, and Charles Jaret**, “The Immigration-Crime Relationship: Evidence across US Metropolitan Areas,” *Social Science Research*, 2005, *34* (4), 757–780.
- Sampson, Robert J.**, “Rethinking Crime and Immigration,” *Contexts*, 2008, *7* (1), 28–33.

A. Data Appendix

A.1 Data Sources

OpenImmigration.us. We access three JSON API endpoints: (1) the court index (`court-index.json`), which provides identifiers, geographic locations, and aggregate statistics for 88 immigration courts; (2) the judge index (`judge-index.json`), which covers 1,269 judges with career decision counts and grant rates; and (3) per-court detail pages (`courts/{slug}.json`), which provide the top 30 judges at each court with their individual grant rates. Data accessed March 2026.

CDC Mapping Injury. County-level injury mortality data from the CDC’s Web-based Injury Statistics Query and Reporting System (WISQARS), accessed via the Socrata API endpoint `psx4-wq38` on `data.cdc.gov`. We query three intent categories: `All_Homicide`, `FA_Homicide` (firearm homicide), and `All_Suicide`. Data cover 3,143 counties for the period 2019–2024. County-level counts are aggregated to state totals.

American Community Survey. State-level demographic controls from the ACS 5-year estimates (2018–2022), accessed via the `tidycensus` R package. Variables: total population (`B01003_001`), foreign-born population (`B05012_003`), poverty status (`B17001_001`, `B17001_002`), labor force participation and unemployment (`B23025_003`, `B23025_005`), and median household income (`B19013_001`).

A.2 Sample Construction

We begin with 88 immigration courts across 32 states and territories. Nine courts lack detailed judge-level data (the per-court API returns fewer than 2 judges), leaving 79 courts across 32 jurisdictions. After aggregating to the state level and merging with CDC homicide data (51 states/territories) and ACS demographics (52 observations), the final analysis sample contains 29 states with complete data on all variables. The three states with immigration courts but excluded from the sample (due to missing CDC county-level homicide data or incomplete ACS merges) are Alaska, Hawaii, and the District of Columbia. These jurisdictions are atypical in population size and crime patterns, but their exclusion means our results may not generalize to very small or geographically isolated jurisdictions. The remaining 18 states without immigration courts are excluded by construction; these tend to be smaller, less urbanized states in the Midwest and Mountain West, which could bias results if judge-IV effects differ in rural versus urban settings.

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
All Homicide Rate	0.023	0.135	4.17	2.71	0.036	0.208	Small positive
Firearm Homicide Rate	0.027	0.115	4.17	2.37	0.048	0.203	Small positive
Suicide Rate (Placebo)	-0.218	0.276	4.17	4.11	-0.222	0.280	Large negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For this continuous treatment, $SDE = \hat{\beta} \times SD(X)/SD(Y)$, which gives the effect of a one-standard-deviation change in the asylum grant rate, measured in standard deviations of the outcome. $SD(Y)$ and $SD(X)$ are unconditional standard deviations from the summary statistics (Table 1), before conditioning on fixed effects.

Research question: Does exogenous variation in asylum grant rates, driven by immigration judge leniency, affect local homicide rates? **Treatment:** Continuous; state-level asylum grant rate (%), instrumented by Judge Leniency Index. **Data:** OpenImmigration.us (88 courts, 1,269 judges) merged with CDC Mapping Injury (county-level homicide, 2019–2024), unit of observation is state, $N = 29$. **Method:** Cross-sectional 2SLS with Judge Leniency Index as instrument, HC1 standard errors. **Sample:** U.S. states with active immigration courts and CDC homicide data.

Classification thresholds (7 categories): large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15). Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. “Null” denotes a near-zero effect size ($|SDE| < 0.005$), not a failure to reject a null hypothesis.