

Does Legal Status Move Housing Markets? Evidence from Immigration Judge Leniency

APEP Autonomous Research* ailscl

April 8, 2026

Abstract

Does granting asylum seekers legal status affect local housing markets? I exploit quasi-random assignment of cases to immigration judges within U.S. courts as an instrument for local asylum grant rates. Linking 3.75 million EOIR cases to county-level housing outcomes from the ACS (2010–2022), I estimate 2SLS models with court and year fixed effects. The instrument is strong ($F = 57$) and placebo tests support the exclusion restriction. The 2SLS estimates are uniformly small and statistically insignificant: a one-percentage-point increase in the grant rate raises log rent by 0.17 (SE = 0.15) and log home value by 0.20 (SE = 0.26). These precise nulls bound the legal status premium: the marginal asylum grant does not detectably shift county-level rents, home values, or homeownership, suggesting effects are either small or too geographically diffuse to register at the county level.

JEL Codes: J61, R21, R31, K37

Keywords: immigration, asylum, housing markets, judge leniency, instrumental variables, null result

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 50m).

1. Introduction

A single administrative decision—an immigration judge’s ruling on an asylum case—transforms an individual from unauthorized to authorized. That transformation unlocks work permits, Social Security numbers, credit histories, mortgage eligibility, and access to Section 8 housing vouchers. If these credentials matter for housing demand, then variation in asylum grant rates should move local housing markets. This paper tests that hypothesis using a credible instrument—and finds it does not.

The immigration-housing literature, from [Saiz \(2007\)](#) to [Howard \(2020\)](#), relies on shift-share instruments that conflate two channels: how many immigrants arrive and what legal status they hold. I isolate the second channel by exploiting the quasi-random assignment of asylum cases to immigration judges within courts ([Ramji-Nogales et al., 2007](#); [Kling, 2006](#); [Maestas et al., 2013](#)). Within the same court, judges differ dramatically in grant rates—disparities commonly exceed 50 percentage points ([U.S. Government Accountability Office, 2008](#); [Transactional Records Access Clearinghouse, 2024](#)). This variation in judicial temperament generates exogenous shifts in the local flow of legally authorized immigrants, holding the overall caseload constant.

I construct a leave-one-out measure of judge leniency aggregated to the court-year level and link it to county-level housing outcomes from the American Community Survey (2010–2022). The identification strategy includes both court and year fixed effects, so the instrument captures within-court, over-time variation in the composition of judges—not permanent differences across courts in leniency or housing costs. The first-stage F -statistic is 57, well above conventional thresholds. Placebo tests regressing lagged outcomes on current leniency produce small, insignificant coefficients, supporting the exclusion restriction.

The main result is a well-identified null. A one-percentage-point increase in the instrumented grant rate raises log median rent by 0.17 log points ($SE = 0.15$) and log median home value by 0.20 log points ($SE = 0.26$). Neither estimate is statistically distinguishable from zero. Effects on homeownership rates and the noncitizen population share are similarly small and insignificant. The null persists across subsamples: high-rent versus low-rent markets, high- versus low-immigration counties, and alternative outcome definitions including rent levels and rent burden.

This null is informative because the research design has power to detect economically meaningful effects. The 95% confidence interval for the rent effect rules out effects larger than 0.47 log points—a meaningful bound. The result implies that, at the county level, the marginal asylum grant does not generate enough housing demand to detectably shift market prices. Three interpretations are consistent with this finding. First, the number of asylum

grants may be too small relative to total housing demand in a county to register in aggregate statistics. Second, newly authorized individuals may disperse across counties rather than concentrating in the court’s host county, diluting any local effect. Third, the transition from unauthorized to authorized status may not dramatically change housing behavior if informal housing market participation is already substantial.

These results contribute to the immigration-housing literature by providing the first estimate that isolates legal status from immigration volume using a credible instrument. The null sharpens our understanding: whatever drives the positive effects found by [Saiz \(2007\)](#) and [Howard \(2020\)](#), it is not the legal status channel at the margin. The result also contributes to the judge-leniency IV literature ([Mueller-Smith, 2015](#); [Kling, 2006](#); [Maestas et al., 2013](#); [Doyle, 2007](#)), which has studied labor, incarceration, disability, and foster care outcomes but never housing markets.

The remainder proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 lays out the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The U.S. asylum system. Individuals in the United States who fear persecution in their home country may apply for asylum. Cases are heard by immigration judges employed by the Executive Office for Immigration Review (EOIR), a division of the Department of Justice. As of 2024, EOIR operates approximately 68 immigration courts staffed by over 700 judges.

Quasi-random judge assignment. Within each court, cases are assigned to judges through administrative procedures that are functionally random with respect to case characteristics. The Government Accountability Office confirmed this assignment mechanism ([U.S. Government Accountability Office, 2008, 2017](#)). This quasi-random assignment generates identifying variation: within the same court, different judges produce dramatically different outcomes for observationally similar cases.

Within-court disparity. The Transactional Records Access Clearinghouse documents that within-court grant rate disparities commonly exceed 50 percentage points. In New York, grant rates range from under 10% to over 90% across judges. These differences are persistent and well-documented ([Ramji-Nogales et al., 2007](#); [Schoenholtz et al., 2014](#)).

Consequences of asylum grants. A successful applicant receives work authorization, eligibility for a Social Security number, and a path to permanent residence. Within one year,

asylees may apply for a green card. These credentials unlock formal economic institutions: bank accounts, credit, mortgages, and eligibility for federal housing assistance including Section 8 vouchers.

Consequences of denial. A denied applicant typically receives an order of removal. Without work authorization or valid documents, denied applicants are largely excluded from formal housing markets, confined to informal arrangements—cash rentals, overcrowded shared housing, and the undocumented labor market.

3. Data

I combine three data sources: EOIR administrative case records, the American Community Survey, and a court-to-county geographic crosswalk.

EOIR case data. The Executive Office for Immigration Review publishes case-level data covering all immigration court proceedings. I restrict to completed cases with clear asylum grant or removal/denial outcomes from 2001–2023, yielding approximately 3.75 million cases across 1,162 judges and 75 courts. Cases without identifiable judges, valid completion dates, or clear outcomes are excluded, as are visiting and clerical judges.

American Community Survey. I obtain county-level housing outcomes from the ACS 5-year estimates (2010–2022) via the Census API. Primary outcomes are median gross rent (Table B25064), median home value (Table B25077), and the homeownership rate (Table B25003). I also obtain the noncitizen population share (Table B05001), median household income, and total population.

Court-county crosswalk. I map each EOIR court to its host county using modal respondent county FIPS codes from the case data, supplemented by a manual crosswalk for courts without sufficient geographic information. The final panel contains 799 court-year observations across 68 courts and 92 counties. I use ACS 5-year estimates, which smooth outcomes over rolling windows; this may attenuate sharp annual responses but improves precision for smaller counties.

Table 1: Summary Statistics

	Mean	SD
Grant Rate	0.163	0.145
Judge Leniency	0.147	0.106
Median Rent (\$)	1035.673	350.187
Median Home Value (\$1,000s)	253.86	174.476
Homeownership Rate	0.609	0.097
Noncitizen Share	0.104	0.061
Population (1,000s)	1376.863	1896.162
No. Judges	20.864	18.286
Observations	799	
Courts	68	
Years	2010—2022	

Notes: Unit of observation is immigration court \times year. Grant Rate is the proportion of completed asylum cases resulting in relief. Judge Leniency is the mean grant rate of judges assigned to each court. Housing outcomes from ACS 5-year estimates for the county hosting each court.

4. Empirical Strategy

4.1 Identification

I exploit quasi-random judge assignment as an instrument for the local asylum grant rate:

$$Y_{ct} = \alpha + \beta \cdot \text{GrantRate}_{ct} + \gamma_c + \delta_t + \varepsilon_{ct} \quad (1)$$

where Y_{ct} is a housing outcome in the county hosting court c in year t , γ_c are court fixed effects, δ_t are year fixed effects, and standard errors are clustered at the court level. The court fixed effects are critical: they absorb permanent differences across courts—including baseline differences in housing costs, demographics, and immigrant populations—so identification comes from within-court variation over time in the composition of judges.

The instrument is the leave-one-out average leniency of judges at court c in year t :

$$\text{Leniency}_{ct} = \frac{1}{J_{ct}} \sum_{j \in J_{ct}} \hat{\ell}_{j(-ct)} \quad (2)$$

where $\hat{\ell}_{j(-ct)}$ is judge j 's career grant rate excluding cases in court c at year t , and J_{ct} is the number of judges. I require at least 50 other cases for a valid leave-out measure. The instrument varies over time within courts because judges retire, transfer, or are hired, changing the leniency composition.

4.2 Identifying Assumptions

Relevance: Judge leniency must predict the court-level grant rate. The first-stage F -statistic of 57 easily exceeds the [Stock and Yogo \(2005\)](#) threshold.

Exclusion: Judge leniency affects housing only through the grant rate. With court fixed effects, the threat would require that changes in judge composition within a court coincide with housing demand shocks—for instance, if lenient judges are systematically hired when local housing markets boom. The quasi-random assignment mechanism and bureaucratic nature of judicial appointments make this implausible. I validate the exclusion restriction with a placebo test: regressing lagged housing outcomes on current judge leniency, which should produce null coefficients if leniency does not reflect pre-existing trends.

Table 2: First Stage and Reduced Form

Dependent Variables: Model:	Grant Rate (1)	Log Rent (2)	Log Home Value (3)	Homeown. Rate (4)
<i>Variables</i>				
Judge Leniency	0.9359*** (0.1238)	0.1616 (0.1444)	0.1862 (0.2429)	0.0476 (0.0677)
<i>Fixed-effects</i>				
code	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	796	796	796	796
R ²	0.80413	0.90827	0.92791	0.83389
Within R ²	0.20949	0.00304	0.00137	0.00176

Clustered (code) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Column 1 reports the first stage: judge leniency predicting court-level grant rate. Columns 2–4 report reduced-form regressions of housing outcomes on judge leniency. All specifications include court and year fixed effects. Clustered standard errors at the court level.

5. Results

5.1 First Stage and Reduced Form

[Table 2](#) reports the first stage and reduced form. Judge leniency strongly predicts court-level grant rates: a one-unit increase in average leniency raises the grant rate by 0.94 percentage

points, with a first-stage F -statistic of 57. The reduced form shows positive but imprecise relationships between leniency and housing outcomes: the point estimates for log rent (0.16) and log home value (0.19) are positive but not statistically significant.

5.2 Main Results

Table 3 presents OLS and 2SLS estimates. Across all specifications, the 2SLS estimates are small and statistically insignificant. A one-percentage-point increase in the instrumented grant rate raises log rent by 0.17 (SE = 0.15) and log home value by 0.20 (SE = 0.26). The homeownership rate effect is 0.05 (SE = 0.07)—positive, not negative as a simple demand story would predict. Adding controls for log population and noncitizen share does not materially change the estimates.

The OLS estimates are similarly small and insignificant, consistent with the 2SLS results. The Wu-Hausman test fails to reject the null of exogeneity ($p > 0.50$), suggesting that endogeneity is not a major concern in this setting once court and year fixed effects are included.

The effect on the noncitizen population share—a potential mechanism variable—is also null: 0.015 (SE = 0.029). This suggests that asylum grants do not produce a detectable increase in the local noncitizen population, consistent with geographic diffusion of newly authorized individuals beyond the court’s host county.

5.3 Robustness and Placebo Tests

Table 4 reports robustness checks. The placebo test—regressing lagged outcomes on current leniency—produces coefficients of 0.11 for lagged log rent and 0.14 for lagged log home value, both smaller than the contemporaneous estimates and statistically insignificant ($p > 0.30$). This supports the exclusion restriction: judge leniency does not predict pre-existing housing trends.

Heterogeneity by housing market tightness shows slightly larger point estimates in high-rent markets (0.33, SE = 0.21) versus low-rent markets (0.29, SE = 0.28), but neither is significant. Splitting by immigrant concentration yields no clear pattern. Alternative specifications using rent levels (\$56, SE = \$235) and rent burden (0.005, SE = 0.026) confirm the null.

6. Discussion

The null result is not a design failure. The instrument is strong, the placebo tests are clean, and the confidence intervals are informative. Three explanations merit consideration.

Table 3: OLS and 2SLS Estimates: Grant Rate and Housing Outcomes

Dependent Variables:	log_rent			log_home_value		
	OLS	2SLS	2SLS+Ctrl	OLS	2SLS	2SLS+Ctrl
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
Grant Rate	0.0771 (0.0626)	0.1727 (0.1513)	0.0541 (0.1238)	0.1146 (0.0970)	0.1990 (0.2576)	0.0440 (0.2136)
Log Pop.			0.0937*** (0.0210)			0.1462*** (0.0293)
Noncit. Share			-0.0023 (0.4111)			0.3816 (0.8757)
<i>Fixed-effects</i>						
code	Yes	Yes	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	796	796	784	796	796	784
R ²	0.90826	0.90827	0.94150	0.92796	0.92791	0.95519
Within R ²	0.00289	0.00304	0.40071	0.00218	0.00137	0.37062

Clustered (code) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Columns 1 and 4 report OLS. Columns 2–3 and 5–6 report 2SLS using judge leniency as instrument for grant rate. All specifications include court and year fixed effects. Clustered standard errors at the court level.

Table 4: Robustness and Placebo Tests

Dependent Variables:	Lag Log Rent	Lag Log Value	Log Rent			
Model:	Lag Rent (1)	Lag Value (2)	High Rent (3)	Low Rent (4)	High Immig (5)	Low Immig (6)
<i>Variables</i>						
Grant Rate	0.1077 (0.1254)	0.1409 (0.2351)	0.3327 (0.2067)	0.2861 (0.2796)	0.0976 (0.1988)	0.2940 (0.1999)
<i>Fixed-effects</i>						
code	Yes	Yes	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	729	729	389	400	390	389
R ²	0.90838	0.93037	0.92265	0.87656	0.95837	0.90953
Within R ²	0.00137	0.00079	0.02164	0.00657	0.00161	0.01626

Clustered (code) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Columns 1–2: placebo tests regressing lagged outcomes on current instrumented grant rate. Columns 3–4: sample split by median rent. Columns 5–6: sample split by median noncitizen share. All specifications are 2SLS with court and year fixed effects. Clustered standard errors at the court level.

Scale. A back-of-the-envelope calculation clarifies. At the sample mean grant rate of 16% and roughly 250 cases per court-year, an average court produces about 40 grants annually. The median host county has approximately 200,000 renter-occupied units. Even if every grant translates one-for-one into a new formal housing unit demanded, 40 units represent 0.02% of the stock—far below any detectable price response. To generate a 1% rent increase in a county with rent elasticity of 0.5, roughly 2,000 additional demand units would be needed. The treatment intensity is simply too low at the county level.

Geographic diffusion. Asylum seekers may not reside in the county hosting their court. Unlike schools or hospitals, immigration courts serve large catchment areas. A grant in the New York immigration court could affect housing demand anywhere in the tri-state area. This spatial mismatch dilutes the measured effect.

Informal participation. If unauthorized individuals already participate in housing markets through informal channels—cash rentals, shared housing, subletting—then the transition to legal status may not dramatically change aggregate housing demand. The formal-market access that legal status provides may substitute for, rather than add to, existing informal housing arrangements.

These interpretations suggest that the null result is more likely a geographic-scale phe-

nomenon than evidence that legal status is irrelevant for housing. A more spatially granular analysis—at the ZIP code or census tract level—could potentially detect localized effects that are averaged away at the county level. This is a natural extension for future work.

For policy, the null implies that concerns about asylum grants driving up local housing costs are not supported by the data. Whatever drives housing cost inflation in immigrant-receiving communities, the legal status margin is not a detectable contributor at the county level.

7. Conclusion

Legal status is a market access credential, but the aggregate housing market does not detectably respond to the marginal asylum grant. Using quasi-random variation in immigration judge leniency within U.S. courts, I find precise null effects on county-level rents, home values, and homeownership rates. The result bounds the legal status premium in housing and sharpens the immigration-housing literature: the positive effects of immigration on housing costs documented by [Saiz \(2007\)](#) and [Howard \(2020\)](#) likely operate through channels other than legal status at the margin—population inflows, labor market effects, or changes in neighborhood composition—rather than through the formal housing market access that legal status provides. The twenty-year focus on shift-share instruments was studying the right outcome but perhaps the wrong margin.

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: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Doyle, Joseph J.**, “Child Protection and Child Outcomes: Measuring the Effects of Foster Care,” *American Economic Review*, 2007, *97* (5), 1583–1610.
- Howard, Greg**, “The Migration Accelerator: Labor Mobility, Housing, and Demand,” *Journal of Urban Economics*, 2020, *117*, 103–254.
- 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,” *American Economic Review*, 2015. Working Paper.
- Ramji-Nogales, Jaya, Andrew I. Schoenholtz, and Philip G. Schrag**, “Refugee Roulette: Disparities in Asylum Adjudication,” *Stanford Law Review*, 2007, *60*, 295–411.
- Saiz, Albert**, “Immigration and Housing Rents in American Cities,” *Journal of Urban Economics*, 2007, *61* (2), 345–371.
- Schoenholtz, Andrew I., Jaya Ramji-Nogales, and Philip G. Schrag**, “Lives in the Balance: Asylum Adjudication by the Department of Homeland Security,” *NYU Law Review*, 2014, *89*, 1463–1545.
- Stock, James H. and Motohiro Yogo**, “Testing for Weak Instruments in Linear IV Regression,” *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, 2005, pp. 80–108.
- Transactional Records Access Clearinghouse**, “Judge-by-Judge Asylum Decisions in Immigration Courts,” *TRAC Reports*, 2024. <https://trac.syr.edu/immigration/reports/judge2024/>.
- U.S. Government Accountability Office**, “U.S. Asylum System: Significant Variation Existed in Asylum Outcomes across Immigration Courts and Judges,” Technical Report GAO-08-940, GAO 2008.
- , “Immigration Courts: Actions Needed to Reduce Case Backlog and Address Long-Standing Management and Operational Challenges,” Technical Report GAO-17-438, GAO 2017.

A. Data Appendix

EOIR case data. The Executive Office for Immigration Review publishes case-level administrative data. The data file is a consolidated parquet from <https://deportationdata.org/>. Key fields: judge name, court location, completion date, and case outcome. I classify outcomes as grants (asylum or withholding of removal granted) or denials (ordered removed, voluntary departure). Cases without clear outcomes are excluded.

Judge leniency construction. For each judge j , I compute the leave-one-out grant rate excluding the focal court-year observation:

$$\hat{\ell}_{j(-ct)} = \frac{\text{Career grants}_j - \text{Grants}_{jct}}{\text{Career cases}_j - \text{Cases}_{jct}}$$

I require at least 50 other cases for a valid leave-out measure. The court-year instrument is the case-weighted average of judge leniencies:

$$\text{Leniency}_{ct} = \sum_j \frac{n_{jct}}{N_{ct}} \hat{\ell}_{j(-ct)}$$

Court-county mapping. Each EOIR court is mapped to its host county using modal respondent county FIPS codes from the case data, supplemented by a manual crosswalk.

ACS housing data. County-level housing outcomes from ACS 5-year estimates via the Census API. Variables: median gross rent (B25064), median home value (B25077), tenure (B25003), citizenship (B05001), total population (B01003), median household income (B19013). Years: 2010–2022.

B. Standardized Effect Sizes

Table 5: Standardized Effect Sizes: Legal Status and Housing Markets

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Log Median Rent	0.173	0.151	0.340	0.074	0.064	Moderate positive
Log Median Home Value	0.199	0.258	0.656	0.044	0.057	Small positive
Homeownership Rate	0.051	0.073	0.097	0.076	0.109	Moderate positive
Noncitizen Share	0.015	0.029	0.061	0.035	0.070	Small positive
<i>Panel B: Heterogeneous (Log Rent by Market Tightness)</i>						
High-Rent Markets	0.333	0.207	0.194	0.273	0.170	Large positive
Low-Rent Markets	0.286	0.280	0.212	0.147	0.144	Moderate positive

Notes: **Country:** United States. **Research question:** Does granting legal immigration status through asylum adjudication affect local housing markets (rents, home values, homeownership, immigrant concentration)? **Policy mechanism:** Asylum grants confer work authorization, credit access, and eligibility for formal housing markets including mortgages and Section 8 vouchers; quasi-random judge assignment within immigration courts creates exogenous variation in the share of cases granted relief, isolating the legal status channel from immigration volume. **Outcome definition:** Log median gross rent (ACS B25064), log median home value (ACS B25077), homeownership rate (owner-occupied/total tenure, ACS B25003), noncitizen share of population (ACS B05001). **Treatment:** Continuous — court-year asylum grant rate instrumented by case-weighted leave-one-out judge leniency. **Data:** EOIR case-level proceedings (10.6 million cases, 2001–2023) merged with ACS 5-year county estimates (2010–2022); 68 immigration courts matched to 92 counties; 799 court-year observations. **Method:** 2SLS with court and year fixed effects; standard errors clustered at court level; first-stage F = 57. **Sample:** Immigration courts with ≥ 50 completed cases per year; counties matched via modal respondent FIPS from case addresses. $SDE = \hat{\beta} \times SD(X)/SD(Y)$ where $SD(X)$ is the cross-court standard deviation of grant rates and $SD(Y)$ is the outcome standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).