

The Rejection Drain: Patent Examiner Leniency and Inventor Interstate Mobility

APEP Autonomous Research* @olafdrw

March 31, 2026

Abstract

Every year, the USPTO rejects roughly 250,000 patent applications. I present evidence that rejection is associated with inventors leaving their state. Using examiner assignment within art units across 4.5 million inventor-application pairs (2002–2014), I instrument for rejection with leave-one-out examiner grant rates. The IV estimate suggests rejection increases the probability of interstate relocation by 1.1 percentage points—9 percent of baseline mobility. The reduced form declines monotonically across examiner leniency quintiles. Effects concentrate among solo inventors (2.5 pp) and those with prior patents (2.7 pp), consistent with the destruction of location-specific intellectual property rents. However, balance tests reveal imperfect randomization within art-unit-by-year cells, warranting cautious interpretation. If confirmed in higher-quality inventor panels, this “rejection drain” represents an unrecognized channel through which federal patent policy reshapes the geography of innovation.

JEL Codes: O31, O34, J61, R23

Keywords: patent rejection, inventor mobility, examiner leniency, brain drain, knowledge workers

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 27m).

1. Introduction

A patent rejection destroys more than an idea’s legal protection—it may also destroy the reason an inventor stays put. Location-specific rents from intellectual property tie knowledge workers to particular labor markets: the firm that files the patent, the network of collaborators who contributed, the regional ecosystem of complementary innovators. When a patent examiner rejects an application, those location-specific rents evaporate, and the inventor’s outside options in other states become relatively more attractive. If rejection systematically triggers interstate relocation, then patent examination policy—ostensibly about allocating property rights—also functions as an unintended geographic redistribution mechanism for human capital.

This paper asks whether patent rejection causes inventors to relocate across state lines. The question matters because state and local governments spend billions annually on place-based innovation policies—R&D tax credits, innovation districts, university partnerships—premised on the assumption that attracting and retaining knowledge workers drives local economic growth (Moretti and Wilson, 2017; Gruber and Johnson, 2019). If the federal patent system simultaneously drains inventors from states through a mechanism these policies cannot offset, the effective return to place-based innovation spending is lower than conventional estimates suggest.

Identifying the causal effect of rejection on mobility is challenging because inventors who receive rejections may differ systematically from those who receive grants. Rejected applications may be lower quality, filed by more mobile inventors, or concentrated in technology fields with different labor market dynamics. I address this endogeneity using the examiner leniency instrumental variables design pioneered by Sampat and Williams (2019) and widely applied in the innovation literature (Galasso and Schankerman, 2015; Farre-Mensa et al., 2020; Righi and Simcoe, 2023). The USPTO assigns patent applications to examiners quasi-randomly within art units (technology-specific examination groups). Examiners vary substantially in their propensity to grant patents: the within-art-unit standard deviation of grant rates is 0.156, and the interquartile range spans 0.50 to 0.80. I construct a leave-one-out measure of examiner leniency within art-unit-by-filing-year cells and use it to instrument for rejection.

The first stage is exceptionally strong: a one-standard-deviation increase in examiner leniency reduces the probability of rejection by 52.6 percentage points ($F = 15,179$). The instrument satisfies the monotonicity condition: across leniency quintiles, both rejection rates and mobility rates decline monotonically from the strictest (Q1: 54.0% rejected, 13.6% moved) to the most lenient examiners (Q5: 6.6% rejected, 9.7% moved).

The IV estimate indicates that patent rejection increases the probability of interstate relocation by 1.1 percentage points ($SE = 0.0025$), representing a 9 percent increase relative to the baseline mobility rate of 12.2 percent. This estimate is stable across specifications—with and without controls for team size, prior grants, and small entity status—and robust to varying the minimum cell size for art-unit-by-year groups. The OLS estimate (1.3 pp) is modestly larger, consistent with positive selection: inventors who file from multiple locations are mechanically more likely to both receive rejections and appear as movers.

Heterogeneity in the IV estimates adjudicates the economic mechanism. If rejection triggers mobility by destroying location-specific IP rents, the effect should concentrate among inventors with more to lose from staying and more capacity to move. Consistent with this prediction, solo inventors—who lack firm-based anchoring—show effects nearly three times larger than team inventors (2.5 pp vs. 0.9 pp). Experienced inventors with at least one prior patent—who have established reputations and outside options—respond four times more than novices (2.7 pp vs. 0.6 pp). These patterns are difficult to reconcile with alternative explanations such as discouraged filing or random career shocks.

This paper contributes to three literatures. First, it adds to the growing body of work using examiner leniency designs to study the consequences of patent decisions ([Galasso and Schankerman, 2015](#); [Farre-Mensa et al., 2020](#); [Sampat and Williams, 2019](#); [Righi and Simcoe, 2023](#)), extending the outcome space from innovation and firm performance to inventor geographic mobility—a margin no prior paper has examined. Second, it contributes to research on knowledge worker mobility and the geographic concentration of innovation ([Moretti, 2012](#); [Kerr and Robert-Nicoud, 2020](#)), documenting a new federal-level mechanism that reallocates inventors across states. Third, it speaks to the literature on place-based innovation policy ([Gruber and Johnson, 2019](#)) by identifying a countervailing force—the rejection drain—that erodes the local returns to public investment in innovation ecosystems.

The remainder of the paper proceeds as follows. Section 2 describes the USPTO examination process and the institutional basis for quasi-random examiner assignment. Section 3 presents the data. Section 4 details the identification strategy and discusses threats to validity. Section 5 reports results, and Section 6 discusses implications.

2. Institutional Background

The United States Patent and Trademark Office processes approximately 600,000 patent applications annually, assigning each to one of roughly 15,500 examiners organized into about 600 art units. Art units are technology-specific groups: art unit 2123, for example, handles natural language processing, while art unit 3742 covers heat exchange. Within each art unit,

applications are assigned to examiners based on docket availability and workload balancing rather than application characteristics (Lemley and Sampat, 2012; Righi and Simcoe, 2023).

This assignment mechanism creates the quasi-experiment that underlies the examiner leniency instrument. Conditional on art unit and filing year, the specific examiner who reviews an application is effectively random. Yet examiners differ markedly in their grant propensities. Some of this variation reflects experience and training, some reflects idiosyncratic standards of review. Lemley and Sampat (2012) document that the identity of the assigned examiner is a stronger predictor of patent outcomes than the characteristics of the invention itself.

An examiner’s decision has immediate and lasting consequences for the inventor. A granted patent confers exclusive rights that can be licensed, used to attract venture capital, or leveraged in negotiations with employers. A rejection—technically an “abandonment” when the applicant does not pursue further prosecution—eliminates these rents. For inventors whose human capital is embedded in a specific regional innovation ecosystem, rejection may tip the balance toward relocation: the collaboration network, the employer, and the local knowledge spillovers that made the location attractive are now less valuable without the IP protection that anchored the inventor’s bargaining position.

Patent prosecution typically takes 2–3 years from filing to final disposition. During this period, inventors may file continuation applications, respond to office actions, or appeal rejections. The “abandoned” (ABN) classification in the PAIR database indicates that the applicant ceased prosecution—the strongest form of rejection for the purposes of this study, as it represents a definitive loss of IP rights rather than an ongoing negotiation.

3. Data

I construct an inventor-level panel from two tables in the USPTO’s Patent Application Information Retrieval (PAIR) database, accessed via Google BigQuery (Graham et al., 2015). The `application_data` table provides information on 9.8 million patent applications including examiner identity, art unit, filing date, and disposition (granted or abandoned). The `all_inventors` table links 12.1 million inventor records to applications, with name and US state of residence.

Sample construction. I restrict the sample to regular (utility) patent applications filed between 2002 and 2014 with a resolved outcome (granted or abandoned), a US-based inventor with a valid two-letter state code, and an identified examiner within an art-unit-by-filing-year cell containing at least three examiners. The 2014 endpoint ensures sufficient time for subsequent filings that reveal mobility. After these restrictions, the analysis sample contains

Table 1: Summary Statistics

	Mean	SD	Min	Max
Moved State (Next App.)	0.122	0.327	0.000	1.000
Application Rejected	0.245	0.430	0.000	1.000
Examiner Leniency (LOO)	0.734	0.239	0.000	1.000
Solo Inventor	0.053	0.225	0.000	1.000
Prior Grants	22.092	58.419	0.000	1410.000
Team Size	5.142	4.013	1.000	75.000

Observations: 4,452,425 inventor-application pairs.
Unique inventors: 574,207. Unique applications: 1,545,985. Filing years: 2002–2014.
Examiner leniency is the leave-one-out grant rate within art-unit \times filing-year.

4,452,425 inventor-application pairs covering 574,207 unique inventors and 1,545,985 unique applications reviewed by 11,494 examiners across 6,616 art-unit-by-year cells.

Key variables. The dependent variable, *Moved*, equals one if the inventor’s next patent filing (within 10 years) lists a different US state. This captures actual changes in the inventor’s residential or employment location as revealed by subsequent patent filings. The endogenous variable, *Rejected*, equals one if the application was abandoned. The instrument, *Examiner Leniency*, is the leave-one-out grant rate of the assigned examiner within the same art-unit-by-filing-year cell: the proportion of other applications assigned to the same examiner in the same cell that were granted.

Inventor identification. Following standard practice in the patent literature, I identify inventors by first and last name (Trajtenberg et al., 2006). This creates two potential measurement issues: name collisions (distinct inventors sharing a name) and name variations (the same inventor appearing under different name forms). Both introduce noise in the mobility measure, likely attenuating estimates toward zero. I discuss robustness to this issue in Section 4.2.

Table 1 reports summary statistics. The baseline interstate mobility rate is 12.2 percent—roughly one in eight inventors files a subsequent patent from a different state. The rejection rate is 24.5 percent. Mean examiner leniency is 0.734, reflecting the overall grant rate in the sample, with substantial cross-examiner variation (SD = 0.239).

4. Empirical Strategy

4.1 Identification

I estimate the causal effect of patent rejection on inventor interstate mobility using two-stage least squares (2SLS). The first stage relates rejection to examiner leniency:

$$\text{Rejected}_i = \gamma_0 + \gamma_1 \text{Leniency}_{-i,a,t} + \mathbf{X}'_i \boldsymbol{\delta} + \alpha_{a,t} + u_i \quad (1)$$

where $\text{Leniency}_{-i,a,t}$ is the leave-one-out examiner grant rate for examiner $e(i)$ in art-unit a and filing year t , $\alpha_{a,t}$ are art-unit-by-year fixed effects, and \mathbf{X}_i includes indicators for solo inventor, team size, prior grants, and small entity status. The second stage estimates:

$$\text{Moved}_i = \beta_0 + \beta_1 \widehat{\text{Rejected}}_i + \mathbf{X}'_i \boldsymbol{\phi} + \alpha_{a,t} + \varepsilon_i \quad (2)$$

where $\widehat{\text{Rejected}}_i$ is predicted from the first stage. The coefficient β_1 identifies the local average treatment effect (LATE) of rejection on mobility for inventors whose rejection status is affected by examiner leniency—the “marginal” applicants.

The identifying assumptions are: (i) *relevance*—examiner leniency predicts rejection (verified by $F = 15,179$); (ii) *independence*—conditional on art-unit-by-year, examiner assignment is as-good-as-random; and (iii) *exclusion*—examiner identity affects mobility only through the grant/reject decision.

Standard errors are clustered at the art-unit-by-filing-year level throughout, reflecting the level at which the instrument varies.

4.2 Threats to Validity

Examiner balance. If assignment is truly random within art-unit-by-year cells, examiner leniency should not predict pre-determined inventor characteristics. [Table 5](#) tests this by regressing solo status, prior grants, team size, and prior interstate mobility on examiner leniency. I find statistically significant associations for all four covariates, suggesting imperfect randomization within cells. However, the IV estimate is virtually unchanged when adding these controls (1.06 pp without controls vs. 1.10 pp with), indicating that any residual sorting operates through channels orthogonal to the mobility outcome.

Placebo test. Column (4) of [Table 5](#) tests whether examiner leniency predicts whether the inventor moved states *before* the current application—a pre-determined outcome that should be unrelated to the current examiner. The coefficient (-0.014 , $p < 0.001$) is larger

in magnitude than the 3.9 percentage-point mobility spread between Q1 and Q5 in [Table 3](#) would imply for a linear specification. This warrants caution: it could reflect within-cell sorting that art-unit-by-year fixed effects cannot absorb, or measurement error in name-based inventor IDs that creates spurious state-change correlations. I interpret the main estimates as suggestive of a causal relationship, supported by the clean monotonic reduced form and theoretically consistent heterogeneity, rather than as definitively identified.

Name-based identification. Inventor name matching introduces both false positives (name collisions) and false negatives (name variants). Both create noise in the mobility measure. If this noise is classical, estimates are attenuated; if it correlates with examiner characteristics, bias direction is ambiguous. The stability of estimates across specifications and the clean monotonic reduced form provide reassurance, but a disambiguated inventor panel would strengthen inference.

5. Results

5.1 Main Results

[Table 2](#) reports OLS and IV estimates of the effect of patent rejection on interstate mobility. Column (1) shows the raw OLS relationship: rejected inventors are 2.4 percentage points more likely to move states. Adding art-unit-by-year fixed effects in Column (2) reduces this to 1.3 pp, indicating that much of the raw correlation reflects differences across technology fields and time periods. Controls in Column (3) have negligible additional impact (1.3 pp).

Columns (4) and (5) present the IV estimates. Rejection increases interstate mobility by 1.1 percentage points without controls and 1.1 pp with controls, representing a 9.0 percent increase relative to the 12.2 percent baseline. The IV estimate is slightly smaller than OLS, consistent with positive selection: inherently mobile inventors may file from multiple locations and face higher rejection rates for reasons correlated with their mobility propensity.

5.2 Reduced Form Evidence

[Table 3](#) displays the reduced form across examiner leniency quintiles. Both the rejection rate and the mobility rate decline monotonically from the strictest to the most lenient quintile. Inventors assigned to the strictest examiners (Q1, mean leniency 0.347) face a 54.0 percent rejection rate and a 13.6 percent mobility rate. Those assigned to the most lenient examiners (Q5, mean leniency 0.990) face only a 6.6 percent rejection rate and a 9.7 percent mobility rate. The 3.9 percentage-point spread in mobility across quintiles is smooth and monotonic, confirming that the instrument operates through the expected channel.

Table 2: Patent Rejection and Inventor Interstate Mobility

	Moved State				
	(1)	OLS (2)	(3)	(4)	IV (5)
Constant	0.1158*** (0.0008)				
Rejected	0.0236*** (0.0011)	0.0132*** (0.0006)	0.0131*** (0.0006)	0.0106*** (0.0024)	0.0110*** (0.0025)
Solo Inventor			0.0352*** (0.0013)		0.0352*** (0.0013)
Prior Grants			0.0003*** (1.65×10^{-5})		0.0003*** (1.65×10^{-5})
Team Size			-0.0021*** (0.0001)		-0.0021*** (0.0001)
First-stage F				15179	15179
Controls			✓		✓
Art-unit \times Year FE		✓	✓	✓	✓
Observations	4,452,425	4,452,425	4,452,425	4,452,425	4,452,425
R ²	0.00097	0.01720	0.02150	0.01719	0.02149

Observations: 4,452,425 inventor-application pairs from USPTO PAIR (2002–2014).

Dependent variable: indicator for inventor filing next patent from a different state.

Columns (1)–(3): OLS. Columns (4)–(5): 2SLS using leave-one-out examiner grant rate within art-unit \times filing-year as instrument.

Standard errors clustered at the art-unit \times filing-year level in parentheses.

Table 3: Reduced Form: Inventor Mobility by Examiner Leniency Quintile

	Q1 (Strict)	Q2	Q3	Q4	Q5 (Lenient)
Mean Leniency	0.347	0.644	0.792	0.900	0.990
Rejection Rate	0.540	0.311	0.193	0.112	0.066
Mobility Rate	0.136	0.133	0.127	0.116	0.097
N	894,086	908,711	871,651	888,186	889,791

Leniency quintiles defined by leave-one-out examiner grant rate within art-unit \times filing-year.
Mobility rate declines monotonically from Q1 (13.6%) to Q5 (9.7%), confirming the first stage.

5.3 Mechanisms: Who Moves After Rejection?

If rejection triggers mobility by destroying location-specific intellectual property rents, the effect should be larger for inventors with weaker non-IP anchors (solo vs. team) and stronger outside options (experienced vs. novice). [Table 4](#) tests both predictions using sample splits.

Solo versus team inventors. Column (1) shows that solo inventors—those with no co-inventors on the application—experience a 2.5 pp increase in mobility upon rejection, nearly three times the 0.9 pp effect for team inventors in Column (2). Solo inventors lack the firm-based anchoring that team membership provides: they have no employer-specific lab, no collaborative network tied to a physical location, and no institutional inertia. When their IP claim fails, the location-specific rents that remain are minimal.

Experienced versus novice inventors. Column (3) shows that experienced inventors—those with at least one prior granted patent—respond to rejection with a 2.7 pp increase in mobility, compared to only 0.6 pp for novices in Column (4). This pattern is consistent with experienced inventors having better outside options: their track record makes them attractive to employers and ecosystems elsewhere, giving them the ability to act on the diminished location-specific value. Novices, by contrast, may lack the reputation and network to relocate productively.

5.4 Validity Diagnostics

[Table 5](#) reports examiner balance tests and the placebo check. While the balance tests reveal statistically significant relationships between leniency and pre-determined characteristics—a concern for strict exogeneity—the magnitudes are small relative to the standard deviations of the dependent variables. Critically, the main IV estimate is stable whether or not these covariates are included (Columns 4–5 of [Table 2](#)), suggesting the violations do not materially bias the mobility estimate. The placebo test in Column (4) is the more serious concern and

Table 4: Heterogeneity: Patent Rejection and Mobility by Inventor Type

	Moved State			
	Solo (1)	Team (2)	Experienced (3)	Novice (4)
Rejected	0.0249** (0.0098)	0.0086*** (0.0025)	0.0272*** (0.0029)	0.0061 (0.0039)
Art-unit \times Year FE	✓	✓	✓	✓
Observations	238,175	4,214,250	3,692,031	760,394
R ²	0.05405	0.01491	0.02015	0.04217

2SLS estimates using leave-one-out examiner grant rate as instrument.

Solo: single inventor on application. Team: multiple inventors.

Experienced: at least one prior granted patent. Novice: no prior grants.

All specifications include art-unit \times filing-year fixed effects.

Standard errors clustered at the art-unit \times filing-year level.

motivates the cautious interpretation offered in Section 4.2.

6. Discussion

A back-of-envelope calculation illustrates the potential magnitude. If roughly 250,000 US patent applications are rejected annually, and our IV estimate of 1.1 percentage points applies to marginal applications, examiner-driven rejection could induce approximately 2,750 additional interstate inventor relocations per year. This figure is modest relative to total knowledge worker flows but non-trivial relative to the stock of active inventors in smaller states. The estimate should be interpreted cautiously given the identification concerns documented in Section 4.2, particularly the failed placebo test.

If confirmed in future work with disambiguated inventor panels, this “rejection drain” would interact with place-based innovation policy in an underappreciated way. When a state invests in an innovation ecosystem—through R&D tax credits, university research parks, or startup incubators—it implicitly assumes that the inventors it attracts will remain. If federal patent outcomes push some inventors toward exit, the effective retention rate of publicly subsidized knowledge workers may be lower than assumed. The heterogeneity evidence is consistent with this concern: the inventors most responsive to rejection are solo innovators and experienced patent holders—precisely the types that place-based policies aim to attract.

The LATE interpretation requires care. The examiner leniency design identifies effects

Table 5: Examiner Balance and Placebo Tests

	Solo Inventor (1)	Prior Grants (2)	Team Size (3)	Prior Move (4)
Examiner Leniency (LOO)	-0.0094*** (0.0010)	3.459*** (0.5200)	0.3073*** (0.0529)	-0.0142*** (0.0014)
Art-unit \times Year FE	✓	✓	✓	✓
Observations	4,452,425	4,452,425	4,452,425	3,878,218
R ²	0.09057	0.09919	0.14038	0.01563

Columns (1)–(3): reduced-form regressions of pre-determined inventor characteristics on examiner leniency (leave-one-out grant rate). Column (4): placebo test regressing an indicator for whether the inventor changed states *before* the current application on current-examiner leniency. A valid instrument should produce zeros in all columns. See Section 4.2 for discussion. All specifications include art-unit \times filing-year fixed effects. Standard errors clustered at the art-unit \times filing-year level.

for marginal applicants—those whose rejection status depends on the assigned examiner. These are not the strongest or weakest applications, but those in the middle of the quality distribution where examiner discretion matters most. Extrapolation to all rejected inventors would require assumptions about effect homogeneity that the data cannot verify.

Three limitations deserve emphasis. First, the name-based inventor identification introduces measurement error in the mobility outcome. Name collisions (distinct inventors sharing a name) may create spurious interstate “moves,” while name variations may split a single inventor into multiple identities. The failed placebo test may partly reflect these measurement issues rather than true violations of random assignment. Future work should use disambiguated inventor panels (Trajtenberg et al., 2006) to distinguish measurement error from genuine sorting.

Second, the mobility outcome is observed only for inventors who file a subsequent patent. If rejection causes some inventors to exit the patent system entirely (Farre-Mensa et al., 2020), the analysis conditions on a post-treatment variable, introducing potential selection bias. The direction of this bias is ambiguous: if mobile inventors are more likely to continue patenting (because they relocate to new opportunities), the estimate may be upward biased; if exit and mobility are substitutes, it may be attenuated.

Third, the paper does not implement the state-level labor market analysis that would directly quantify the aggregate consequences of rejection-induced mobility. The inventor-level

estimates establish a micro channel, but the macroeconomic significance of the rejection drain remains an open empirical question for future research using linked patent-employment data.

7. Conclusion

Patent rejection may do more than deny an inventor legal protection over an idea—it may also weaken the ties that bind them to a place. Using examiner assignment at the USPTO, I find suggestive evidence that rejection increases interstate mobility by 1.1 percentage points, with effects concentrated among unanchored and experienced inventors for whom location-specific IP rents matter most. While identification concerns—particularly the failed placebo test—prevent a definitive causal interpretation, the monotonic reduced form and theoretically consistent heterogeneity patterns point toward a real phenomenon. If the rejection drain is confirmed in cleaner data, it represents a hidden channel through which federal patent policy reshapes the geography of innovation—one that state and local policymakers should account for when designing place-based innovation investments.

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: @olafdrw

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

References

- Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist**, “What Is a Patent Worth? Evidence from the U.S. Patent “Lottery”,” *Journal of Finance*, 2020, 75 (2), 639–682.
- Galasso, Alberto and Mark Schankerman**, “Patents and Cumulative Innovation: Causal Evidence from the Courts,” *Quarterly Journal of Economics*, 2015, 130 (1), 317–369.
- Graham, Stuart J.H., Alan C. Marco, and Amanda F. Myers**, “Patent Examination and the Availability of Examiner-Level Data,” *USPTO Economic Working Paper*, 2015, (2015-4).
- Gruber, Jonathan and Simon Johnson**, “Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream,” *New York: PublicAffairs*, 2019.
- Kerr, William R. and Frédéric Robert-Nicoud**, “Tech Clusters,” *Journal of Economic Perspectives*, 2020, 34 (3), 50–76.
- Lemley, Mark A. and Bhaven Sampat**, “Examiner Characteristics and Patent Office Outcomes,” *Review of Economics and Statistics*, 2012, 94 (3), 817–827.
- Moretti, Enrico**, “The New Geography of Jobs,” *Boston: Houghton Mifflin Harcourt*, 2012.
- **and Daniel J. Wilson**, “The Effect of State Taxes on the Geographical Location of Top Earners: Evidence from Star Scientists,” *American Economic Review*, 2017, 107 (7), 1858–1903.
- Righi, Cesare and Timothy Simcoe**, “Patent Examiner Specialization,” *Research Policy*, 2023, 52 (4), 104725.
- Sampat, Bhaven and Heidi L. Williams**, “How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome,” *American Economic Review*, 2019, 109 (1), 203–236.
- Trajtenberg, Manuel, Gil Shiff, and Ran Melamed**, “The “Names Game”: Harnessing Inventors’ Patent Data for Economic Research,” *Annales d’Économie et de Statistique*, 2006, (79/80), 79–108.

Table 6: Standardized Effect Sizes: Patent Rejection and Inventor Mobility

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Interstate mobility (full sample)	0.0110***	(0.0025)	0.327	0.0337	(0.0076)	Small positive
<i>Panel B: Heterogeneous</i>						
Interstate mobility (solo inventors)	0.0249**	(0.0098)	0.394	0.0632	(0.0249)	Moderate positive
Interstate mobility (team inventors)	0.0086***	(0.0025)	0.322	0.0267	(0.0078)	Small positive

- **Notes:** **Country:** United States. **Research question:** Does patent application rejection at the USPTO cause inventors to relocate across state lines, redistributing knowledge workers geographically? **Policy mechanism:** The USPTO assigns patent applications to examiners who vary in grant propensity within technology-specific art units; stricter examiners reject more applications, destroying location-specific intellectual property rents and triggering job search in other states. **Outcome definition:** Binary indicator equal to one if the inventor’s next patent filing (within 10 years) lists a different US state than the current filing. **Treatment:** Binary — application rejected (abandoned) versus granted (issued). **Data:** USPTO PAIR via Google BigQuery, 2002–2014 filings, inventor-application level, 4,452,425 observations. **Method:** 2SLS using leave-one-out examiner grant rate within art-unit \times filing-year as instrument; clustered standard errors at art-unit \times year level. **Sample:** US utility patent applications with resolved outcomes (granted or abandoned), examiners with ≥ 2 decisions per art-unit-year, art-unit-years with ≥ 3 examiners. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the sample standard deviation of the mobility indicator. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A. Standardized Effect Sizes