The Patent Payroll Illusion: Examiner Leniency, Patent Grants, and Local Employment

APEP Autonomous Research* @olafdrw

March 30, 2026

Abstract

The positive correlation between patenting and local employment is often invoked to justify public investment in innovation. We test whether this relationship is causal using a Bartik shift-share instrument that combines pre-determined county patent-technology shares with quasi-random examiner leniency shocks at the USPTO. In a panel of 876 U.S. counties over 2004–2012, OLS yields a positive and significant elasticity of employment with respect to patent grants ($\hat{\beta} = 0.010$, $p < 0.01$). Instrumenting with examiner leniency, the 2SLS estimate is economically zero and statistically insignificant ($\hat{\beta} = -0.018$, SE $= 0.044$; Anderson-Rubin $p = 0.69$). The sign reversal from OLS to IV — what we call the *patent payroll illusion* — survives LIML estimation, state-by-year fixed effects, sector splits, and a clean pre-trend placebo. We find no detectable employment effect of examiner-induced patent grants in invention-active counties.

JEL Codes: O34, J23, R11

Keywords: patents, local employment, examiner leniency, Bartik instrument, null result

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 52m).

1. Introduction

The United States Patent and Trademark Office operates on an annual budget exceeding \$4 billion, and state governments spend billions more on innovation incentives premised on the assumption that patents create jobs (Lerner, 2009). The positive correlation between patenting activity and local employment is well documented and frequently cited by policymakers (Moretti, 2021). But does granting a patent actually cause job creation in the inventor's county — or does this correlation simply reflect the fact that productive places both patent more and hire more?

This paper exploits the quasi-random assignment of patent applications to examiners at the USPTO to answer this question. Within a given technology area (art unit) and year, applications are assigned to examiners through a process that is effectively random (Lemley and Sampat, 2012; Righi and Simcoe, 2023). Because examiners vary substantially and persistently in their propensity to grant patents — some approve over 90% of applications while others reject nearly half — the identity of the assigned examiner generates exogenous variation in whether an application receives a patent (Sampat and Williams, 2019; Farre-Mensa et al., 2020). We aggregate this examiner-level variation to the county level using a shift-share (Bartik) instrument: pre-determined shares capture each county's fixed technology composition, while leave-one-out examiner leniency shocks provide the exogenous year-to-year variation (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022).

Our main finding is a credible null. OLS estimates suggest that a 10% increase in patent grants raises county employment by 0.1% ($p < 0.01$) — a seemingly modest but statistically significant elasticity. When we instrument patent grants with examiner leniency, however, the 2SLS estimate is -0.018 with a standard error of 0.044 ($t = -0.40$). The Anderson-Rubin confidence set, which is robust to weak instruments, comfortably includes zero ($p = 0.69$). The OLS and IV estimates have opposite signs: the positive OLS coefficient reflects selection — counties with more patents are more productive for reasons unrelated to the patents themselves — while the causal effect of an exogenous patent grant on local employment is indistinguishable from zero.

We label this pattern the *patent payroll illusion*: the widely observed positive correlation between patents and jobs is not causal but rather a selection artifact. The instrument is strong (first-stage $F = 16.1$) and the reduced-form effect of the Bartik instrument on employment is also zero ($\hat{\beta} = -0.026$, SE $= 0.065$), confirming that the null is not an artifact of weak instrumentation.

This null survives LIML estimation (-0.019), state-by-year fixed effects (-0.004), contemporaneous employment (-0.012), IHS transformation (-0.016), and leave-one-out state exercises (range: -0.048 to $+0.009$). A pre-trend placebo returns 0.017 (SE = 0.072), confirming the instrument does not predict pre-existing employment trajectories. Sector splits reveal no offsetting effects: neither exposed sectors (-0.046, SE = 0.092) nor local-service sectors (-0.012, SE = 0.036) show significant responses.

This paper contributes to three literatures. First, prior work using examiner leniency finds that patents matter for the firms that receive them — affecting startup growth (Farre-Mensa et al., 2020), follow-on innovation (Galasso and Schankerman, 2015), and gene-patent research (Sampat and Williams, 2019). We ask whether these firm-level effects translate into *local* employment gains. Second, we contribute to innovation and local labor markets (Moretti, 2021; Bloom et al., 2013; Jaffe et al., 1993). Our shift-share approach, closest to Autor et al. (2013)'s China-shock Bartik design, shows that the marginal patent grant does not cause employment growth in the county of invention. Third, we demonstrate the Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2022) Bartik framework in a setting with textbook-quality exogenous shocks.

2. Institutional Background

The USPTO receives roughly 600,000 utility patent applications per year, each routed to one of approximately 300 technology-specific "art units." Within an art unit, applications are assigned to individual examiners quasi-randomly based on availability and workload (Lemley and Sampat, 2012; Righi and Simcoe, 2023). Examiners vary substantially in their grant propensities: some approve over 90% of applications while others grant fewer than 50%, and these differences persist over time (Cockburn et al., 2003; Lemley and Sampat, 2012). Sampat and Williams (2019) formalized this into an IV strategy: because assignment is quasi-random within art-unit-year cells, examiner leniency instruments for whether an application is granted. We extend this to a county-level Bartik shift-share instrument.

The channel from patent grants to local employment could operate through firm-level hiring enabled by commercialization (Farre-Mensa et al., 2020) or through local multiplier effects (Moretti, 2010). We test the aggregate prediction without specifying the exact mechanism.

3. Data

We construct a county-by-year panel combining patent data from the USPTO with employment data from the Census Bureau. The sample covers 876 U.S. counties observed annually from 2004 through 2012.

3.1 Patent Data

We obtain patent examination records from two sources. The USPTO Patent Examination Research (PatEx) dataset, accessed via Google BigQuery, provides the universe of utility patent applications filed since 2001, including examiner identifiers, art unit assignments, filing dates, and final dispositions (granted, abandoned, or pending). We link applications to counties using the PatentsView geocoded inventor data, which assigns each inventor to a county based on the address recorded on the patent or application.

For each application, we identify the first-listed inventor and assign the application to that inventor's county. We restrict to applications with at least one U.S.-based inventor. The share-construction window uses applications filed during 2001–2003 to compute each county's technology composition, measured as its share of applications within each art unit. We require counties to have at least 20 patent applications during this window, yielding 876 counties that represent the bulk of U.S. patenting activity.

For the treatment variable, we count the number of applications *granted* in each county-year during 2004–2012, where the grant year is the calendar year of patent issuance. The endogenous regressor in the second stage is log county patent grants.

An important limitation is that PatentsView provides geocoded inventor locations only for *granted* patents, not for the full universe of applications. The 2001–2003 technology composition shares s_{ca} are therefore computed from granted-patent locations rather than from the application universe. These shares are predetermined in the sense that they precede the estimating sample by at least one year, but they are not exogenous: counties that received more grants during 2001–2003 (partly due to examiner leniency in that period) contribute more to the shares. Under the Borusyak et al. (2022) shock-exogeneity framework, the identifying variation comes from the randomness of post-2003 examiner leniency draws, not from the shares themselves. We rely on this argument plus the clean pre-trend placebo, leave-one-out stability, and robustness to state-by-year fixed effects as evidence that grant-based shares do not contaminate the instrument. Nevertheless, this limitation means our design is somewhat less clean than one using application-level geography would be, and readers should weight the results accordingly.

3.2 Examiner Leniency and the Bartik Instrument

The instrument is constructed in two steps. First, for each examiner j in art unit a and year t, we compute a leave-one-out grant rate: the share of applications granted by examiner j in art-unit-year (a, t), excluding the focal application. This leave-one-out construction avoids mechanical correlation between the instrument and the treatment. We then average across all applications in each art-unit-year cell to obtain the leniency shock z_{at}.

Second, we interact these shocks with pre-determined county technology shares:

$$\text{Bartik}_{ct} = \sum_a s_{ca} \times z_{at} \tag{1}$$

where s_{ca} is the fraction of county c's total patent applications filed in art unit a during 2001–2003 (a within-county technology composition share), and z_{at} is the leave-one-out examiner leniency shock in art unit a in year t. The shares are fixed at their 2001–2003 values and do not vary with contemporaneous economic conditions.

3.3 Employment Data

Employment outcomes come from the Census Bureau's Quarterly Workforce Indicators (QWI), which provide county-level employment, hires, and earnings derived from state unemployment insurance records covering approximately 98% of private-sector employment. We aggregate quarterly data to annual frequency: employment is the mean of four quarters, and new hires are summed across quarters. Monthly earnings are averaged across quarters.

For the mechanism analysis, we split employment into two sector groups. "Exposed" sectors — manufacturing (NAICS 31–33), information (51), and professional/scientific/technical services (54) — are directly connected to patenting activity. "Local-service" sectors — retail (44–45), accommodation and food services (72), and other services (81) — serve local demand and should respond only through indirect spillovers.

3.4 Summary Statistics

Table 1 presents summary statistics for the estimation sample. The average county-year has 119 patent grants, with substantial variation (SD = 457), reflecting the skewed distribution of inventive activity across U.S. counties. Mean county employment is 108,000, again with large cross-county variation.

4. Empirical Strategy

4.1 Estimating Equations

Our primary specification is a two-stage least squares (2SLS) regression at the county-year level:

First stage:

$$\log(\text{Grants}_{ct}) = \alpha_c + \gamma_t + \pi \cdot \text{Bartik}_{ct} + u_{ct} \tag{2}$$

Table 1: Summary Statistics

	Mean	SD	Min	Max	N
Panel A: Patent Variables					
Patent grants (county-year)	118.5	456.8	0	11,865	7,824
Bartik instrument	0.7	0.1	0	1	7,824
Panel B: Employment Outcomes ($t+1$)					
Employment (thousands)	107.7	228.7	2	3,854	7,824
New hires (thousands, annual)	69.5	146.3	0	2,845	7,824
Monthly earnings ($)	3,338.3	716.1	1,924	9,141	7,824
Panel C: Sector Employment ($t+1$, thousands)					
Exposed sectors (Mfg + Info + Prof/Sci)	22.8	54.3	0	1,054	7,824
Local services (Retail + Accom + Other)	29.4	59.4	0	1,059	7,824

Notes: Sample consists of 876 US counties with at least 20 patent applications during the 2001–2003 share-construction window, observed annually over 2004–2012. Patent grants are counted by first-inventor county using PatentsView geocoded locations linked to USPTO Patent Examination Research data. Employment outcomes are from the Census Bureau Quarterly Workforce Indicators (QWI), aggregated from quarterly to annual frequency (employment: mean of four quarters; hires: sum). The Bartik instrument is the county-level shift-share instrument constructed from pre-determined (2001–2003) county art-unit patent application shares interacted with leave-one-out examiner leniency shocks.

Second stage:

$$\log(\text{Emp}_{c,t+1}) = \alpha_c + \gamma_t + \beta \cdot \log(\text{Grants}_{ct}) + \varepsilon_{ct} \tag{3}$$

where α_c and γ_t are county and year fixed effects. The outcome is measured at $t+1$ to allow time for patents to affect commercialization and hiring. Standard errors are clustered at the county level to account for serial correlation within counties (Bertrand et al., 2004).

The OLS estimate of β in Equation (3) is biased if unobserved county-year shocks (e.g., local demand booms, arrival of a major employer) jointly drive patenting and employment. The Bartik instrument isolates variation in patent grants attributable to the quasi-random examiner assignment process, which is orthogonal to local economic conditions by construction.

4.2 Identification and Threats

The validity of the Bartik instrument rests on the exogeneity of the shocks (Borusyak et al., 2022). Three conditions must hold. First, *shock exogeneity*: within an art-unit-year, examiner assignment is quasi-random (Lemley and Sampat, 2012), so average leniency is uncorrelated with any county's application quality. Second, *no dominant shares*: with 876 counties and roughly 300 art units, no single county dominates, verified through leave-one-out state exercises. Third, the *exclusion restriction* requires the instrument to affect employment only through patent grants; we test this with a pre-trend placebo at $t-1$.

The main threat is that pre-determined technology shares correlate with county employment trends. We address this with: (i) the pre-trend placebo; (ii) state-by-year fixed effects absorbing state-level confounders; and (iii) leave-one-out state exercises. Weak instruments are a secondary concern: the first-stage $F = 16.1$ exceeds the Stock and Yogo (2005) threshold of 10, and we report LIML estimates and Anderson-Rubin confidence sets as additional safeguards.

5. Results

5.1 First Stage

Column (1) of Table 2 reports the first-stage regression of log patent grants on the Bartik instrument. The coefficient is 1.504 (SE = 0.375), and the F-statistic is 16.1, well above conventional thresholds. A one-unit increase in the Bartik instrument — corresponding to a shift from uniformly average to uniformly lenient examiners across a county's technology portfolio — increases patent grants by approximately 150 log points. The instrument has substantial predictive power.

5.2 The Patent Payroll Illusion: OLS versus IV

Table 2 presents the core result. OLS yields a positive elasticity of 0.010 (SE $= 0.002$, $p < 0.01$), consistent with the conventional wisdom. The 2SLS estimate reverses sign: $\hat{\beta} = -0.018$ (SE $= 0.044$, $t = -0.40$), economically negligible and statistically indistinguishable from zero. The reduced form confirms: the Bartik instrument's direct effect on employment is -0.026 (SE $= 0.065$), also zero.

This sign reversal — positive OLS, zero IV — is the patent payroll illusion. OLS is upward-biased because productive counties simultaneously patent more and hire more. The IV estimate removes this selection and reveals that the marginal examiner-induced patent grant has no detectable effect on local employment.

Table 2: Patent Grants and Local Employment: OLS and IV Estimates

	First Stage $\log(\text{Grants}_{ct})$ (1)	OLS $\log(\text{Emp}_{c,t+1})$ (2)	2SLS $\log(\text{Emp}_{c,t+1})$ (3)	Reduced Form $\log(\text{Emp}_{c,t+1})$ (4)
Bartik instrument	1.5041***			-0.0264
	(0.3747)			(0.0654)
$\log(\text{Grants}_{ct})$		0.0099***	-0.0175	
		(0.0022)	(0.0437)	
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
First-stage F	16.1		16.1	
Observations	7,824	7,824	7,824	7,824
Counties	876	876	876	876

Notes: The dependent variable in columns (2)–(4) is log county employment in year $t + 1$, measured as the mean of four quarterly QWI observations. Column (1) reports the first stage: the Bartik instrument predicts log patent grants. The Bartik instrument interacts pre-determined (2001–2003) county patent-technology shares with leave-one-out examiner leniency shocks within art-unit-year cells. Column (2) is OLS. Column (3) is 2SLS using the Bartik instrument. Column (4) is the reduced-form effect of the instrument on employment. Standard errors clustered at the county level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.3 Robustness

Table 3 presents five specifications probing the null's stability. LIML, which is median-unbiased under weak instruments, yields -0.019 (SE $= 0.022$; AR $p = 0.687$). State-by-year fixed effects, which absorb all state-level time-varying confounders and exploit only within-state variation, produce -0.004 (SE $= 0.048$). Contemporaneous employment gives -0.012 (SE $= 0.047$); the IHS transformation gives -0.016 (SE $= 0.040$). Leave-one-out state exercises (not tabulated) show the point estimate ranges from -0.048 to $+0.009$, never approaching significance.

Table 3: Robustness of the Null Employment Effect

	Baseline (1)	LIML (2)	State×Year FE (3)	Contemp. (t) (4)	IHS(Grants) (5)
log(Grants)	-0.0175	-0.0191	-0.0039	-0.0115	
	(0.0437)	(0.0220)	(0.0481)	(0.0472)	
IHS(Grants)					-0.0161
					(0.0403)
AR p-value ($\beta = 0$)	0.687				
County FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes		Yes	Yes
State×Year FE			Yes		
First-stage F	16.1		35.6	16.1	18.3
Observations	7,824	7,824	7,824	7,811	7,824

Notes: All specifications instrument log patent grants with the Bartik examiner-leniency instrument. Column (1) reproduces the baseline from Table 2. Column (2) uses LIML estimation with heteroskedasticity-robust standard errors. Column (3) replaces year fixed effects with state×year fixed effects, exploiting only within-state variation across counties. Column (4) uses contemporaneous employment (t) rather than $t + 1$. Column (5) uses the inverse hyperbolic sine transformation of grants. The Anderson-Rubin p-value tests the null $\beta = 0$ and is robust to weak instruments. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.4 Mechanisms and Sector Heterogeneity

Table 4 tests whether the aggregate null masks offsetting sector-level effects. The exposed-sector estimate is -0.046 (SE $= 0.092$) and the local-service estimate is -0.012 (SE $= 0.036$); neither is significant. The pre-trend placebo (column 4) returns 0.017 (SE $= 0.072$), validating the identifying assumption. Additional margins show a similar pattern: new hires yield -0.106 (SE $= 0.086$), and monthly earnings yield -0.073 (SE $= 0.030$), a marginally significant negative effect that we do not interpret strongly given multiple testing across outcomes.

Table 4: Mechanism: Sector Heterogeneity and Pre-Trend Placebo

	All Sectors $\log(\text{Emp}_{c,t+1})$ (1)	Exposed $\log(\text{Emp}_{c,t+1})$ (2)	Local Service $\log(\text{Emp}_{c,t+1})$ (3)	Placebo ($t-1$) $\log(\text{Emp}_{c,t-1})$ (4)
$\log(\text{Grants}_{ct})$	-0.0175	-0.0459	-0.0123	
	(0.0437)	(0.0920)	(0.0355)	
Bartik instrument				0.0169
				(0.0723)
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	7,824	7,824	7,824	6,935

Notes: Columns (1)–(3) report 2SLS estimates of the effect of log patent grants on log employment in year $t+1$, separately for all sectors, exposed sectors (manufacturing, information, professional/scientific/technical services), and local-service sectors (retail, accommodation/food, other services). Column (4) reports the reduced-form effect of the Bartik instrument on log employment at $t-1$ (pre-treatment placebo). A zero in column (4) supports the identifying assumption that the instrument does not predict pre-existing employment trends. Standard errors clustered at the county level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.5 Interpretation

Three observations sharpen the interpretation. First, the null persists at longer horizons: the 2SLS estimate at $t+2$ is -0.015 (SE $= 0.043$, $N = 7{,}041$), ruling out the possibility that commercialization lags mask a positive effect that simply takes more than one year to materialize.

Second, our design identifies the effect of the *marginal* patent — the application that an unusually lenient examiner grants but an average examiner would reject. These marginal patents may be lower-quality or closer to the patentability threshold, and the null does not extend to infra-marginal breakthrough inventions. Prior work finds that firm-level examiner leniency affects startup growth (Farre-Mensa et al., 2020) and follow-on innovation (Galasso and Schankerman, 2015); our contribution is showing that these firm-level effects do not aggregate into detectable local employment gains.

Third, the design is sufficiently powered to detect economically meaningful effects. With SE $= 0.044$, the minimum detectable elasticity at 80% power is approximately 0.12. This means a 10% increase in examiner-induced patent grants would need to raise county employment by about 1.2% for us to detect it. The null therefore rules out large employment multipliers from patent granting decisions, while remaining agnostic about effects below this threshold.

6. Conclusion

Policymakers routinely invoke the positive correlation between patenting activity and local prosperity to justify public investment in innovation infrastructure. We show that this correlation does not survive causal scrutiny. Using quasi-random examiner assignment at the USPTO to construct a Bartik shift-share instrument, we find no detectable employment effect of examiner-induced patent grants in 876 invention-active U.S. counties. The OLS estimate is positive and significant; the IV estimate is zero. We call this sign reversal the patent payroll illusion.

This null result does not imply that patents are useless. Patents clearly matter for the firms that receive them (Farre-Mensa et al., 2020), and the innovation they protect may generate substantial welfare gains through consumer surplus and knowledge spillovers (Galasso and Schankerman, 2015). The lesson is narrower but sharper: place-level prosperity is not evidence that marginal patent-office grant decisions create local employment. The correlation between patents and jobs reflects the sorting of innovative activity into productive places, not a causal effect of granting decisions on the local economy.

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

Autor, David H, David Dorn, and Gordon H Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Bloom, Nick, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Cockburn, Iain M, Samuel Kortum, and Scott Stern, "Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes," *Patents in the Knowledge-Based Economy*, 2003, pp. 19–53.

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.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson, "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics*, 1993, *108* (3), 577–598.

Lemley, Mark A and Bhaven Sampat, "Examiner Characteristics and Patent Office Outcomes," *Review of Economics and Statistics*, 2012, *94* (3), 817–827.

Lerner, Josh, *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed—and What to Do About It*, Princeton University Press, 2009.

Moretti, Enrico, "Local Multipliers," *American Economic Review*, 2010, *100* (2), 373–377.

_ , "The Effect of High-Tech Clusters on the Productivity of Top Inventors," *American Economic Review*, 2021, *111* (10), 3328–3375.

Righi, Cesare and Timothy Simcoe, "Patent Examiner Specialization," *Research Policy*, 2023, *52* (4), 104714.

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.

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.

A. Standardized Effect Sizes

Table 5: Standardized Distributional Effects

	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Employment ($t+1$)	-0.0175	0.0437	1.207	-0.0145	0.0362	Small
New hires ($t+1$)	-0.1061	0.0857	1.241	-0.0855	0.0691	Moderate
Monthly earnings ($t+1$)	-0.0735	0.0302	0.192	-0.3823	0.1569	Large
Panel B: Heterogeneous (Sector Splits)						
Exposed sectors ($t+1$)	-0.0459	0.0920	1.210	-0.0379	0.0760	Small
Local-service sectors ($t+1$)	-0.0123	0.0355	1.202	-0.0102	0.0295	Small

Country: United States. **Research question:** Do examiner-induced patent grants cause local employment growth? **Policy mechanism:** USPTO patent examination quasi-random examiner assignment. **Outcome definition:** Log county-level QWI employment, hires, and earnings at $t+1$. **Treatment:** Log patent grants instrumented by Bartik examiner-leniency shock. **Data:** USPTO PatEx (BigQuery) linked to PatentsView for county geography; Census QWI (LEHD). **Method:** 2SLS with Bartik shift-share IV (pre-determined county art-unit shares $\times$ LOO examiner leniency shocks). **Sample:** 876 US counties with $\geq$20 patent applications in 2001–2003, observed 2004–2012. Classification refers to magnitude, not statistical significance. SDE $= \hat{\beta}/\text{SD}(Y)$. Large: $|\text{SDE}| > 0.15$; Moderate: 0.05–0.15; Small: 0.005–0.05; Null: < 0.005.