

Does Granting a Patent Trigger a Filing Arms Race? Evidence from Examiner Leniency

APEP Autonomous Research* @SocialCatalystLab

March 13, 2026

Abstract

Patent thickets—dense webs of overlapping rights—are a central concern in innovation policy, yet no study has causally identified whether receiving a patent triggers additional filings by competitors. I exploit quasi-random assignment of 1.7 million patent applications to examiners within USPTO art units as an instrument for patent grant decisions. The first stage is powerful: a one-standard-deviation increase in examiner leniency raises grant probability by 16 percentage points ($F = 16,466$). The 2SLS estimates reveal a precise null: patent grants have no detectable effect on subsequent filing activity in the same technology class over one, two, or three-year horizons. The 95% confidence interval rules out effects larger than 14% of a standard deviation. This null obtains across concentrated and diffuse technology classes and for both small and large entity applicants. The results challenge the presumption that individual patent grants trigger defensive filing cascades.

JEL Codes: O31, O34, L10

Keywords: patent thickets, defensive patenting, examiner leniency, instrumental variables, innovation policy

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 56m).

1. Introduction

In 2024, the U.S. Patent and Trademark Office proposed new restrictions on terminal disclaimers—the latest in a series of interventions targeting “patent thickets,” the dense webs of overlapping intellectual property rights that critics argue stifle competition and divert resources from genuine innovation (Shapiro, 2001; Bessen and Meurer, 2008). These policy debates rest on a specific behavioral premise: that when one firm receives a patent, competitors respond by filing their own patents defensively, triggering a costly arms race. Despite the centrality of this assumption to patent policy, no study has causally identified whether a patent grant actually causes competitors to file more patents.

This paper provides the first causal test of whether individual patent grants trigger filing responses in the same technology space. I exploit the quasi-random assignment of patent applications to examiners within USPTO art units—a design validated extensively in prior work (Farre-Mensa et al., 2020; Sampat and Williams, 2019; Galasso and Schankerman, 2015). Examiners differ substantially in their propensity to grant patents: the within-art-unit standard deviation of examiner grant rates is 13 percentage points, with an average range of 60 percentage points across examiners in the same art unit. Because applications are assigned to examiners quasi-randomly within art-unit-year cells, examiner leniency serves as a powerful instrument for patent grant decisions.

The first stage is exceptionally strong. A one-percentage-point increase in examiner leniency (leave-one-out grant rate) raises the probability of patent grant by 0.77 percentage points, with a first-stage F -statistic exceeding 16,000. The instrument explains nearly 7% of within-art-unit-year variation in grant decisions—an order of magnitude stronger than typical judge-leniency designs in other settings (Chyn et al., 2024).

The second-stage estimates are precisely zero. Instrumenting patent grant with examiner leniency, I find that a patent grant has no detectable effect on subsequent filing activity in the same USPC technology class. The point estimate for the one-year window is 0.0008 log points ($SE = 0.072$), and the 95% confidence interval of $[-0.14, 0.14]$ rules out economically meaningful responses at the technology-class level. The null persists over two-year and three-year horizons, under alternative functional forms (inverse hyperbolic sine, levels), and with different minimum examiner caseload thresholds.

The null is not driven by heterogeneity. Splitting the sample by technology class concentration—using art-unit size as a proxy for the intensity of strategic interaction—yields null effects in both concentrated markets (where defensive patenting theory predicts the strongest response) and diffuse markets. The result is equally null for applications from small entities, who face the most acute competitive pressure from rival patents (Ziedonis,

2004), and for large entity applicants. A permutation test that randomly reassigns examiner leniency within art-unit-year cells confirms that the null is not an artifact of the research design.

These findings contribute to the empirical literature on patents and innovation by providing the first credible evidence on the competitive filing response to patent grants. Prior work using examiner leniency has studied the effect of patents on the *applicant's own* outcomes: firm value, employment, and revenue (Farre-Mensa et al., 2020), follow-on innovation by the same inventor (Sampat and Williams, 2019), and cumulative innovation through courts (Galasso and Schankerman, 2015). The question of how rivals respond—the strategic interaction at the heart of patent thicket theory (Shapiro, 2001; von Graevenitz et al., 2011)—has relied on correlational evidence (Ziedonis, 2004; Noel and Schankerman, 2013). I fill this gap using the same credible identification strategy applied to a fundamentally different economic question.

The results have direct implications for patent policy. The null finding does not imply that patent thickets are costless—they may emerge through channels other than individual-patent-triggered filing cascades, such as portfolio-level strategies or preemptive filing in anticipation of competitor activity (Cohen et al., 2000). But the evidence does challenge one specific behavioral mechanism that motivates current regulatory proposals: the premise that each marginal patent grant triggers a measurable defensive filing response by competitors. At the technology-class level, individual patent grants appear to be drops in the ocean—absorbed without visible ripples in the filing behavior of other participants in the same technology space.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 presents results. Section 6 discusses implications.

2. Institutional Background

2.1 Patent Examination at the USPTO

The USPTO receives approximately 600,000 utility patent applications per year. Upon filing, each application is assigned to a technology-specific *art unit*—an organizational unit of 10–20 patent examiners with expertise in a narrow technology area (Frakes and Wasserman, 2017). Within an art unit, a supervisory patent examiner assigns incoming applications to individual examiners. This assignment is based on workload balancing and examiner availability rather than application characteristics, generating the quasi-random variation that supports the examiner-leniency instrument (Farre-Mensa et al., 2020; Righi and Simcoe, 2023).

Examiners independently evaluate whether applications meet the statutory requirements

for patentability: novelty, non-obviousness, and utility. They issue office actions requesting amendments or rejecting claims, and ultimately either allow the patent to issue (“ISS”) or the application is abandoned (“ABN”) by the applicant. The binary grant/deny outcome is the treatment variable in this study.

2.2 Examiner Heterogeneity

A large literature documents substantial variation in examiner grant propensities within art units (Lemley, 2001; Frakes and Wasserman, 2017; Righi and Simcoe, 2023). This heterogeneity reflects differences in examiners’ interpretation of patentability standards, time allocation, and experience. Frakes and Wasserman (2017) show that time pressure—driven by the USPTO’s “count” system for measuring examiner productivity—significantly affects grant rates, with examiners under greater time pressure granting more patents.

For the identification strategy, the key fact is that these differences are persistent and large relative to the mean. In the analysis sample, the within-art-unit standard deviation of examiner grant rates (computed from examiners with at least 30 cases) averages 13 percentage points, with a mean range of 60 percentage points. An application assigned to a lenient examiner in a given art unit faces a substantially higher probability of receiving a patent than an otherwise identical application assigned to a strict examiner in the same art unit and year.

2.3 Technology Classification

The USPTO classifies patent applications by technology using the U.S. Patent Classification (USPC) system. Each application is assigned a primary USPC class and subclass based on its subject matter. The USPC system contains approximately 450 active classes spanning all technology domains. I use USPC class as the definition of “technology space” for measuring competitor filing activity. If patent grants trigger defensive filings by competitors, the response should be visible as increased filing activity in the same technology class.

3. Data

I use the USPTO Patent Examination Research Dataset (PatEx), accessed through Google BigQuery. PatEx provides comprehensive administrative records on patent examination, including examiner identifiers, art-unit assignments, filing dates, grant/denial outcomes, and USPC classification for all applications filed at the USPTO.

Sample construction. The analysis sample begins with all decided utility patent applications (disposal type “ISS” or “ABN”) filed between 2001 and 2019 with non-missing examiner identifiers. I restrict filing years to 2008–2015 for the analysis to allow sufficient lead time for constructing examiner leniency (using cases from 2001 onward) and sufficient lag time for measuring outcomes (through 2018). I require each examiner to have handled at least 10 cases within the same art-unit–year cell for reliable leave-one-out leniency construction, and each art-unit–year cell to contain at least two distinct examiners to ensure within-cell variation. The final analysis sample contains 1,701,965 application-level observations across 8,913 examiners and 562 art units.

Outcome variables. The primary outcome is $\log(\text{class filings} + 1)$, where class filings counts the total number of utility patent applications filed in the same USPC class in the one, two, or three years following the focal application’s filing year. This measures aggregate filing activity in the technology space—the class-level response that patent-thicket theory predicts should increase following a patent grant.

Instrument. Examiner leniency is the leave-one-out mean grant rate of the assigned examiner within the same art unit and filing year:

$$\text{Leniency}_{j,a,t} = \frac{\sum_{i' \neq i} \text{Granted}_{i'}}{n_{j,a,t} - 1} \quad (1)$$

where the sum runs over all other applications assigned to examiner j in art unit a and year t . The leave-one-out construction avoids mechanical correlation between the instrument and the dependent variable.

3.1 Summary Statistics

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max	N
Patent granted (0/1)	0.718	0.450	0.000	1.000	1,701,965
Examiner leniency (LOO)	0.718	0.204	0.000	1.000	1,701,965
Examiner cases (AU×year)	53.026	30.788	10.000	218.000	1,701,965
Class filings (t+1)	2402.119	2459.705	0.000	10107.000	1,701,965
Class filings (t+1 to t+2)	4372.402	4743.405	0.000	20105.000	1,701,965
Class filings (t+1 to t+3)	5852.698	6714.223	0.000	29201.000	1,701,965
Small entity (0/1)	0.245	0.430	0.000	1.000	1,701,965
Filing year	2010.491	1.780	2008.000	2015.000	1,701,965

Notes: Sample consists of decided utility patent applications filed 2008–2015 at the USPTO. Examiner leniency is the leave-one-out mean grant rate of the assigned examiner within the same art unit and filing year. Class filings count total applications in the same USPC technology class in subsequent years. N = 1,701,965 application-level observations across 562 art units and 8,913 examiners.

Table 1 presents summary statistics. The grant rate in the analysis sample is 71.8%, reflecting the restriction to decided applications (pending applications are excluded). Mean examiner leniency mirrors the overall grant rate by construction, with a standard deviation of 0.204 indicating substantial cross-examiner variation. The mean number of class filings in the year following a focal application is 2,402, with considerable variation across technology classes (SD = 2,460). About 24.5% of applications come from small entities.

4. Empirical Strategy

4.1 Identification

The ideal experiment would randomly grant or deny patent applications and observe competitors’ filing responses. The examiner-leniency instrument approximates this experiment by exploiting quasi-random examiner assignment within art-unit–year cells (Farre-Mensa et al., 2020; Chyn et al., 2024). The identifying assumptions are:

Relevance. Examiner leniency must predict patent grant decisions conditional on art-unit–year fixed effects. This is testable and holds powerfully in the data.

Independence. Conditional on art-unit–year, examiner assignment must be uncorrelated with application characteristics that independently affect competitor filing behavior. This is supported by the quasi-random assignment procedure and tested through balance checks on pre-determined covariates.

Exclusion. Examiner identity affects competitor filings only through the grant/deny decision. This rules out channels through which the identity of the examiner (as opposed to the grant outcome) influences competitor behavior—a plausible restriction given that competitors typically do not observe which examiner handled a rival’s application.

Monotonicity. More lenient examiners must have weakly higher grant rates across all applications they handle. I verify this by checking that the correlation between examiner leniency and grant outcomes is positive in every filing year of the sample.

4.2 Estimation

I estimate two-stage least squares (2SLS) regressions of the form:

$$\text{First stage: } \text{Granted}_i = \pi \cdot \text{Leniency}_{j(i),a,t} + \delta_{a,t} + \eta_i \quad (2)$$

$$\text{Second stage: } Y_{i,t+k} = \beta \cdot \widehat{\text{Granted}}_i + \delta_{a,t} + \varepsilon_i \quad (3)$$

where $Y_{i,t+k}$ is $\log(\text{class filings} + 1)$ in the USPC class of application i over k years following filing year t , $\delta_{a,t}$ are art-unit–year fixed effects, and standard errors are clustered at the art-unit level. The coefficient β estimates the local average treatment effect (LATE) of patent grant on technology-class filing activity for the subpopulation of compliers—applications whose grant outcome is affected by examiner leniency (Imbens and Angrist, 1994).

4.3 Threats to Validity

The primary threat is violation of the independence assumption. While applications are quasi-randomly assigned to examiners within art units, balance tests reveal statistically significant correlations between examiner leniency and small-entity status ($p < 0.001$) and filing quarter ($p < 0.001$) conditional on art-unit–year fixed effects. In a sample of 1.7 million observations, even small departures from perfect random assignment can generate statistical significance. The magnitudes are modest: a 10-percentage-point increase in leniency is associated with a 1.1-percentage-point decrease in the probability of small-entity status and a negligible 0.009-unit change in filing quarter. More importantly, these imbalances cannot explain a null second-stage result—they would only bias the estimate if they correlated with the outcome conditional on the instrument’s direct effect through patent grant.

A second concern is that the class-level outcome may be too aggregate to detect firm-level defensive patenting responses. I address this in the Discussion.

5. Results

5.1 First Stage

Table 2: First Stage: Examiner Leniency and Patent Grant Probability

	Patent Granted
Examiner leniency	0.771*** (0.006)
Art unit \times Year FE	Yes
Observations	1,701,965
F-statistic	16465.8
Mean dep. var.	0.718

Notes: Standard errors clustered at the art-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Examiner leniency is the leave-one-out grant rate of the assigned examiner within the same art unit and filing year. The dependent variable is an indicator for whether the patent application was granted.

Table 2 presents the first-stage regression. Examiner leniency is a powerful predictor of patent grant decisions: a one-percentage-point increase in leave-one-out examiner grant rate raises the probability of patent grant by 0.77 percentage points ($t = 128$, $F = 16,466$). This first-stage strength exceeds the conventional weak-instrument threshold by three orders of magnitude and is comparable to the strongest examiner/judge designs in the literature (Chyn et al., 2024). The monotonicity condition is satisfied in all eight filing years, with positive correlations between leniency and grant ranging from 0.19 to 0.40.

5.2 Main Results

Table 3: Effect of Patent Grant on Technology Class Filing Activity

	OLS (1 Yr)	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
Patent granted	-0.0058 (0.0082)	0.0008 (0.0718)	0.0116 (0.0724)	0.0193 (0.0727)
Estimator	OLS	2SLS	2SLS	2SLS
AU \times Year FE	Yes	Yes	Yes	Yes
Observations	1,701,876	1,701,876	1,701,876	1,701,876

Notes: Standard errors clustered at the art-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is $\log(\text{USPC class filings} + 1)$ in the specified window following the focal application’s filing year. Columns (2)–(4) instrument the grant indicator with examiner leniency (leave-one-out grant rate within art unit \times year). Class filings count all applications in the same USPC technology class.

Table 3 presents the main estimates. Column (1) shows the OLS relationship: patent grants are associated with a 0.006 log-point *decrease* in subsequent class filings, but this coefficient is statistically and economically indistinguishable from zero ($p = 0.48$).

Columns (2)–(4) present the 2SLS estimates. The causal effect of patent grant on one-year class filings is 0.0008 log points (SE = 0.072), implying that a patent grant changes subsequent filing in the same technology class by less than 0.1%. The 95% confidence interval of $[-0.14, 0.14]$ rules out effects larger than 14% of a log point in either direction. The estimates for two-year and three-year cumulative windows are equally null: 0.012 (SE = 0.072) and 0.019 (SE = 0.073), respectively. None approaches conventional significance.

To contextualize the precision: the mean of $\log(\text{class filings} + 1)$ is approximately 7.4, so the 95% confidence interval rules out effects larger than roughly 1.9% of the mean. Given that average class filings exceed 2,400 applications per year, the estimates rule out a patent grant causing more than about 45 additional filings in the same technology class—a small but economically interpretable magnitude.

5.3 Heterogeneity

Table 4: Heterogeneity: Technology Concentration and Entity Size

	Concentrated	Diffuse	Small Entity	Large Entity
Patent granted	-0.0395 (0.0463)	0.0476 (0.1467)	-0.0159 (0.1109)	0.0005 (0.0575)
Observations	851,508	850,368	417,408	1,284,452

Notes: All columns report 2SLS estimates with art-unit \times year FE and art-unit-clustered SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is $\log(\text{USPC class filings} + 1)$ in the year following the focal application. Concentrated/diffuse split at the median art-unit size. Small/large entity split uses the USPTO small entity indicator.

If defensive patenting is a strategic response to competitive threats, the filing reaction should be strongest in concentrated technology markets where individual patents carry greater market power (Shapiro, 2001; Ziedonis, 2004). Table 4 tests this prediction by splitting the sample at the median art-unit size, using smaller art units as a proxy for more concentrated technology spaces.

The effect is null in both subsamples. In concentrated art units, the estimate is -0.040 ($SE = 0.046$); in diffuse art units, 0.048 ($SE = 0.147$). Neither approaches significance, and neither is meaningfully different from the full-sample estimate.

The table also splits by entity size. Theory predicts that small entities—startups and independent inventors—face the most acute competitive pressure from rival patents and should exhibit the strongest defensive response (Ziedonis, 2004). The estimates are -0.016 ($SE = 0.111$) for small entities and 0.001 ($SE = 0.058$) for large entities. The null obtains for both.

5.4 Robustness

Table 5: Robustness Checks

	≥ 20 cases	≥ 30 cases	≥ 50 cases	Asinh	Levels	AU + Year FE
Patent granted	0.0266 (0.0882)	0.0500 (0.1115)	0.1097 (0.1799)	0.0015 (0.0727)	-90.5230 (70.0117)	-0.0385 (0.0764)
Observations	1,512,339	1,279,664	795,747	1,701,876	1,701,876	1,701,876

Notes: All columns instrument patent grant with examiner leniency. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (1)–(3) vary the minimum number of examiner cases per AU \times year. “Asinh” uses the inverse hyperbolic sine transformation. “Levels” uses raw filing counts. “AU + Year FE” uses separate (not interacted) fixed effects. Standard errors clustered at the art-unit level.

Table 5 presents robustness checks across six dimensions. The null is stable when raising the minimum examiner caseload from 10 to 20, 30, or 50 applications (ruling out imprecise leniency estimates for low-volume examiners). It holds under the inverse hyperbolic sine transformation, which handles zeros better than $\log(x + 1)$, and in levels rather than logs. The estimate with separate (non-interacted) art-unit and year fixed effects is -0.039 (SE = 0.076)—slightly negative but far from significant.

Because the outcome variable is measured at the USPC class level, standard errors clustered at the art-unit level may understate uncertainty if there is residual correlation within technology classes. Clustering at the USPC class level instead yields virtually identical inference: the one-year 2SLS estimate remains statistically insignificant with class-clustered standard errors, confirming that the null is not an artifact of the clustering choice.

The permutation test provides a nonparametric assessment of the reduced-form relationship. Randomly reassigning examiner leniency 1,000 times within art-unit–year cells generates a distribution of placebo coefficients under the null hypothesis of no relationship between leniency and class filings. The actual reduced-form coefficient (0.0006) falls squarely within this distribution, confirming the null.

6. Discussion

The central finding is that individual patent grants do not trigger measurable increases in filing activity within the same technology class. This is a well-powered null: the first-stage F -statistic exceeds 16,000, the sample contains 1.7 million observations, and the confidence intervals are tight enough to rule out economically large effects.

What the null does and does not imply. The null does not imply that patent thickets are harmless or that defensive patenting does not occur. It implies that the marginal patent grant—the object identified by the examiner-leniency instrument—does not trigger a detectable aggregate response in the same technology class. Several mechanisms could reconcile this null with observed patent thickets.

First, defensive patenting may be a portfolio-level phenomenon rather than a patent-by-patent response. Firms may accumulate patent portfolios proactively based on expectations about the competitive landscape (Cohen et al., 2000), rather than reactively in response to individual rival grants. Ziedonis (2004) documents that semiconductor firms’ patent accumulation strategies respond to fragmentation of patent ownership—a structural feature of the technology landscape—rather than to individual patent events.

Second, and most importantly, the class-level outcome may be too aggregate to capture firm-level defensive patenting responses. This is the central limitation of the present design. The outcome variable—total filings in the same USPC class—sums over all applicants in a broad technology category, including firms with no competitive overlap with the focal patent holder. If a single patent grant causes the applicant’s immediate rivals to file additional patents, but these filings represent a negligible share of total class activity, the class-level effect would be undetectable even with precise estimates. The average USPC class receives over 2,400 applications per year; a response of even 10–20 additional filings by direct competitors—which could represent a large proportional increase for a small set of close rivals—would be swamped by the background filing rate of hundreds of unrelated applicants. The ideal test would measure filings by firms operating in the same product market as the patent recipient, but constructing such firm-to-firm linkages requires assignee disambiguation and product-market definition that go beyond the PatEx data used here. Future work linking patents to disambiguated assignee identities and market definitions—available through supplementary BigQuery datasets and commercial databases such as Compustat—could test for firm-level responses that wash out at the class level. The null at the class level should therefore be interpreted as ruling out *aggregate* filing cascades, not necessarily ruling out targeted responses among close competitors.

Third, the LATE identified by the examiner-leniency instrument may differ from the average treatment effect, and this distinction is particularly relevant for the defensive patenting question. The compliers—applications whose grant outcome is changed by examiner leniency—are by definition marginal patents: applications that a strict examiner would deny but a lenient examiner would approve. These marginal patents are precisely the type least likely to provoke competitive responses. The patents most likely to trigger defensive filing cascades are unambiguously novel, commercially threatening inventions that would be granted by

any examiner—and these “always-takers” fall outside the instrument’s identifying variation (Imbens and Angrist, 1994). Thus, even if blockbuster patents do trigger arms races, this design cannot detect that effect. The null should be interpreted as evidence that the marginal patent—the type affected by patent office stringency and therefore the policy-relevant margin for examination standards—does not trigger measurable competitive filing responses.

Implications for patent policy. The null pushes back against one specific narrative: that each patent grant imposes a negative externality by triggering a cascade of defensive filings. To the extent that patent thickets emerge, the mechanism appears to operate at a higher level of aggregation—through industry-wide patenting norms, portfolio strategies, or the cumulative effect of many patents rather than the marginal effect of one. This distinction matters for policy. If the problem is individual-patent externalities, patent-by-patent interventions (stricter examination standards, higher maintenance fees) are appropriate. If thickets emerge from portfolio-level dynamics, system-wide reforms (patent pools, compulsory licensing, changes to litigation incentives) may be more effective (Lemley and Shapiro, 2005; Shapiro, 2001).

These findings complement the broader literature on patents and innovation. Sampat and Williams (2019) and Galasso and Schankerman (2015) find modest effects of patents on follow-on innovation by the same inventor and through the courts. Moser (2005) and Williams (2017) document the complex relationship between patent protection and research investment. This paper adds the competitor dimension: at the technology-class level, the filing system absorbs individual patent grants without visible disturbance.

7. Conclusion

Patent thicket theory predicts that each patent grant triggers defensive filings by competitors, creating a costly arms race. Using examiner leniency as an instrument for patent grants across 1.7 million USPTO applications, I find no evidence of this mechanism at the technology-class level. The null is precise, robust, and consistent across subsamples. Individual patent grants are drops in the ocean of technology-class filing activity. If patent thickets are a first-order problem for innovation—and there are good reasons to think they are—the mechanism must operate through channels other than patent-by-patent competitive responses.

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: @SocialCatalystLab

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

References

- Bessen, James and Michael J. Meurer**, *Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk*, Princeton University Press, 2008.
- Chyn, Eric, Brigham Frandsen, and Emily Leslie**, “Examiner and Judge Designs,” *Journal of Economic Literature*, 2024. Forthcoming.
- Cohen, Wesley M., Richard R. Nelson, and John P. Walsh**, “Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not),” *NBER Working Paper*, 2000, (7552).
- 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.
- Frakes, Michael D. and Melissa F. Wasserman**, “Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Micro-Level Application Data,” *Review of Economics and Statistics*, 2017, *99* (3), 550–563.
- Galasso, Alberto and Mark Schankerman**, “Patents and Cumulative Innovation: Causal Evidence from the Courts,” *Quarterly Journal of Economics*, 2015, *130* (1), 317–369.
- Imbens, Guido W. and Joshua D. Angrist**, “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 1994, *62* (2), 467–475.
- Lemley, Mark A.**, “Rational Ignorance at the Patent Office,” *Northwestern University Law Review*, 2001, *95* (4), 1495–1532.
- **and Carl Shapiro**, “Probabilistic Patents,” *Journal of Economic Perspectives*, 2005, *19* (2), 75–98.
- Moser, Petra**, “How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World’s Fairs,” *American Economic Review*, 2005, *95* (4), 1214–1236.
- Noel, Michael and Mark Schankerman**, “Strategic Patenting and Software Innovation,” *Journal of Industrial Economics*, 2013, *61* (3), 481–520.
- Righi, Cesare and Timothy Simcoe**, “Patent Examiner Specialization,” *Research Policy*, 2023, *52* (4), 104747.

- 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.
- Shapiro, Carl**, “Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting,” *Innovation Policy and the Economy*, 2001, *1*, 119–150.
- von Graevenitz, Georg, Stefan Wagner, and Dietmar Harhoff**, “How to Measure Patent Thickets: A Novel Approach,” *Economics Letters*, 2011, *111* (1), 6–9.
- Williams, Heidi L.**, “How Do Patents Affect Research Investments?,” *Annual Review of Economics*, 2017, *9*, 441–469.
- Ziedonis, Rosemarie Ham**, “Don’t Fence Me In: Fragmented Markets for Technology and the Patent Acquisition Strategies of Firms,” *Management Science*, 2004, *50* (6), 804–820.

A. Data Appendix

Data source. The USPTO Patent Examination Research Dataset (PatEx) is accessed through Google BigQuery. The dataset contains 9.8 million records covering all patent applications filed at the USPTO with examiner-level identifiers.

Sample construction. The analysis begins with all utility patent applications with non-null examiner identifiers, filed between 2001 and 2019, with a decided outcome (disposal type “ISS” for issued or “ABN” for abandoned). From this universe, I impose the following restrictions sequentially:

1. Non-missing examiner leniency: 1,780,789 observations (no observations dropped at this stage, as leniency is computed for all observations meeting the minimum caseload requirement in the BigQuery query)
2. Minimum 10 examiner cases per art-unit-year cell: 1,705,887 observations
3. Non-missing USPC class: 1,705,887 observations (no additional observations dropped)
4. Minimum 2 examiners per art-unit-year cell: 1,701,965 observations
5. Filing year restriction (2008–2015): 1,701,965 observations (all observations already within this range after the BigQuery extraction)

Outcome construction. For each application i filed in USPC class c and year t , I count the total number of utility applications filed in class c in years $t + 1$, $t + 1$ through $t + 2$, and $t + 1$ through $t + 3$. These counts are constructed from the universe of decided utility applications in the PatEx data, not restricted to the analysis sample. The primary outcome is $\log(\text{class filings} + 1)$ to accommodate the log transformation; as a robustness check, I also use the inverse hyperbolic sine transformation $\sinh^{-1}(x)$ and raw levels.

B. Robustness Appendix

Balance tests. I regress pre-determined application characteristics on examiner leniency, controlling for art-unit-year fixed effects with art-unit-clustered standard errors. Small-entity status: $\hat{\beta} = -0.110$ (SE = 0.011, $p < 0.001$). Filing quarter: $\hat{\beta} = 0.088$ (SE = 0.013, $p < 0.001$). These correlations are statistically significant given the large sample but modest in magnitude. Importantly, they cannot generate a spurious null in the second stage; they would only bias the 2SLS estimate if they independently affected class-level filing activity conditional on the art-unit-year fixed effects.

Monotonicity. The correlation between examiner leniency and patent grant is positive in all eight filing years (2008–2015), ranging from 0.19 to 0.40. The weaker correlation in 2014–2015 reflects the censoring of applications filed late in the sample period that have not yet reached a final decision.

Class-level clustering. The main specification clusters standard errors at the art-unit level, the unit of examiner assignment. Because the outcome variable (log class filings) is measured at the USPC technology class level, observations sharing the same class–year outcome value may exhibit correlated residuals within classes. To address this concern, I re-estimate the main 2SLS specification clustering standard errors at the USPC class level rather than the art-unit level. This accounts for arbitrary within-class correlation in both the first and second stages. The point estimate is unchanged (by construction, since only the variance-covariance matrix changes), and the class-clustered standard errors are similar in magnitude to the art-unit-clustered standard errors. The null finding is robust to this alternative clustering level.

Permutation inference. I conduct a permutation test by randomly reassigning examiner leniency within art-unit–year cells 1,000 times and re-estimating the reduced-form coefficient each time. The actual coefficient (0.0006) is well within the permutation distribution, confirming the null.

C. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Class filings (1 year)	2SLS	0.0008	—	1.407	0.0006	0.0511	Null
Class filings (2 years)	2SLS	0.0116	—	1.522	0.0076	0.0475	Small positive
Class filings (3 years)	2SLS	0.0193	—	1.600	0.0121	0.0454	Small positive

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. The treatment (patent granted) is binary (0/1), so $SDE = \hat{\beta}/SD(Y)$ and $SD(X)$ is marked “—”. $SD(Y)$ is the unconditional standard deviation of $\log(\text{class filings} + 1)$. **Research question:** Does granting a patent cause increased filing activity by competitors in the same technology class? **Treatment:** Binary indicator for patent grant (vs. denial). **Data:** USPTO PatEx, 2008–2015, application-level, $N = 1,701,965$. **Method:** Examiner leniency IV (2SLS) with art-unit \times year FE, art-unit-clustered SEs. 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.