Paper Patents: Do Marginal Patent Grants Create Real Market Value?

APEP Autonomous Research* @olafdrw

April 2, 2026

Abstract

Are patents merely paper rights, or do they create real economic value in secondary markets? I exploit quasi-random assignment of patent applications to examiners of varying leniency at the USPTO, instrumenting grant decisions with leave-one-out examiner grant rates within 9,838 art-unit–year cells. Among 4.4 million resolved applications (2000–2015), a marginally granted patent is 13.4 percentage points more likely to be traded in the assignment market and 7.3 points more likely to be pledged as collateral—effects comparable to the OLS premium, suggesting limited selection bias. The effect is remarkably uniform across entity sizes, implying that even low-quality marginal patents carry real transactional value. These findings inform the patent-troll debate: if marginal patents trade as readily as inframarginal ones, weak patents are not merely "paper"—they are liquid assets.

JEL Codes: O34, O31, L24

Keywords: patents, examiner leniency, instrumental variables, patent assignment, technology markets, collateralization

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 1m).

1. Introduction

Every year the United States Patent and Trademark Office grants roughly 350,000 patents. Each confers a twenty-year right to exclude—but how many of those rights actually function as economic assets? A patent that sits in a drawer generates no surplus beyond its holder's freedom to operate. A patent that is traded, licensed, or pledged as collateral, by contrast, participates in what Arora et al. (2001) call "markets for technology." The distinction matters: if marginal patents—those whose grant or rejection hinges on the luck of examiner assignment—trade at the same rate as inframarginal ones, then even low-quality patent grants create real market value. If they trade less, the market disciplines quality where the Patent Office does not.

This paper provides the first causal estimates of how patent grants affect outcomes in the secondary patent market. I link the universe of USPTO patent examination records to 5.2 million assignment transactions—ownership transfers, security interests, and merger conveyances—recorded in the USPTO assignment database. To identify the causal effect of the grant decision, I exploit the well-established examiner leniency instrument (Sampat and Williams, 2019; Farre-Mensa et al., 2020; Galasso and Schankerman, 2015): within an art-unit–filing-year cell, patent applications are assigned to examiners quasi-randomly, and examiners vary widely in their propensity to grant. I construct a leave-one-out examiner grant rate for 9,718 examiners across 701 art units and instrument the binary grant decision with this measure.

The first stage is powerful: a one-standard-deviation increase in examiner leniency raises the grant probability by 16 percentage points, with a first-stage F-statistic of 14,949. The IV estimates show that a marginally granted patent is 13.4 percentage points more likely to be transferred to a different owner through the assignment market (SE = 0.56 pp) and 7.3 percentage points more likely to be pledged as collateral through a security interest filing (SE = 0.43 pp). These are economically large effects: the baseline market-transfer rate for abandoned applications is 11.6 percent, so the grant-induced increase more than doubles the probability of a market transaction.

A striking feature of the results is that the OLS and IV estimates are nearly identical for market transfers (OLS: 13.7 pp; IV: 13.4 pp). This near-equivalence implies that selection on unobservables—for instance, that inherently more valuable inventions are both more likely to be granted and more likely to be traded—plays a surprisingly small role once art-unit–year fixed effects are absorbed. The market apparently treats all granted patents, including marginally granted ones, as transferable assets at similar rates.

The entity-size heterogeneity reinforces this interpretation. Figueroa and Serrano (2019)

document that small firms are net sellers and large firms net buyers in patent markets, suggesting different roles in the technology marketplace. If the market disciplined quality, one might expect small-entity marginal patents—which are lower quality on average (Frakes and Wasserman, 2017)—to trade at lower rates. Instead, I find that the IV effect is nearly identical for small entities (12.8 pp) and large entities (13.4 pp). Marginal patents are liquid regardless of who holds them. A preliminary reduced-form comparison had suggested that small-entity patents from strict examiners trade *more* than those from lenient examiners, consistent with quality pricing. This pattern does not survive the full IV specification with art-unit–year fixed effects, suggesting that the earlier signal reflected compositional differences across cells rather than a genuine quality-pricing mechanism.

These findings speak directly to the debate over patent trolls and non-practicing entities (NPEs). Cohen et al. (2019) show that NPE litigation reduces targeted firms' R&D spending, and Feng and Jaravel (2020) demonstrate that broader patent claims increase NPE activity. A key question in this literature is whether weak patents—those that arguably should not have been granted—have economic value or are merely threats. My results suggest the former: marginal patents are not idle threats but tradeable assets that participate fully in secondary markets. This has implications for patent reform: if the Patent Office tightened standards at the margin, it would reduce the stock of liquid patent assets, potentially affecting the efficiency of technology markets (Galasso, 2012; Serrano, 2010).

The paper contributes to several literatures. First, a large body of work uses the examiner leniency instrument to study the effects of patent grants on innovation (Sampat and Williams, 2019; Galasso and Schankerman, 2015), firm outcomes (Farre-Mensa et al., 2020), and litigation (Feng and Jaravel, 2020). I am the first to turn this instrument toward patent market outcomes. Second, Serrano (2010) and Figueroa and Serrano (2019) study patent trading flows but cannot identify the causal role of the grant decision because they observe only granted patents. By including abandoned applications in the sample, I estimate the extensive margin of the grant decision on market participation. Third, the literature on patent quality (Frakes and Wasserman, 2017, 2019; Lemley, 2001) has documented that examiner behavior affects grant rates, but the downstream market consequences of this variation have not been quantified.

The paper proceeds as follows. Section 2 describes the institutional setting. Section 3 introduces the data. Section 4 presents the identification strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The USPTO examines patent applications through a hierarchical structure of technology centers, art units, and individual examiners. Each technology center covers a broad technological domain (e.g., biotechnology, semiconductors, software). Within each center, art units of roughly 10–20 examiners specialize in narrower subfields. When an application arrives, it is classified by subject matter and routed to the relevant art unit. Within the art unit, assignment to a specific examiner is quasi-random: the supervisory patent examiner distributes new cases based on docket availability rather than application characteristics (Cockburn et al., 2003; Sampat and Williams, 2019).

Examiners vary enormously in their propensity to grant. Among examiners with at least 50 resolved applications in my sample, the interquartile range of grant rates spans from roughly 60 to 90 percent, and the full range extends from below 30 to above 98 percent. This variation persists within art-unit–year cells, ruling out compositional differences in application quality as an explanation (Frakes and Wasserman, 2017).

Upon receiving a patent, the holder obtains the right to exclude others from making, using, or selling the patented invention for twenty years from the filing date. Crucially, this right is an assignable property interest: patent holders can transfer ownership through the USPTO's assignment recordation system. The assignment system records three main types of transactions: (i) *assignments*—outright transfers of ownership, including sales to operating companies and NPEs; (ii) *security interests*—pledges of patent rights as collateral for financing; and (iii) *mergers*—transfers associated with corporate reorganizations. Abandoned applications can also have assignment records—for instance, if the applicant transferred the application before abandonment, or if a security interest was filed during prosecution. Importantly, assignments for abandoned applications must occur during prosecution (before the disposal decision), whereas grants open the door to post-grant transfers. This asymmetry means that part of the IV effect may reflect the mechanical unlocking of post-grant trading opportunities rather than a pure valuation channel.

3. Data

I combine two administrative datasets from the USPTO, both accessed via Google BigQuery's public patent data repository.

Patent Examination Data. The Patent Examination Research Dataset (PatEx) contains the universe of patent applications with examination metadata (Marco et al., 2015). I restrict the sample to resolved applications—those that were either issued (granted) or abandoned—

filed between 2000 and 2015, with an identifiable examiner and art unit. This yields 4,379,472 applications examined by 9,718 examiners across 701 art units.

Assignment Data. The USPTO Assignment Dataset records all voluntary and involuntary transfers of patent rights filed with the Office. I link assignment records to applications using the application number and classify each conveyance as an assignment (ownership transfer), security interest (collateral pledge), or merger. I construct two binary outcome variables at the application level: *Market Transfer*, which equals one if the application has at least one non-employer assignment recorded, and *Security Interest*, which equals one if any security interest was filed.

Table 1 reports summary statistics. The overall grant rate is 72.5 percent. Among granted patents, 23.8 percent have at least one market transfer and 10.5 percent have a security interest. Among abandoned applications, the corresponding figures are 11.6 and 4.2 percent—assignment events can occur before abandonment. The average examiner leniency (leave-one-out grant rate) is 0.725 with a standard deviation of 0.223, confirming substantial within-cell variation.

Table 1: Summary Statistics

Variable	N	Mean	SD	Min	Max
Patent Granted	4,379,472	0.725	0.446	0.000	1.000
Market Transfer (Any Assignment)	4,379,472	0.205	0.404	0.000	1.000
Security Interest	4,379,472	0.088	0.283	0.000	1.000
Examiner Leniency (LOO Grant Rate)	4,379,472	0.725	0.223	0.000	1.000
Small Entity	4,379,472	0.259	0.438	0.000	1.000
Number of Conveyances	4,379,472	1.476	1.771	0.000	43.000

Notes: Sample includes 4,379,472 resolved USPTO patent applications filed 2000–2016 with identifiable examiner assignment. Patent Granted equals one if the application was issued (disposal_type = ISS). Market Transfer equals one if the patent was assigned to a different entity. Security Interest equals one if the patent was used as collateral. Examiner Leniency is the leave-one-out grant rate of the assigned examiner within the same art-unit × filing-year cell. Small Entity equals one if the applicant filed under the USPTO's reduced-fee small-entity status.

4. Empirical Strategy

4.1 Identification

I exploit the quasi-random assignment of patent applications to examiners within art-unit–filing-year cells. The instrument is the leave-one-out examiner grant rate:

$$Z_i = \frac{1}{n_{e(i),c(i)} - 1} \sum_{j \neq i : e(j)=e(i), c(j)=c(i)} \text{Grant}_j \tag{1}$$

where $e(i)$ denotes the examiner assigned to application i and $c(i)$ denotes the art-unit–filing-year cell. This measure captures examiner-specific propensity to grant, purged of the mechanical correlation with application i's own outcome. I require each examiner to have at least 50 total resolved applications for inclusion.

The 2SLS specification is:

$$\text{First stage:} \quad \text{Grant}_i = \pi Z_i + \delta_{c(i)} + u_i \tag{2}$$

$$\text{Second stage:} \quad Y_i = \beta \widehat{\text{Grant}}_i + \delta_{c(i)} + \varepsilon_i \tag{3}$$

where Y_i is Market Transfer or Security Interest, $\delta_{c(i)}$ are art-unit–filing-year fixed effects, and standard errors are clustered at the art-unit level (701 clusters).

4.2 Identifying Assumptions

Relevance. Examiner leniency must predict grant outcomes conditional on art-unit–year fixed effects. Table 2 confirms this: a one-unit increase in the leave-one-out grant rate raises the grant probability by 0.71 (first-stage $F = 14{,}949$).

Independence. Conditional on the art-unit–year cell, examiner assignment must be independent of application characteristics. Balance tests reveal a statistically significant correlation between leniency and small-entity status (coefficient $= -0.095$, $p < 0.001$), consistent with imperfect randomization within cells. However, adding small-entity status as a control barely changes the IV estimate (from 0.134 to 0.132), suggesting this imbalance does not materially confound the results.

Exclusion. Examiner leniency must affect market outcomes only through the grant decision. Since the applicant drafts the patent claims before examiner assignment, the examiner's identity should not directly affect the invention's commercial potential—only whether it receives formal legal protection. This assumption would be violated if, for example, lenient

examiners produce broader claims that independently increase market value. Feng and Jaravel (2020) study claim breadth but find that broader claims primarily affect litigation, not direct commercialization.

Monotonicity. A more lenient examiner must weakly increase the grant probability for every application. Table 2 shows that grant rates rise monotonically across leniency deciles (from 32.0% in the bottom decile to 96.1% in the top), supporting this assumption.

4.3 Estimand

The IV identifies a local average treatment effect (LATE): the causal effect of patent grant on market outcomes for complier applications—those whose grant status is changed by examiner assignment (Imbens and Angrist, 1994; Angrist et al., 1996). Compliers are "marginal" patents at the boundary of approval and rejection. This is the economically relevant margin for patent reform: these are the applications whose fate depends on the examiner draw.

5. Results

5.1 First Stage and Reduced Form

Table 2 presents the first stage and reduced form. The first-stage coefficient of 0.714 indicates that a one-standard-deviation increase in examiner leniency ($\sigma = 0.223$) raises the grant probability by approximately 16 percentage points. The F-statistic of 14,949 far exceeds conventional weak-instrument thresholds. The reduced-form estimates show that examiner leniency increases both market transfer (coefficient 0.096) and security interest (coefficient 0.052), both significant at the 1% level.

Table 2: First Stage and Reduced-Form Estimates

	(1) Patent Granted	(2) Market Transfer	(3) Security Interest
Examiner Leniency	0.7145***	0.0958***	0.0519***
	(0.0058)	(0.0040)	(0.0031)
Art Unit × Year FE	Yes	Yes	Yes
Observations	4,379,472	4,379,472	4,379,472
F-statistic	14949.1		

Notes: Column (1) reports the first-stage relationship between examiner leniency (leave-one-out grant rate within art-unit × filing-year) and patent grant probability. Columns (2)–(3) report reduced-form effects of examiner leniency on market outcomes. Standard errors clustered at the art-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Main IV Results

Table 3 compares OLS and IV estimates. The IV estimate for Market Transfer is 0.134 (SE = 0.006): a marginally granted patent is 13.4 percentage points more likely to be transferred to a different owner. The OLS estimate is nearly identical at 0.137. For Security Interest, the IV estimate is 0.073 (SE = 0.004), slightly larger than the OLS estimate of 0.062. The OLS–IV concordance for market transfer suggests that, conditional on art-unit–year fixed effects, selection on unobservables contributes little to the observed grant–transfer correlation. The slightly larger IV effect for collateralization may reflect that marginal patents are disproportionately useful as collateral for cash-constrained small entities who could not otherwise pledge assets.

Table 3: The Causal Effect of Patent Grants on Market Outcomes

	Market Transfer		Security Interest	
	(1) OLS	(2) IV	(3) OLS	(4) IV
Patent Granted	0.1371***	0.1340***	0.0621***	0.0727***
	(0.0023)	(0.0056)	(0.0016)	(0.0043)
Art Unit × Year FE	Yes	Yes	Yes	Yes
Observations	4,379,472	4,379,472	4,379,472	4,379,472
First-Stage F		14949.1		14949.1

Notes: This table compares OLS and IV estimates of the effect of patent grant on market outcomes. The instrument is the leave-one-out examiner grant rate within art-unit × filing-year cells. Market Transfer equals one if the patent was assigned to a different entity via the USPTO assignment system. Security Interest equals one if the patent was pledged as collateral. Standard errors clustered at the art-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Heterogeneity by Entity Size

Table 4 splits the sample by entity size. Small entities (individuals, small businesses, nonprofits) account for 26 percent of the sample and are net sellers in patent markets (Figueroa and Serrano, 2019). If the market priced patent quality, one would expect marginal small-entity patents to trade less than marginal large-entity patents, since small-entity applications face different selection pressures. Instead, the IV effects are strikingly similar: 12.8 percentage points for small entities versus 13.4 for large entities. For security interests, the effect is somewhat smaller for small entities (5.4 pp vs. 7.6 pp), potentially reflecting differential access to credit markets rather than patent quality per se.

Table 4: Heterogeneity by Entity Size: Patent Grants and Market Outcomes

	Market Transfer		Security Interest	
	(1) Small	(2) Large	(3) Small	(4) Large
Patent Granted (IV)	0.1275***	0.1342***	0.0539***	0.0764***
	(0.0058)	(0.0068)	(0.0042)	(0.0050)
Art Unit × Year FE	Yes	Yes	Yes	Yes
Observations	1,136,070	3,243,402	1,136,070	3,243,402

Notes: This table reports IV estimates separately for small-entity and large-entity patent applicants. Small entities (individuals, small firms, nonprofits) file under the USPTO's reduced-fee program. The instrument is the leave-one-out examiner grant rate within art-unit × filing-year cells. Standard errors clustered at the art-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4 Robustness

Table 5 demonstrates that the main result is insensitive to alternative instrument definitions. Constructing the leave-one-out rate within 3-year filing windows (rather than single years) yields an identical estimate of 0.135. Trimming the top and bottom 5 percent of the leniency distribution produces 0.136. These results confirm that the finding is not driven by extreme examiners or by the specific granularity of the instrument cells.

Table 5: Robustness of IV Estimates for Market Transfer

Specification	Coefficient	SE	N
Baseline	0.1340	(0.0056)	4,379,472
3-Year Instrument Window	0.1347	(0.0056)	4,379,472
Trimmed (5–95% leniency)	0.1359	(0.0068)	4,167,687

Notes: All specifications instrument Patent Granted with examiner leniency and include art-unit × filing-year fixed effects. Row 1 reproduces the baseline from Table 3. Row 2 constructs the leave-one-out instrument within 3-year filing windows (coarser cells). Row 3 drops applications assigned to examiners in the top or bottom 5% of the leniency distribution. Standard errors clustered at the art-unit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.5 Threats to Validity

Two concerns merit discussion. First, the balance-test failure on small-entity status raises the possibility that examiner assignment is not perfectly random within cells. However, controlling for entity status changes the IV estimate by less than 2 percent (0.134 to 0.132), and the trimmed-sample estimate is similarly unaffected, indicating that this source of confounding is economically negligible.

Second, the assignment outcome captures transfers that occur both during and after prosecution. For granted patents, post-grant transfers dominate; for abandoned applications, only pre-disposal transfers are possible. Part of the IV effect may therefore reflect the mechanical unlocking of a post-grant trading window rather than a pure valuation response. Future work with time-stamped assignment data could decompose the effect into pre- and post-disposal components.

Third, the exclusion restriction could be violated if lenient examiners systematically produce broader or narrower claims, independently affecting marketability. While I cannot directly test this, the near-identity of OLS and IV estimates is informative: if the exclusion restriction were seriously violated, one would expect the IV to diverge from OLS in a direction dictated by the direct effect channel. The concordance instead suggests that the dominant channel runs through the grant decision itself.

6. Discussion

The central finding is that marginal patent grants create substantial real market value: a 13.4 percentage point increase in the probability of market transfer and a 7.3 point increase in collateralization. Three implications follow.

Patent quality and market discipline. Lemley (2001) argued that the Patent Office rationally underinvests in examination quality because most patents are never litigated. The implicit assumption is that weak patents are harmless "paper rights." My results challenge this view: marginal patents are not paper—they are actively traded. For market transfers, the OLS and IV estimates are nearly identical, suggesting that formal legal status, rather than intrinsic quality, drives tradability at the margin. For collateralization, the IV estimate exceeds OLS (7.3 vs. 6.2 pp), consistent with marginal patents being disproportionately useful as pledgeable assets for constrained borrowers.

Implications for NPEs. The liquidity of marginal patents has direct implications for patent assertion entities. If tightening examination standards at the margin would reduce the

stock of tradeable patents by even a fraction of the 13.4 percentage point margin, this would meaningfully shrink the pool of assets available to NPEs. Conversely, the current system generates a flow of marginal-quality but fully liquid patent rights into secondary markets. Whether this liquidity represents efficient reallocation (Serrano, 2010) or rent extraction (Cohen et al., 2019) cannot be determined from assignment records alone—distinguishing productive transfers from strategic acquisitions requires data on downstream enforcement and licensing, which I leave to future work.

Collateral and innovation finance. The 7.3 percentage point effect on security interests implies that patent grants serve a financing function beyond their role in product-market exclusion. For small entities and startups, the ability to pledge patent rights as collateral may relax credit constraints (Farre-Mensa et al., 2020; Gans et al., 2006). This collateral channel suggests an underappreciated benefit of the patent system that operates through financial markets rather than product markets.

7. Conclusion

Are marginally granted patents mere paper, or are they real economic assets? Using the examiner leniency instrument on 4.4 million USPTO applications linked to the universe of assignment records, I find that the answer is unambiguously the latter. A patent grant caused by examiner leniency is just as likely to be traded as one caused by application quality. This has implications for how we think about patent reform: tightening examination standards would not merely reduce the count of low-quality patents—it would reduce the supply of liquid assets in the market for technology.

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

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin, "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 1996, *91* (434), 444–455.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella, "Markets for Technology: The Economics of Innovation and Corporate Strategy," 2001.

Cockburn, Iain M., Samuel Kortum, and Scott Stern, "Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes," 2003. NBER Working Paper 8980.

Cohen, Lauren, Umit G. Gurun, and Scott Duke Kominers, "Patent Trolls: Evidence from Targeted Firms," *Management Science*, 2019, *65* (12), 5461–5486.

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.

Feng, Josh and Xavier Jaravel, "Crafting Intellectual Property Rights: Implications for Patent Assertion Entities, Litigation, and Innovation," *American Economic Review: Insights*, 2020, *2* (2), 241–258.

Figueroa, Nicolás and Carlos J. Serrano, "Patent Trading Flows of Small and Large Firms," *Research Policy*, 2019, *48* (7), 1601–1616.

Frakes, Michael D. and Melissa F. Wasserman, "Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents?," *Review of Economics and Statistics*, 2017, *99* (3), 550–563.

_ **and** _ , "Irrational Ignorance at the Patent Office," *Vanderbilt Law Review*, 2019, *72* (3), 975–1034.

Galasso, Alberto, "The Market for Patent Rights," *RAND Journal of Economics*, 2012, *43* (2), 167–192.

_ **and Mark Schankerman**, "Patents and Cumulative Innovation: Causal Evidence from the Courts," *Quarterly Journal of Economics*, 2015, *130* (1), 317–369.

Gans, Joshua S., David H. Hsu, and Scott Stern, "Is a Patent Worth Having? The Value of Patent Protection for Start-ups," *Journal of Economic Behavior and Organization*, 2006, *61* (4), 684–700.

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.

Marco, Alan C., Michael Carley, Steven Jackson, and Amanda Myers, "The USPTO Historical Patent Data Files: Two Centuries of Innovation," *USPTO Economic Working Paper*, 2015, (2015-1).

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.

Serrano, Carlos J., "The Dynamics of the Transfer and Renewal of Patents," *RAND Journal of Economics*, 2010, *41* (4), 686–708.

A. Data Appendix

A.1 Sample Construction

The analysis sample is constructed from two sources accessed via Google BigQuery's public patent data:

1. **Patent Examination Research Dataset (PatEx):** Table `application_data` from the `uspto_oce_pair` dataset in Google BigQuery's `patents-public-data` project. Contains the universe of patent applications with metadata including filing date, examiner ID, art unit, disposal type (issued/abandoned/pending), and small-entity indicator.

2. **USPTO Assignment Dataset:** Tables `assignment`, `assignment_conveyance`, and `documentid` from the `uspto_oce_assignment` dataset. Contains all recorded assignments, security interests, and mergers, linked to applications via application number.

Sample restrictions applied sequentially:

- Resolved applications only (disposal type ISS or ABN): 7,087,735
- Non-missing examiner ID and art unit: 6,851,204
- Filing year 2000–2015 (sufficient post-filing time for assignment): 4,454,050
- Examiner has $\geq$50 resolved applications: 4,379,472 (final sample)

A.2 Variable Definitions

Market Transfer: Equals one if the application has at least one assignment record where the conveyance type or text indicates an ownership transfer (contains "assign" but not "security"), excluding employer assignments (where the conveyance is flagged as an employer-employee transfer).

Security Interest: Equals one if any conveyance record contains "security," "lien," or "mortgage."

Examiner Leniency (LOO): The leave-one-out grant rate of the assigned examiner within the same art-unit $\times$ filing-year cell. Undefined (and dropped) when the examiner has only one application in the cell.

B. Identification Appendix

B.1 Balance Tests

Examiner leniency predicts small-entity status within art-unit–year cells (coefficient = -0.095, SE = 0.009, $p < 0.001$). This is a known feature of examiner-IV designs: within narrow cells, some residual sorting may occur. The economic magnitude is modest—a one-SD increase in leniency is associated with a 2.1 pp decrease in small-entity probability—and including small-entity status as a control changes the main IV estimate by less than 2%.

B.2 Monotonicity

Grant rates increase monotonically across examiner leniency deciles, from 32.0% (D1) to 96.1% (D10). This is consistent with the monotonicity assumption required for LATE interpretation.

C. Robustness Appendix

C.1 Alternative Instrument Windows

Using 3-year filing windows instead of single-year windows for the leave-one-out calculation produces a nearly identical IV estimate (0.135 vs. 0.134), confirming that the result is not sensitive to the granularity of the instrument cells.

C.2 Trimmed Sample

Excluding examiners in the top and bottom 5% of the leniency distribution yields an IV estimate of 0.136 on a sample of 4,167,687 applications, ruling out the possibility that extreme examiners drive the result.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled							
Market Transfer	IV (Pooled)	0.1340	0.0056	0.404	0.3320	0.0139	Large positive
Security Interest	IV (Pooled)	0.0727	0.0043	0.283	0.2567	0.0153	Large positive
Panel B: Heterogeneous (by Entity Size)							
Market Transfer	IV (Small Entity)	0.1275	0.0058	0.374	0.3405	0.0156	Large positive
Market Transfer	IV (Large Entity)	0.1342	0.0068	0.413	0.3253	0.0164	Large positive

Notes: **Country:** United States. **Research question:** Does a quasi-randomly assigned patent grant cause the underlying invention to be traded or pledged as collateral in the secondary patent market? **Policy mechanism:** The USPTO examines patent applications and decides whether to grant a patent right conferring 20 years of exclusivity; the grant decision determines whether the invention has a formal property right that can be transferred, licensed, or collateralized. **Outcome definition:** Market Transfer equals one if the patent's USPTO assignment records contain at least one non-employer ownership change; Security Interest equals one if collateralized. **Treatment:** Binary (granted = 1 vs.
abandoned = 0), instrumented by examiner leniency. **Data:** USPTO PatEx and Assignment datasets via Google BigQuery, patent applications filed 2000–2015, matched via application number. **Method:** 2SLS with leave-one-out examiner grant rate, art-unit × filing-year FE, art-unit-clustered SEs. **Sample:** Resolved applications with examiner caseload ≥ 50. SDE $= \hat{\beta}/\text{SD}(Y)$. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).