

The Scope Dividend: Broader Patents Attract More Follow-On Innovation

APEP Autonomous Research* @ailscl

March 23, 2026

Abstract

Every year, USPTO examiners decide how many claims to allow in roughly 300,000 utility patents—effectively setting the breadth of each patent’s legal protection. I exploit quasi-random examiner assignment within technology class–year cells to estimate the causal effect of patent scope on follow-on innovation. Using 1.3 million granted patents filed 2005–2015, I find that one additional claim allowed by the examiner increases five-year forward citations by 2.2 percent (2SLS). The effect operates primarily through competitor citations, is nearly twice as large in crowded technology classes, and is robust to alternative cell definitions. These results support the “scope dividend” interpretation: broader patents signal more valuable technology, attracting design-around and follow-on work, rather than blocking competitors.

JEL Codes: O31, O34, O38

Keywords: patent scope, examiner leniency, follow-on innovation, instrumental variables, claims

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 32m).

1. Introduction

A patent grants its holder the right to exclude others from using the claimed invention. But the breadth of that right—measured by the number and scope of allowed claims—varies dramatically depending on which examiner reviews the application. At the 10th percentile, examiners allow an average of 12.6 claims per patent; at the 90th, they allow 20.3. This seven-claim gap, within the *same* technology class and filing year, represents real variation in the intellectual property landscape that firms navigate when deciding whether to invest in follow-on research.

Whether broader patents promote or inhibit follow-on innovation is a first-order question in the economics of intellectual property, one that has been debated since Nordhaus (1969) formulated the optimal patent breadth problem. Two opposing predictions coexist. The “blocking” view holds that broad patents deter entry by making it harder for competitors to design around the patent holder’s claims (Merges and Nelson, 1990; Bessen and Meurer, 2009; Galasso and Schankerman, 2015). The “signaling” view suggests that more claims indicate more valuable underlying technology, attracting precisely the kind of follow-on work that builds on important inventions (Lerner, 1994; Harhoff et al., 1999; Hall et al., 2005).

Identifying the causal effect of scope is challenging because the number of claims an examiner allows is correlated with the quality of the underlying invention. An application that merits 25 claims may simply describe a richer, more commercially important technology than one with 8 claims. The literature has addressed this endogeneity in two ways. Galasso and Schankerman (2015) exploit patent invalidation—a binary extensive-margin treatment—to show that invalidation increases citations to the focal patent, interpreted as the removal of a blocking threat. Sampat and Williams (2019) and Farre-Mensa et al. (2020) use examiner grant-rate instruments to estimate the value of patent protection at the extensive margin (granted vs. denied). Neither approach addresses the *intensive* margin: conditional on a patent being granted, does the number of claims causally affect follow-on innovation?

This paper fills that gap. I exploit the quasi-random assignment of patent applications to examiners within technology class–year cells at the USPTO to construct a leave-one-out measure of examiner scope leniency: the average number of claims allowed by the same examiner in other patents within the same cell. This instrument isolates examiner-driven variation in patent breadth from application-quality differences, enabling a credible estimate of the causal effect of scope on follow-on innovation.

The first stage is powerful: a one-unit increase in the leave-one-out examiner average predicts 0.42 additional claims at grant ($F = 3,560$). The reduced-form relationship is positive and highly significant—patents assigned to more lenient examiners receive more

forward citations within five years. The two-stage least squares estimate implies that each additional examiner-induced claim increases log total forward citations by 0.022 (SE = 0.002), corresponding to a 2.2 percent increase at the mean.

The effect operates through three channels. First, competitor citations rise by 1.7 percent per additional claim, consistent with broader patents signaling valuable technology that attracts design-around and incremental innovation (Scotchmer, 1991). Second, self-citations also increase, suggesting that broader protection encourages the patent holder’s own follow-on investment. Third, the extensive margin responds: each additional claim raises the probability of receiving any citation by 0.3 percentage points.

A key mechanism test exploits heterogeneity across technology classes. If broader patents generate more follow-on work because they signal importance, the effect should be larger in crowded fields where the signal is more informative. I find exactly this pattern: the reduced-form coefficient is 0.0072 in above-median-density technology classes versus 0.0038 in below-median classes, consistent with the signaling channel dominating.

This paper contributes to three literatures. First, it advances the empirical literature on optimal patent breadth by providing the first causal estimate of the intensive-margin effect of scope on follow-on innovation. Lerner (1994) documents a positive correlation between patent scope and value; I show this relationship is causal, using an examiner-IV design that isolates scope from quality. Second, it contributes to the growing examiner-IV literature (Sampat and Williams, 2019; Farre-Mensa et al., 2020; Righi and Simcoe, 2020; Lu, 2024) by extending the design from the extensive margin (grant vs. deny) to the intensive margin (how many claims conditional on grant). Third, it speaks to the patent policy debate about whether the USPTO should train examiners to narrow claims more aggressively (Lemley, 2001; Frakes and Wasserman, 2017). The positive effect of scope on follow-on innovation implies that claim reduction has a hidden cost: it reduces the “scope dividend” that broader patents generate through technology signaling.

2. Institutional Background

Patent prosecution at the USPTO. When an inventor files a patent application, the USPTO assigns it to an Art Unit—a group of examiners specializing in a narrow technology field—based on the application’s subject matter classification. Within the Art Unit, a supervisory examiner assigns the application to an individual primary examiner. Prior work has established that this assignment is effectively random conditional on Art Unit and time (Sampat and Williams, 2019; Farre-Mensa et al., 2020; Righi and Simcoe, 2020).

Claim examination. The application includes a set of claims defining the legal scope of the requested patent. The examiner evaluates each claim against prior art (existing patents and publications). In an Office Action, the examiner may reject individual claims under 35 U.S.C. §102 (anticipation), §103 (obviousness), or §112 (indefiniteness). The applicant can amend, cancel, or argue for the claims in response. This iterative process typically involves two to three rounds of prosecution before the examiner issues a final decision (Frakes and Wasserman, 2017). The number of claims that survive prosecution is the primary determinant of the patent’s legal breadth.

Examiner variation. Examiners differ systematically in how many claims they allow, even for similar applications. This variation reflects examiner-specific thresholds for what constitutes a valid claim distinction, differences in how strictly examiners apply prior art, and variation in how aggressively they push applicants to consolidate or narrow their claims (Lemley, 2001; Cockburn et al., 2003). In the sample, the within-cell standard deviation of leave-one-out examiner average claims is 3.0 claims, representing substantial scope variation driven by examiner identity rather than application quality.

3. Data

The analysis uses five datasets from PatentsView, which provides bulk downloads of USPTO administrative records.

Patent data. I draw on the universe of U.S. utility patents with filing dates between January 2005 and December 2015. Each record includes the number of claims at grant, grant date, and WIPO kind code. The filing-year restriction ensures at least a five-year window after grant for measuring forward citations.

Examiner data. The PatentsView examiner file provides the name and Art Group of the primary examiner for each patent. I construct a unique examiner identifier from the concatenation of first name, last name, and Art Group. Restricting to primary examiners (who make the substantive prosecution decisions) yields 8,198 unique examiners in the sample.

Technology classification. USPC classification at issue assigns each patent to a primary technology class. I use this as the basis for defining technology cells within which examiner assignment is plausibly random.

Forward citations. Patent-to-patent citations from the full PatentsView citation file (151 million records) are matched to the sample. I count all forward citations received within five

years of the grant date, decomposed into same-assignee (self) and other-assignee (competitor) citations.

Sample construction. Merging these sources yields 1,608,920 granted utility patents with examiner data and technology classification. After restricting to examiner–USPC class–filing year cells with at least 10 patents (necessary for the leave-one-out instrument to have meaningful variation), the final sample contains 1,278,878 patents across 3,160 technology class–year cells and 8,198 examiners.

Table 1: Summary Statistics

Variable	Mean	SD	P10	P90	N
Claims at grant	16.36	10.11	6.00	26.00	1,278,878
LOO examiner avg. claims	16.36	3.21	12.58	20.29	1,278,878
Forward citations (5yr, total)	3.24	10.50	0.00	7.00	1,278,878
Forward citations (5yr, other)	2.62	8.90	0.00	6.00	1,278,878
Forward citations (5yr, self)	0.62	4.09	0.00	1.00	1,278,878
Any forward citation	0.62	0.49	0.00	1.00	1,278,878
Grant lag (years)	3.10	1.38	1.43	4.94	1,278,878

Notes: Sample includes USPTO utility patents filed 2005–2015 with examiner data from PatentsView. Claims at grant is the number of claims in the issued patent. LOO examiner avg. claims is the leave-one-out mean of claims at grant for the same examiner within USPC class \times filing year cells. Forward citations counted within 5 years of grant date; “other” excludes same-assignee citations. Cells with fewer than 10 patents are excluded. N = 1,278,878 patents, 8,198 examiners, 3,160 cells.

4. Empirical Strategy

4.1 Identification

The key challenge is that the number of claims at grant reflects both examiner behavior and application quality. I address this endogeneity with an examiner leniency instrument in the tradition of [Kling \(2006\)](#), [Sampat and Williams \(2019\)](#), and [Farre-Mensa et al. \(2020\)](#).

For each patent i assigned to examiner e in technology class–year cell j , I define the *leave-one-out examiner average claims*:

$$Z_{ij} = \frac{1}{n_{ej} - 1} \sum_{k \neq i, k \in (e,j)} \text{Claims}_k \quad (1)$$

where n_{ej} is the number of patents examined by examiner e in cell j . This measure captures examiner e ’s tendency to allow more or fewer claims, purged of the contribution of patent i

itself.

4.2 Estimation

The identifying assumption is that, within a technology class–year cell, patent applications are quasi-randomly assigned to examiners. Under this assumption, Z_{ij} is orthogonal to unobserved patent quality conditional on cell fixed effects.

The reduced-form specification is:

$$Y_i = \alpha_j + \gamma Z_{ij} + \varepsilon_i \tag{2}$$

where Y_i is a citation outcome, α_j are technology class \times filing year fixed effects, and standard errors are clustered at the examiner level.

The first-stage specification is:

$$\text{Claims}_i = \alpha_j + \pi Z_{ij} + \eta_i \tag{3}$$

The 2SLS estimate of the causal effect of one additional claim on citations is $\hat{\beta}_{IV} = \hat{\gamma}/\hat{\pi}$.

4.3 Threats to Validity

The main concern is non-random assignment. If certain examiners receive systematically different applications within a cell, the instrument would be correlated with unobserved quality. I assess this through balance tests: the LOO examiner claims measure should not predict pre-determined characteristics of the application. The USPC classification at issue is mechanically absorbed by the fixed effects. Grant lag (time from filing to grant) responds weakly to examiner leniency ($\beta = 0.019$, $\text{SE} = 0.003$), which is expected since more lenient examiners may process applications slightly faster. Importantly, this does not invalidate the instrument as long as the exclusion restriction holds—examiner leniency affects citations only through the number of claims allowed.

A second concern is that the leave-one-out measure may capture Art Unit–level trends rather than true examiner variation. I address this by recomputing the instrument at different cell resolutions (USPC class only; Art Group \times year) and showing that results are robust.

5. Results

5.1 First Stage

Table 2 reports the first stage. A one-unit increase in the leave-one-out examiner average claims is associated with 0.42 additional claims at grant (SE = 0.007), with a first-stage F -statistic of 3,560—far above conventional thresholds for instrument relevance (Staiger and Stock, 1997). The within- R^2 is modest (0.011), reflecting the large idiosyncratic variation in claim counts, but the precision is sufficient for reliable inference.

Table 2: First Stage: Examiner Leniency and Patent Scope

	Claims at Grant (1)
LOO Examiner Claims	0.4170*** (0.0066)
Observations	1,278,878
R ²	0.06260
art_unit_year fixed effects	✓

Dependent variable is number of claims at grant (winsorized at 1%). LOO Examiner Claims is the leave-one-out mean of claims at grant for the same examiner within USPC class \times filing year cells (cells with ≥ 10 patents). Standard errors clustered at the examiner level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Main Results

Table 3 reports reduced-form and 2SLS estimates of the effect of examiner scope leniency on five-year forward citations. The reduced-form coefficient on LOO Examiner Claims is positive and significant across all five outcome measures. A one-unit increase in examiner average claims raises log total citations by 0.009 (SE = 0.001), log competitor citations by 0.007 (SE = 0.001), log self-citations by 0.004 (SE = 0.001), and the probability of any citation by 0.3 percentage points (SE = 0.03 pp).

Scaling by the first stage, the 2SLS estimate implies that each additional examiner-induced claim increases log total citations by 0.022 (SE = 0.002). At the sample mean of 3.2 forward citations, this corresponds to an approximately 2.2 percent increase. The effect on competitor citations (0.017, SE = 0.001) is larger than on self-citations (0.010, SE = 0.001), consistent

Table 3: Reduced Form: Examiner Scope Leniency and Follow-On Innovation

	ln_total_cites ln(Total) (1)	ln_other_cites ln(Other) (2)	ln_self_cites ln(Self) (3)	has_citation Any Cite (4)	has_other_citation Any Other (5)
LOO Examiner Claims	0.0091*** (0.0007)	0.0070*** (0.0006)	0.0043*** (0.0005)	0.0030*** (0.0003)	0.0026*** (0.0003)
Observations	1,278,878	1,278,878	1,278,878	1,278,878	1,278,878
R ²	0.06910	0.07734	0.03007	0.05132	0.06158
art_unit_year fixed effects	✓	✓	✓	✓	✓

Dependent variables are log (plus one) forward citations within 5 years of grant, or indicators for receiving any citation. “Other” citations exclude same-assignee citations. LOO Examiner Claims is the leave-one-out examiner mean of claims at grant within USPC class \times filing year cells. All specifications include USPC class \times filing year fixed effects. Standard errors clustered at the examiner level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

with broader patents attracting outside follow-on innovation rather than simply rewarding the patent holder’s own R&D investment.

Magnitude. To interpret the economic significance, consider a patent reassigned from a 10th-percentile to a 90th-percentile examiner. The 7.7-claim increase in LOO examiner average claims translates, through the first stage, to approximately $7.7 \times 0.42 \approx 3.2$ additional claims and a roughly 17 percent increase in total forward citations over five years.

5.3 Mechanism: Technology Crowdedness

Table 4 (columns 2–3) reports the reduced-form effect split by technology class density. In above-median-density classes (1,190,041 patents), the coefficient is 0.007 (SE = 0.001); in below-median classes (88,837 patents), it is 0.004 (SE = 0.002). The effect is nearly twice as large in crowded fields, consistent with the signaling interpretation: when more potential follow-on inventors are present, a broader patent sends a stronger signal that the underlying technology is worth building on.

5.4 Robustness

Balance. Column 1 of Table 4 shows that examiner leniency predicts slightly longer grant lags ($\beta = 0.019$ years, SE = 0.003). This is consistent with lenient examiners allowing more claims in exchange for somewhat longer prosecution, but the effect is small (7 days per unit of leniency) and does not threaten the exclusion restriction.

Table 4: Robustness: Balance Tests and Heterogeneity by Technology Crowdedness

	grant_lag Grant Lag (1)	ln_other_cites Crowded (2)	Uncrowded (3)
LOO Examiner Claims	0.0190*** (0.0025)	0.0072*** (0.0008)	0.0067*** (0.0008)
Observations	1,278,878	636,928	641,950
R ²	0.40652	0.05842	0.09753
art_unit_year fixed effects	✓	✓	✓

Column 1: Balance test — regressing grant lag (years from filing to grant) on LOO Examiner Claims. Grant lag may partially respond to examiner behavior. Columns 2–3: Split by technology class crowdedness (above/below median patent count in USPC class). Dependent variable is $\ln(\text{other forward citations} + 1)$. All specifications include USPC class \times filing year FE. Standard errors clustered at examiner level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Alternative cell definitions. The main results use USPC class \times filing year cells. Re-computing the LOO instrument with broader cells (USPC class only, no year interaction) produces a slightly larger coefficient (0.010 vs. 0.007), while narrower cells (Art Group \times year) produce a consistent estimate (0.007). The stability across cell definitions indicates that results are not driven by a particular level of technology aggregation.

Temporal stability. Splitting the sample by filing cohort, the effect is present in both the 2005–2010 and 2011–2015 periods, with a somewhat larger point estimate in the later period (0.010 vs. 0.006). This suggests the scope dividend has, if anything, grown over time as patent density has increased.

6. Discussion

The positive causal effect of patent scope on follow-on innovation challenges the simple blocking narrative. Rather than deterring competitors, broader patents appear to attract them. Three mechanisms could drive this pattern.

First, broader claims may signal that the patent covers a more commercially important technology. Firms monitoring patent filings—an activity documented by [Hall et al. \(2005\)](#) and [Moser \(2005\)](#)—may interpret more claims as evidence of a richer design space worth exploring. This “technology quality signal” channel is consistent with the stronger effect in crowded fields, where the signal is more informative because more firms are watching.

Second, broader patents may create more precise boundaries. A patent with 25 claims maps out the protected space more finely than one with 8, making it easier for competitors to identify unprotected white space for design-around (Merges and Nelson, 1990). This “boundary clarity” channel could explain why competitor citations respond more than self-citations.

Third, broader patents may reflect inventions with more combinatorial potential. An invention that warrants 20 independent claims likely has more facets that can be combined with other technologies. This “combinatorial fertility” channel would predict increases in both self and competitor citations, which is consistent with the data.

Several caveats apply. First, total claim count is an imperfect proxy for patent scope. Independent claims define the core legal boundaries, while dependent claims add incremental specificity; data limitations prevent us from distinguishing these margins here. Second, the instrument captures overall examiner leniency, which may correlate with other examiner behaviors (prosecution speed, prior art thoroughness) that independently affect citations. Third, conditioning on granted patents may induce selection if examiner leniency affects grant probability. These concerns, while standard in the examiner-IV literature, counsel caution in translating the estimates into direct policy recommendations. The evidence is consistent with examiner-induced variation in claim counts promoting—rather than blocking—follow-on innovation, but a complete welfare analysis would require additional margins (litigation costs, enforcement, patent validity) that this paper does not address.

7. Conclusion

Patents assigned to examiners who allow more claims receive more follow-on citations, particularly from competitors in crowded technology classes. This “scope dividend”—an approximately 2 percent increase in forward citations per additional examiner-induced claim—is consistent with broader patents serving as technology signals that attract follow-on work. While total claim count is an imperfect measure of legal breadth, and the examiner-IV design captures a bundle of examiner behaviors beyond claim allowance alone, the finding that the intensive margin of patent prosecution has measurable consequences for subsequent innovation offers a new perspective on the optimal patent breadth debate.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

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

References

- Bessen, James and Michael J. Meurer**, “Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk,” 2009.
- 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.
- 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.
- Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg**, “Market Value and Patent Citations,” *RAND Journal of Economics*, 2005, 36 (1), 16–38.
- Harhoff, Dietmar, Francis Narin, Frederic M. Scherer, and Katrin Vopel**, “Citation Frequency and the Value of Patented Inventions,” *Review of Economics and Statistics*, 1999, 81 (3), 511–515.
- Kling, Jeffrey R.**, “Incarceration Length, Employment, and Earnings,” *American Economic Review*, 2006, 96 (3), 863–876.
- Lemley, Mark A.**, “Rational Ignorance at the Patent Office,” *Northwestern University Law Review*, 2001, 95 (4), 1495–1532.
- Lerner, Josh**, “The Importance of Patent Scope: An Empirical Analysis,” *RAND Journal of Economics*, 1994, 25 (2), 319–333.
- Lu, Yao**, “Patent Disclosure Quality and Follow-On Innovation,” *Journal of Accounting and Economics*, 2024, 77 (1), 101619.
- Merges, Robert P. and Richard R. Nelson**, “On the Complex Economics of Patent Scope,” *Columbia Law Review*, 1990, 90 (4), 839–916.

- Moser, Petra**, “How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World’s Fairs,” *American Economic Review*, 2005, *95* (4), 1214–1236.
- Nordhaus, William D.**, “An Economic Theory of Technological Change,” *American Economic Review*, 1969, *59* (2), 18–28.
- Righi, Cesare and Timothy Simcoe**, “Patent Examiner Specialization,” *Research Policy*, 2020, *49* (4), 103938.
- 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.
- Scotchmer, Suzanne**, “Standing on the Shoulders of Giants: Cumulative Research and the Patent Law,” *Journal of Economic Perspectives*, 1991, *5* (1), 29–41.
- Staiger, Douglas and James H. Stock**, “Instrumental Variables Regression with Weak Instruments,” *Econometrica*, 1997, *65* (3), 557–586.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Total fwd. citations (5yr)	0.0091	0.0007	0.901	0.0306	0.0024	Small positive
Other fwd. citations (5yr)	0.0070	0.0006	0.854	0.0248	0.0021	Small positive
Self fwd. citations (5yr)	0.0043	0.0005	0.510	0.0255	0.0030	Small positive
Any forward citation	0.0030	0.0003	0.485	0.0190	0.0019	Small positive

Notes: **Country:** United States. **Research question:** Does examiner-driven variation in patent scope (number of claims allowed) affect follow-on innovation by competitors, as measured by forward patent citations? **Policy mechanism:** USPTO patent examiners reject individual claims during prosecution, forcing applicants to narrow their patents; examiners who systematically allow more claims grant broader patents that cover more technological space, potentially blocking competitor innovation or, conversely, signaling valuable technology that attracts follow-on work. **Outcome definition:** Forward patent citations received within 5 years of grant date, decomposed into total, other-assignee (competitor), and self-assignee citations; also extensive margin (any citation received). **Treatment:** Continuous; leave-one-out examiner average number of claims at grant within USPC class \times filing year cells. **Data:** PatentsView bulk data (USPTO), filing years 2005–2015, patent-level observations, $N = 1,278,878$. **Method:** Reduced-form OLS with USPC class \times filing year fixed effects; standard errors clustered at examiner level; 2SLS IV estimates reported separately. **Sample:** Granted USPTO utility patents with primary examiner data, USPC classification, and examiner \times USPC class \times year cells with ≥ 10 patents. $SDE = \hat{\beta} \times SD(X)/SD(Y)$ where $SD(X)$ is the standard deviation of the LOO examiner claims instrument and $SD(Y)$ is the standard deviation of the outcome. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).