

Does Examiner Leniency Affect Cumulative Green Innovation?

Evidence from USPTO Application-Level Data

APEP Autonomous Research* @SocialCatalystLab

March 6, 2026

Abstract

Does assignment to a more permissive patent examiner affect cumulative green innovation? Using application-level data from the Patent Examination Research Dataset—grants and abandonments—I exploit quasi-random examiner assignment within USPTO art units among 640,845 green-technology applications filed 2001–2012. A one standard deviation increase in examiner grant rate raises grant probability by 15.1 percentage points ($F > 13,000$). Because follow-on patenting varies only at the CPC subclass-by-filing-year level (96 cells), I estimate the reduced form there: a one standard deviation increase in average permissiveness is associated with 19 percent less follow-on Y02 patenting ($p = 0.025$). However, this collapsed estimate does not inherit application-level random assignment, and collapsing to the art-unit-by-filing-year level produces a null. The evidence does not support the conclusion that marginal patent-office decisions bind cumulative green innovation.

JEL Codes: O34, O31, Q55, Q58

Keywords: green patents, clean energy innovation, patent examiner, instrumental variables, technology diffusion, climate change mitigation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

Solar panel costs have fallen 99 percent since 1976. Wind turbine efficiency has tripled since the 1990s. Lithium-ion battery energy density doubles roughly every decade. Behind each of these achievements lies a chain of patented inventions—the Y02 classification in the Cooperative Patent Classification system covers hundreds of thousands of US patents across energy generation, buildings, transportation, carbon capture, and industrial production. A persistent policy question is whether the intellectual property system that incentivized these inventions simultaneously impedes the diffusion of the technologies needed to avert climate catastrophe.

The theoretical tension is well understood. [Nordhaus \(1969\)](#) established the case for patent protection as an incentive for private R&D. But [Scotchmer \(1991\)](#) and [Heller and Eisenberg \(1998\)](#) showed how sequential innovation creates a fundamental conflict: if today’s patent blocks tomorrow’s improvement, cumulative progress may slow even as the inventor captures returns. In clean energy, where the IPCC estimates that investment must triple by 2030 ([IPCC, 2023](#)), the stakes are acute. Whether the patent system promotes or hinders follow-on green innovation is an empirical question—one that requires exogenous variation in patent outcomes.

I exploit quasi-random assignment of patent applications to USPTO examiners within technology-specific art units—a design validated by [Lemley and Sampat \(2012\)](#) and deployed in genomics by [Sampat and Williams \(2019\)](#) and for startup patents by [Farre-Mensa et al. \(2020\)](#). A crucial feature of this study is the use of application-level data from the Patent Examination Research Dataset (PatEx), which contains the universe of USPTO utility patent applications including both grants and abandonments. This avoids conditioning on a post-treatment outcome (grants-only analyses restrict to applications that were in fact granted, discarding variation from examiner strictness) and enables a direct first stage: does examiner grant rate predict whether an application is granted?

Among 640,845 green-technology patent applications filed 2001–2012 in Y02-technology art units, a one standard deviation increase in examiner grant rate raises the probability of grant by 15.1 percentage points ($F > 13,000$). The first stage is the paper’s cleanest contribution: by including both granted and abandoned applications, I avoid the selection problem inherent in grants-only examiner IV designs.

The downstream analysis faces a fundamental design challenge. The available follow-on patenting outcome—new Y02 patents filed in the same CPC subclass within five years—varies only at the subclass-by-filing-year level: 640,845 application observations share just 96 unique outcome values (8 subclasses \times 12 years). Application-level regressions therefore

involve massive pseudo-replication, and standard errors clustered at the examiner level overstate precision. I address this by collapsing the data to the level at which the outcome genuinely varies and present both subclass-by-year (96 cells) and art-unit-by-year (980 cells) aggregations.

The subclass-by-year collapse produces a negative estimate (-0.193 , $p = 0.025$): subclass-year cells with more permissive average examiner assignment see less subsequent patenting. But this collapsed design does not inherit the application-level random-assignment guarantee, because examiner assignment is quasi-random within art-unit-by-year, not within subclass-by-year. The art-unit-by-year collapse—which better matches the assignment structure—yields a null (-0.011 , $SE = 0.014$). The sign and significance of the reduced-form relationship depend on the aggregation level, and I do not privilege one over the other.

Forward citations show a large positive response, but this is substantially mechanical: granted patents exist in the citation database while abandoned applications do not. I report citation results for transparency but do not interpret them causally.

The paper’s contribution is methodological and descriptive. It demonstrates how application-level PatEx data solves the first-stage selection problem in examiner IV designs for green technology, but the downstream follow-on evidence is too fragile—across aggregation levels, clustering choices, and outcome definitions—to support a firm causal conclusion about blocking. This aligns with [Sampat and Williams \(2019\)](#), who found no effect of gene patent grants on follow-on research. The forces shaping green innovation trajectories—carbon prices, R&D subsidies, deployment mandates, and market demand ([Acemoglu et al., 2012](#); [Popp, 2002](#))—appear to dominate the marginal examiner decision.

2. Related Literature

This paper contributes to three literatures: the economics of patents and cumulative innovation, the measurement and determinants of green innovation, and the examiner instrumental variable methodology.

2.1 Patents and Cumulative Innovation

The theoretical literature on sequential innovation has long recognized the tension between incentivizing initial inventions and enabling follow-on improvement. [Scotchmer \(1991\)](#) formalized the problem: if the first innovator captures all the surplus, second-generation inventors lack incentive; if the follower captures too much, the pioneer underinvests. [Green and Scotchmer \(1995\)](#) extended this framework to show how the division of profit between sequential innovators depends on patent breadth and licensing efficiency. [Heller and Eisenberg \(1998\)](#)

introduced the “anticommons” concept—fragmented IP rights can impede research when multiple permission holders must be negotiated with—and applied it to biomedical research.

Empirical evidence on whether patents block or facilitate cumulative innovation remains mixed. [Williams \(2013\)](#) found that Celera Genomics’ IP claims on the human genome reduced subsequent scientific research by 20–30 percent. [Galasso and Schankerman \(2015\)](#) exploited random assignment of patent validity cases to Federal Circuit judges and found that patent invalidation increased subsequent citations, particularly by small firms. [Murray et al. \(2007\)](#) found that placing biological materials under restrictive agreements reduced follow-on research.

However, [Sampat and Williams \(2019\)](#) found no significant effect of gene patent grants on follow-on innovation using a patent examiner IV. [Farre-Mensa et al. \(2020\)](#) used the same design for startup patents and found large effects on firm growth but did not examine follow-on innovation by third parties. The present paper extends this literature to climate change mitigation technologies using application-level data that enables a proper first stage and eliminates the grants-only sample selection concern.

2.2 Green Innovation and Climate Policy

A large literature studies the determinants of green innovation. [Popp \(2002\)](#) established that energy prices induce directed technical change. [Acemoglu et al. \(2012\)](#) formalized the role of directed technical change in environmental policy. [Aghion et al. \(2016\)](#) provided evidence from the auto industry that carbon taxes redirect patenting from dirty to clean technologies. On the supply side, [Johnstone et al. \(2010\)](#) documented that renewable energy policies increase green patenting, and [Calel and Dechezleprêtre \(2016\)](#) showed that the EU Emissions Trading System increased low-carbon patenting by regulated firms.

These papers study the *direction* of inventive effort, taking the patent system as given. A separate question—addressed here—is whether the patent system itself promotes or hinders cumulative innovation within green technologies. [Budish et al. \(2015\)](#) showed that fixed patent terms create distortions in pharmaceutical R&D; analogous distortions may exist in green technology where development cycles are long.

2.3 Examiner IV Designs

The examiner IV methodology exploits quasi-random assignment of cases to decision-makers who vary in their propensity to approve. [Lemley and Sampat \(2012\)](#) validated the design for patent examination, documenting large variation in examiner grant rates within art units. [Farre-Mensa et al. \(2020\)](#) deployed the design at scale with F-statistics exceeding 100. [Frakes](#)

and Wasserman (2017) showed that time pressure induces examiners to grant invalid patents.

Chyn et al. (2024) provide a comprehensive treatment of examiner and judge IV designs, addressing many-examiner bias, the monotonicity assumption, and LATE interpretation. Their framework motivates the application-level design used here: with both granted and abandoned applications in the sample, I can estimate a proper first stage and compute the examiner’s actual grant rate rather than relying on grant share among observed grants. A particular challenge they highlight—shared outcomes across observations within decision-maker cells—is directly relevant here, as the follow-on patenting outcome varies at the CPC subclass-by-year level rather than the application level. The exclusion restriction is also demanding in examiner settings: Frakes and Wasserman (2017) show that examiner behavior varies in prosecution duration, office action patterns, and claim narrowing, all of which could affect downstream innovation independently of the binary grant decision.

3. Institutional Background

3.1 Patent Examination at the USPTO

The United States Patent and Trademark Office receives approximately 650,000 utility patent applications per year. Each application is classified into one of approximately 475 “art units”—small groups of 10–20 examiners specializing in a narrow technology area. Within each art unit, applications are assigned to examiners through a process that is functionally random: the Supervisory Patent Examiner routes incoming cases based on workload, seniority rotation, and first-in-first-out queues, without regard to application quality or applicant identity (Lemley and Sampat, 2012; Cockburn et al., 2003).

Examiners differ substantially in their grant propensity. Lemley and Sampat (2012) documented that some examiners grant over 90 percent of their applications while others grant fewer than 30 percent, even within the same art unit. Farre-Mensa et al. (2020) showed that assignment to a one-standard-deviation more permissive examiner increases the probability of grant by approximately 10 percentage points, with first-stage F-statistics exceeding 100. The examination process typically takes 2–3 years from filing to final disposition (grant or abandonment).

3.2 Y02 Classification: Climate Change Mitigation Technologies

The Cooperative Patent Classification system includes a dedicated scheme for climate change mitigation technologies under the Y02 tag, developed as part of an OECD initiative. Y02 classifications are cross-cutting: a single patent may carry both a primary technology

classification and a Y02 tag. The scheme covers six major domains:

- **Y02E**: Energy generation, transmission, and distribution (solar, wind, hydro, geothermal, fuel cells, smart grids)
- **Y02B**: Climate change mitigation in buildings (insulation, lighting, HVAC)
- **Y02P**: Climate change mitigation in production and processing (chemical, metals, minerals)
- **Y02T**: Climate change mitigation in transportation (road, rail, marine, aerospace, EVs)
- **Y02C**: Carbon capture and storage/sequestration
- **Y02W**: Climate change mitigation in waste management

These domains differ in innovation dynamics. Solar and wind technologies exhibit strong cumulative innovation where each generation builds on prior inventions. Carbon capture involves more discrete inventions. This heterogeneity provides a natural test: if patent protection operates through knowledge disclosure, effects should be strongest in cumulative domains.

4. Data

4.1 Patent Examination Research Dataset (PatEx)

The primary data source is the Patent Examination Research Dataset (PatEx), available through Google BigQuery (`patents-public-data.uspto_oce_pair.application_data`). PatEx contains the universe of USPTO utility patent applications with examiner assignments, filing dates, USPC classifications, and—crucially—disposition outcomes including both grants (ISS) and abandonments (ABN). This is the key advance over PatentsView, which contains only granted patents.

I extract all utility patent applications filed between 2001 and 2012 with resolved dispositions and valid examiner assignments. The sample period ends at 2012 because later filings have incomplete disposition data in PatEx. The resulting dataset contains approximately 3.6 million resolved applications: roughly 2.5 million grants and 1.1 million abandonments, assigned to over 12,000 unique examiners across 718 art units.

4.2 PatentsView CPC Data

I link PatEx applications to PatentsView’s CPC classification data to identify green-technology patents. For granted applications, the link is direct via patent number. For abandoned applications (which have no patent number in PatentsView), I use an art-unit mapping strategy: I identify “Y02 art units” as those where more than 10 percent of resolved applications with matched CPC data carry Y02 codes, with a minimum of 50 Y02 grants. All applications—both granted and abandoned—in these art units constitute the analysis population.

This approach has a clear rationale: applications in the same art unit address the same technology space and are reviewed by the same pool of examiners. An abandoned application in an art unit where the majority of grants are Y02-classified is, with high probability, also in the green-technology space. The Y02 CPC subclass for abandoned applications is imputed using the modal subclass among grants in the same art unit.

4.3 Examiner Grant Rate Construction

For each application i assigned to examiner j in art unit a during filing year t , I construct a leave-one-out examiner grant rate:

$$Z_i = \frac{G_{j(i),a,t} - D_i}{N_{j(i),a,t} - 1} \quad (1)$$

where $G_{j(i),a,t}$ is the number of grants by examiner j in art unit a and filing year t , D_i is the focal application’s own grant status (1 if granted, 0 if abandoned), and $N_{j(i),a,t}$ is the total number of applications handled by examiner j in that cell. The leave-one-out adjustment removes the mechanical correlation between the instrument and the focal application’s outcome.

The instrument is defined within *filing-year* cells (when applications enter the system), and the regression fixed effects also use filing-year cells, ensuring consistency. I require at least three examiners per art-unit-by-filing-year cell to ensure meaningful cross-examiner variation. For regressions, I standardize Z_i to mean zero and unit standard deviation, so all reported coefficients represent the effect of a one-standard-deviation increase in examiner grant rate. Table 1 reports summary statistics for the raw (unstandardized) instrument.

Comparison with grant share. The grant rate Z_i uses the examiner’s own applications as the denominator (grants per application), while a grant share uses the cell’s total grants as the denominator (examiner grants divided by cell grants). The grant rate directly measures the examiner’s propensity to grant among applications assigned to them, which is the conceptually

correct instrument for the quasi-random assignment design. Grant share can conflate grant propensity with workload—an examiner handling more cases mechanically has a higher grant share even with identical propensity.

4.4 Follow-on Innovation Outcomes

The primary outcome is the count of new Y02 patents in the same CPC subclass whose filing dates fall within a specified window (3, 5, or 10 years) after the focal application’s filing date. Using filing date as the base is important for two reasons: abandoned applications have no grant date, and the 5-year and 10-year windows extend well beyond the typical 2–3 year examination period, ensuring that the majority of the outcome is realized after the grant decision. The 3-year window includes some pre-disposition follow-on activity, which attenuates the estimated effect; I emphasize the 5-year and 10-year results accordingly. The reduced-form estimand captures the intention-to-treat effect of examiner assignment from filing onward, which includes both the direct grant decision and any prosecution-period effects (e.g., publication of the application at 18 months).

Application controls—claims and backward citations—are available in PatentsView for granted patents only. For abandoned applications, which have no patent number in PatentsView, these variables are coded as zero. This is a data limitation, not a substantive claim that abandoned applications have zero claims. Including these controls in the full-sample regression creates a “bad control” problem: because examiner grant rate predicts grant status, and grants are the only observations with nonzero claims/citations, the controls mechanically absorb treatment-related variation. This is visible in the first stage, where adding these controls reduces the coefficient from 0.151 to 0.018 (Table 2, columns 1 vs. 2). I present the specification without controls as the baseline and include controls as a sensitivity check.

Forward citations serve as a secondary outcome. Because only granted patents appear in the citation database, this outcome is mechanically zero for all abandoned applications. Any positive effect of examiner permissiveness on forward citations conflates genuine knowledge spillovers with the mechanical fact that granted patents exist as citable documents in the prior art database.

4.5 Summary Statistics

Table 1 presents summary statistics. Panel A describes the full application sample of 640,845 observations; Panel B restricts to grants. The full sample includes both granted and abandoned applications in Y02-technology art units. The examiner grant rate instrument

Table 1: Summary Statistics: Green Patent Applications, 2001–2012

Variable	N	Mean	SD	P25	Median	P75
<i>Panel A: Full Sample (Applications)</i>						
Application Granted	640,845	0.690	0.462	0.000	1.000	1.000
Examiner Grant Rate (LOO)	640,845	0.690	0.204	0.571	0.735	0.848
Number of Claims	640,845	11.3	12.1	0	9	18
Backward Citations	640,845	10.6	30.7	0	4	10
Follow-on Y02 Patents (3yr)	640,845	8196.9	5793.8	4,519	6,797	9,684
Follow-on Y02 Patents (5yr)	640,845	15426.5	10998.3	8,450	11,822	18,213
Follow-on Y02 Patents (10yr)	640,845	40729.5	25828.8	21,071	32,220	54,752
Forward Citations	640,845	7.5	29.0	0	1	6
<i>Panel B: Granted Patents Only</i>						
Examiner Grant Rate (LOO)	442,292	0.741	0.173	0.645	0.779	0.870
Number of Claims	442,292	16.3	11.4	9	15	20
Backward Citations	442,292	15.4	35.9	3	7	15
Follow-on Y02 Patents (3yr)	442,292	8068.2	5704.2	4,633	6,797	8,640
Follow-on Y02 Patents (5yr)	442,292	15132.1	10824.2	8,450	11,787	17,476
Follow-on Y02 Patents (10yr)	442,292	39845.9	25404.4	21,071	30,555	52,388
Forward Citations	442,292	10.9	34.4	1	4	10

Notes: Panel A includes all utility patent applications filed 2001–2012 in Y02-technology art units (those where >10% of resolved applications receive Y02 CPC codes) with resolved disposition (granted or abandoned) and valid examiner assignment. Panel B restricts to granted applications. Examiner Grant Rate (LOO) is the leave-one-out grant rate of the assigned examiner within the art-unit \times filing-year cell. Follow-on counts measure new Y02 patents in the same CPC subclass filed within the specified window.

shows substantial within-cell variation, confirming the relevance condition for the IV design.

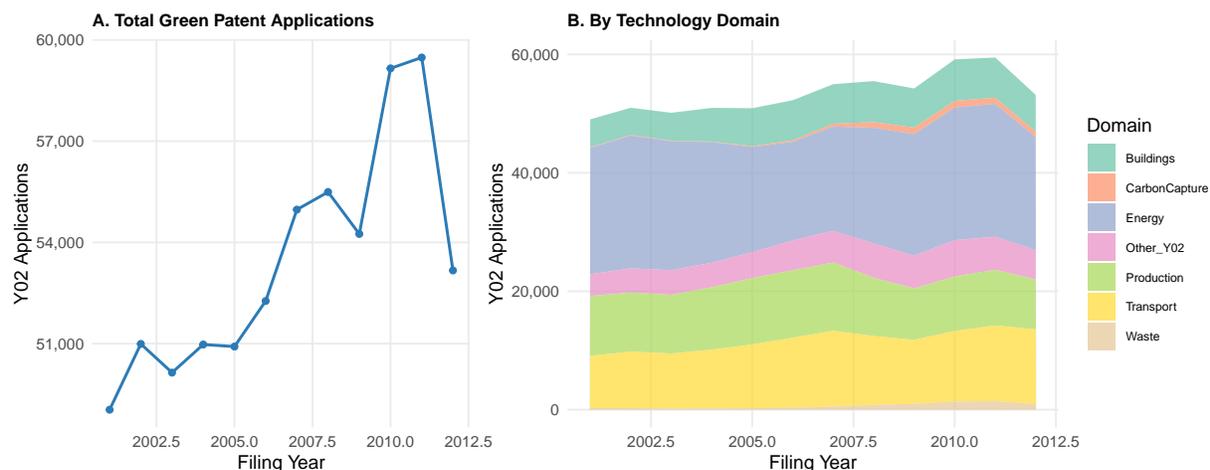


Figure 1: Y02 Green Patent Application Trends, 2001–2012

Notes: Panel A shows total Y02 patent applications by year (both granted and abandoned). Panel B shows composition across Y02 technology domains. Energy dominates the sample, followed by Transport and Production.

4.6 Sample Construction Details

The sample construction proceeds in three stages. First, I extract all utility patent applications from PatEx with resolved dispositions (ISS or ABN) filed between 2001 and 2012. Pending applications (PEND) are excluded, as they lack a known outcome. This yields approximately 3.6 million applications assigned to over 12,000 examiners across 718 art units.

Second, I identify Y02 art units by linking granted applications to PatentsView CPC data via patent number and computing the Y02 share of grants within each art unit. An art unit qualifies as a “Y02 art unit” if more than 10 percent of its granted applications carry a Y02 CPC code and at least 50 grants have Y02 classifications. This threshold balances inclusiveness with specificity—lowering it would include many art units with only incidental green technology activity, while raising it would exclude units with significant but minority green technology focus. The resulting 94 Y02 art units span all six Y02 domains.

Third, I select all applications (both granted and abandoned) in the 94 Y02 art units and impose the requirement of at least three examiners per art-unit-by-filing-year cell, yielding the final sample of 640,845 observations. The grant rate in this sample is 69.0 percent, with substantial variation across examiners: the leave-one-out examiner grant rate has a mean of 0.69 and a standard deviation of 0.20.

For abandoned applications, CPC subclass is imputed using the modal subclass among grants in the same art unit. This is necessary because PatentsView only contains CPC codes for granted patents. The imputation is conservative: it assigns all abandoned applications to the dominant technology class of their art unit, which reduces within-art-unit variation in the outcome measure and biases the estimated effect toward zero.

4.7 Statistical Power

With 640,845 observations and a standard error of 0.0018 on the baseline five-year follow-on coefficient, the 95% confidence interval of $[-0.010, -0.003]$ constrains the true effect to be small. The minimum detectable effect (at 80% power) is approximately 0.004 log points. This is well below the 20–30 percent blocking effects found by [Williams \(2013\)](#) in the genome context. However, this precision may be overstated due to the shared-outcome structure ([Section 7.5](#)).

5. Empirical Strategy

5.1 Identification

The fundamental challenge is endogeneity: applications for high-quality inventions are more likely to be granted and more likely to stimulate follow-on work. I address this using quasi-random examiner assignment as a source of exogenous variation.

Relevance. Examiners must differ in grant propensity within art-unit-by-year cells. The first stage confirms this: a one standard deviation increase in examiner grant rate raises the probability of grant by 15.1 percentage points ($F > 13,000$) without controls, or 1.8 percentage points ($F = 594$) with application controls that partly proxy grant status.

Independence. Conditional on art-unit-by-filing-year fixed effects, examiner identity must be as-good-as-random with respect to application quality and expected follow-on potential. This is supported by the institutional process (assignment based on workload and rotation) and the balance tests reported below, though the available pre-treatment covariates are limited.

Interpretation of the reduced form. The reduced-form coefficient captures the intention-to-treat effect of assignment to a more permissive examiner on cumulative downstream activity. This is a bundle: more permissive examiners not only grant more frequently but may also differ in prosecution duration, claim scope, continuation strategies, and other behaviors ([Frakes and Wasserman 2017](#)). The reduced form does not isolate the binary grant decision.

5.2 Estimation

The first stage estimates whether examiner assignment predicts the grant decision:

$$D_i = \alpha + \beta Z_i + X_i' \gamma + \delta_{a(i),t(i)} + \varepsilon_i \quad (2)$$

where D_i equals one if application i is granted and zero if abandoned, Z_i is the standardized leave-one-out examiner grant rate, X_i includes application controls ($\log(1 + \text{claims})$, $\log(1 + \text{backward citations})$), and $\delta_{a(i),t(i)}$ are art-unit-by-filing-year fixed effects.

The reduced-form equation estimates the effect of examiner assignment on innovation outcomes:

$$\log(1 + Y_i) = \alpha + \pi Z_i + X_i' \gamma + \delta_{a(i),t(i)} + \varepsilon_i \quad (3)$$

where Y_i is the follow-on innovation count or forward citation count for application i . A positive π indicates that applications assigned to more permissive examiners generate more follow-on activity; zero indicates that examiner permissiveness is irrelevant for the innovation trajectory.

The instrumental variables specification combines the first stage and reduced form:

$$\log(1 + Y_i) = \alpha + \lambda D_i + X_i' \gamma + \delta_{a(i),t(i)} + \nu_i \quad (4)$$

where D_i is instrumented by Z_i . The IV coefficient λ estimates the local average treatment effect (LATE) of the grant decision under the additional assumption that examiner leniency affects outcomes *only* through the binary grant decision. This exclusion restriction is demanding: examiners who grant more may also allow broader claims, resolve cases faster, and shape continuation behavior differently. I report IV estimates for completeness but treat the reduced form as the primary estimand.

Standard errors are clustered at the examiner level in the baseline specification. Because the follow-on patenting outcome is a CPC subclass-by-filing-year aggregate shared across many applications, I also report standard errors clustered at the CPC subclass-by-filing-year level, at the art-unit level, and from collapsed regressions at both the subclass-by-year and art-unit-by-year levels.

5.3 Threats to Validity

Random assignment. I test this with balance regressions of application characteristics on examiner grant rate conditional on fixed effects. The available pre-treatment covariates from PatEx are limited: small entity status has no within-cell variation (absorbed by art-unit-by-

year FE), so the main balance tests use claims and backward citations among the grants-only subsample.

Examiner effects beyond grant/deny. Examiners may affect outcomes through prosecution duration, claim narrowing, or applicant behavior. The reduced form captures the full bundle of examiner behaviors correlated with grant rate. The IV specification isolates the grant decision channel only under strict exclusion, which I view as demanding.

Pseudo-replication. The follow-on patenting outcome varies at the CPC subclass-by-filing-year level, but the regression is estimated at the application level (640,845 observations sharing 96 unique outcome values). This inflates the effective sample size and may overstate precision. I address this by (a) clustering at the subclass-by-year level, (b) collapsing to the subclass-by-year level, and (c) collapsing to the art-unit-by-year level.

Application controls. Claims and backward citations are available only for granted patents; they are coded as zero for abandoned applications, then transformed as $\log(1+x)$ for regression use. Because examiner grant rate predicts grant status, these controls are partly endogenous: including them absorbs grant-related variation and attenuates the first-stage coefficient from 0.151 to 0.018. I present the specification without controls as the baseline.

Many-instrument bias. With thousands of examiners, standard 2SLS may exhibit finite-sample bias. The reduced form avoids this concern.

6. Results

6.1 First Stage

Examiner assignment strongly predicts whether an application is granted. A one standard deviation increase in examiner grant rate raises the probability of grant by 15.1 percentage points ($F = 13,006$; [Table 2](#), column 1). Adding $\log(1+\text{claims})$ and $\log(1+\text{backward citations})$ as controls (column 2) reduces the coefficient to 1.8 percentage points ($F = 594$). This dramatic attenuation occurs because claims and backward citations are zero-coded for abandoned applications, making them partial proxies for grant status. The uncontrolled specification (column 1) is the appropriate first stage for the baseline reduced form, and the instrument is exceptionally strong. Column (3) adds domain-by-year fixed effects with no material change relative to column (2).

Table 2: First Stage: Examiner Grant Rate Predicts Application Granted

Dependent Variable:	granted		
	(1)	(2)	(3)
Model:	(1)	(2)	(3)
<i>Variables</i>			
Examiner Grant Rate (1 SD)	0.1514*** (0.0013)	0.0182*** (0.0007)	0.0182*** (0.0007)
<i>Fixed-effects</i>			
Art-Unit \times Filing-Year	Yes	Yes	Yes
Domain \times Filing-Year			Yes
<i>Fit statistics</i>			
Observations	640,845	640,845	640,845
R ²	0.14244	0.85823	0.85829

Clustered (examiner_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Dependent variable: application granted (1 = ISS, 0 = ABN). Examiner grant rate is the leave-one-out grant rate of the assigned examiner within the art-unit \times filing-year cell. Column (1): art-unit \times filing-year FE only. Column (2): adds $\log(1 + \text{claims})$ and $\log(1 + \text{backward citations})$. Column (3): adds Y02 domain \times filing-year FE. Standard errors clustered at the examiner level.

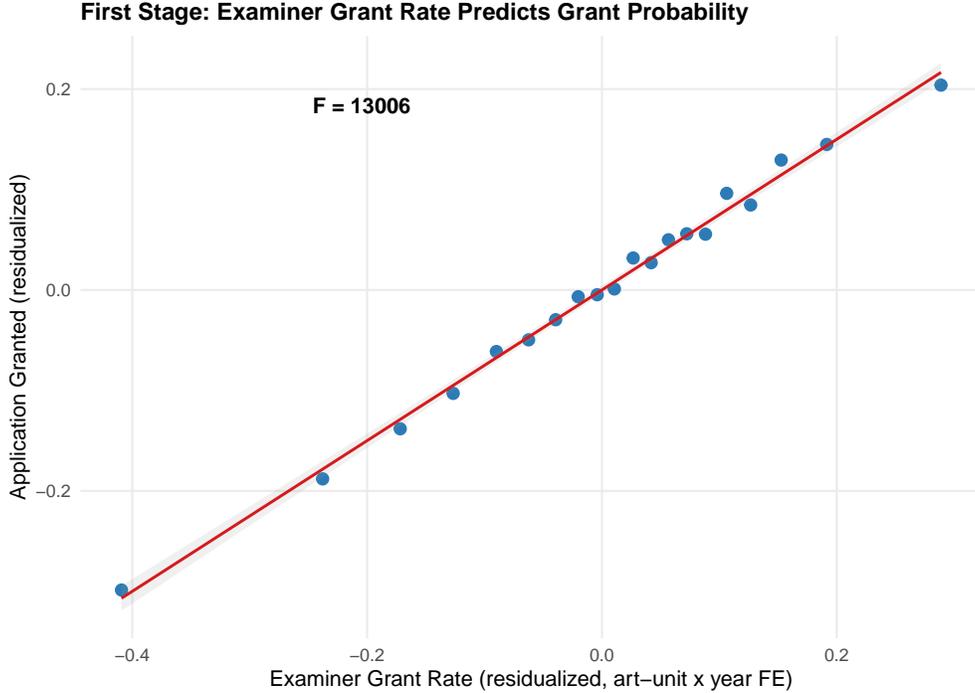


Figure 2: First Stage: Examiner Grant Rate Predicts Grant Probability

Notes: Binscatter of residualized grant probability against residualized examiner grant rate, after absorbing art-unit \times filing-year fixed effects. Each point represents 5% of the data. The fitted line corresponds to the first-stage regression.

6.2 Balance Tests

Table 3 presents balance tests. Because claims and backward citations are zero-coded for abandoned applications, running balance tests on the full sample would produce a mechanical positive correlation with examiner grant rate. Instead, I use the grants-only subsample, which conditions on a post-treatment outcome—a limitation I acknowledge. Under quasi-random assignment, examiner grant rate should not predict application characteristics among grants, conditional on cell fixed effects. The coefficient on log backward citations is small and insignificant (-0.008 , $p = 0.28$). The coefficient on log claims is small in magnitude (0.041) but statistically significant, suggesting mild sorting on claim complexity within art-unit-by-year cells. This is a common finding in examiner IV designs (Lemley and Sampat 2012).

The balance evidence is limited by the available covariates. Small entity status—the only pre-treatment application characteristic observable for all applications in PatEx—has no within-cell variation and is absorbed by art-unit-by-filing-year fixed effects. Stronger validation would require application-level covariates such as continuation status, applicant

Table 3: Balance Test: Application Characteristics on Examiner Grant Rate (Grants Only)

Dependent Variables: Model:	log_claims (1)	log_bwd_cite (2)
<i>Variables</i>		
Examiner Grant Rate (1 SD)	0.0408*** (0.0029)	-0.0082 (0.0076)
<i>Fixed-effects</i>		
Art-Unit \times Filing-Year	Yes	Yes
<i>Fit statistics</i>		
Observations	442,287	442,287
R ²	0.03063	0.10721

Clustered (examiner_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Grants-only subsample (442,292 granted applications; 5 singleton observations removed by fixed effects). Claims and citations are only observed for granted patents; variables are $\log(1 + x)$. Standard errors clustered at the examiner level. Under random assignment, coefficients should be zero.

origin, or inventor count, which are not available in the current extract. The quasi-random assignment assumption rests primarily on the institutional process (workload-based routing) rather than a comprehensive empirical demonstration.

6.3 Reduced Form: Follow-on Innovation

Assignment to a more permissive examiner slightly reduces follow-on innovation in the application-level regression. A one standard deviation increase in examiner grant rate is associated with a 0.67 percent reduction in five-year follow-on Y02 patenting in the same CPC subclass (-0.0067 , $SE = 0.0018$, $p < 0.001$; Table 4, column 1, preferred baseline without controls). Adding zero-coded application controls (column 2) attenuates the coefficient to -0.0040 ($SE = 0.0017$, $p = 0.018$), as these controls partly absorb grant-status variation. However, these application-level p-values should be interpreted cautiously: the follow-on outcome is a CPC subclass-by-filing-year aggregate (640,845 observations sharing 96 unique values), so application-level standard errors overstate precision.

A critical caveat: the follow-on outcome is an aggregate measured at the CPC subclass-by-filing-year level—640,845 applications share only 96 unique outcome values. The high R^2 values in columns (1)–(2) (0.87) reflect this shared-outcome structure: art-unit-by-filing-year fixed effects map closely to the 96 subclass-by-year cells, absorbing most outcome

Table 4: Effect of Examiner Grant Rate on Follow-on Innovation and Citations

Dependent Variables:	log_followon_5yr Follow-on 5yr			log_fwd_citations Forward Citations		
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
Examiner Grant Rate (1 SD)	-0.0067*** (0.0018)	-0.0040** (0.0017)		0.2524*** (0.0048)	0.0205*** (0.0036)	
log(1 + Claims)		-0.0024*** (0.0009)	0.0568** (0.0252)		0.3281*** (0.0032)	0.0223 (0.0560)
log(1 + Bwd. Citations)		-0.0056*** (0.0010)	0.0065 (0.0053)		0.2893*** (0.0039)	0.2268*** (0.0121)
Granted (IV)			-0.2181** (0.0937)			1.127*** (0.2058)
<i>Fixed-effects</i>						
Art-Unit × Filing-Year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	640,845	640,845	640,845	640,845	640,845	640,845
R ²	0.87024	0.87042	0.87042	0.14608	0.47524	0.47524
F-test (1st stage), granted			3,887.8			3,887.8

Clustered (examiner_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Columns (1)–(2) and (4)–(5): reduced-form effect of examiner grant rate. Column (1)/(4): baseline without application controls (preferred). Column (2)/(5): adds log(1 + claims) and log(1 + backward citations) as controls. Note: claims and citations are coded as zero for abandoned applications (unavailable in PatentsView), then log-transformed; these controls therefore partly proxy grant status. Columns (3) and (6): exploratory IV estimates (exclusion restriction demanding; not causally interpretable as grant effects). The IV first-stage F (Wald statistic from joint IV estimation) differs from the standalone F in Table 2 due to the inclusion of endogenous regressors in the joint specification. All specifications include art-unit × filing-year FE. Standard errors clustered at the examiner level. The high R^2 in follow-on columns reflects the shared-outcome structure: follow-on patenting is a CPC subclass × filing-year aggregate (96 unique values), so AU × FY fixed effects absorb nearly all outcome variation.

variation mechanically rather than through model fit. The application-level standard errors may therefore overstate precision. I address the pseudo-replication concern directly in the aggregation analysis below.

6.4 Forward Citations (Mechanical Contamination)

Forward citations (columns 4–6) show a large positive response. The baseline reduced form (column 4) yields a coefficient of 0.252 ($p < 0.001$). This large effect is almost entirely mechanical: abandoned applications cannot receive citations because they do not exist as citable documents in the patent database. After adding controls that partly absorb grant-status variation (column 5), the coefficient falls to 0.021 ($p < 0.001$). Even this residual effect likely conflates genuine knowledge transfer with database visibility. I report these results for transparency but do not interpret them as evidence of a disclosure channel.

6.5 Exploratory IV Estimates

The IV estimates in columns (3) and (6) mechanically scale the reduced form by the first stage. I include them for completeness but emphasize that they should *not* be interpreted as causal effects of the grant decision. The IV implies that a marginal grant reduces same-subclass follow-on patenting by 0.22 log points ($p = 0.02$) while increasing forward citations by 1.13 log points ($p < 0.001$). The first-stage F-statistic reported in the IV columns ($F = 3,888$) is the Wald F from the `fixest` IV specification, which is computed differently from the standalone first-stage F in [Table 2](#) ($F = 13,006$ without controls, $F = 594$ with controls) due to the joint estimation of endogenous and exogenous regressors. These estimates require the demanding exclusion restriction that examiner permissiveness affects outcomes *only* through the binary grant decision—ruling out effects through prosecution duration, claim scope, continuation strategies, and claim narrowing ([Frakes and Wasserman, 2017](#)). Given strong prior reasons to doubt this restriction, the IV magnitudes are not policy-relevant in their current form.

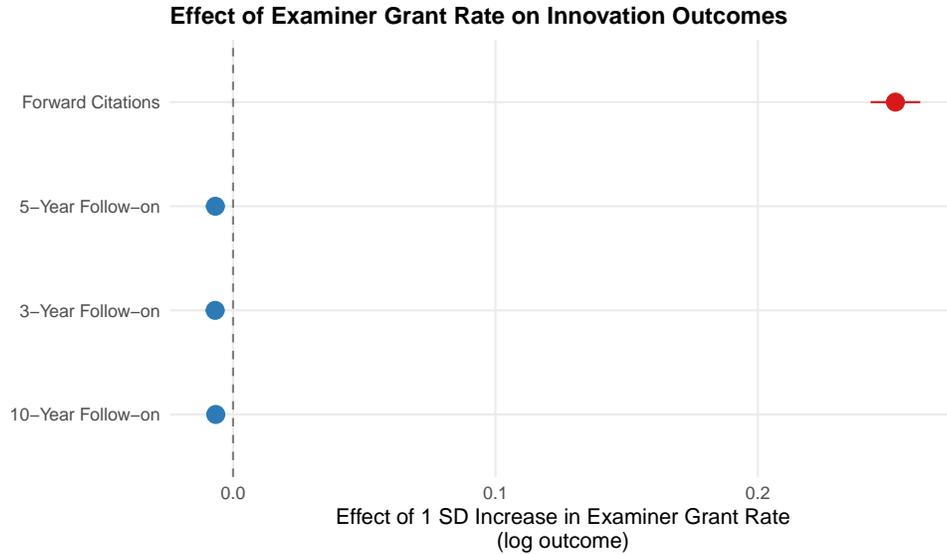


Figure 3: Effect of Examiner Grant Rate on Innovation Outcomes

Notes: Each point represents the coefficient on standardized examiner grant rate from a separate regression with art-unit-by-filing-year fixed effects and no application controls (baseline specification). Blue points: follow-on patenting outcomes. Red point: forward citations (mechanically contaminated). Lines show 95% confidence intervals.

6.6 Heterogeneity by Technology Domain

The impact of a patent grant may differ for a solar panel versus a carbon scrubber. [Table 5](#) reports heterogeneity across Y02 technology domains using forward citations as the outcome (the follow-on patenting outcome is a CPC-subclass-level aggregate that lacks within-cell variation in domain subsamples). The citation effect is positive across most domains, with the strongest effects in Transport (0.048) and Carbon Capture (0.080). Given the mechanical citation contamination discussed above, this heterogeneity should be interpreted cautiously—it may reflect differences in the propensity to cite recently-granted patents across technology fields rather than differences in knowledge transfer.

Table 5: Heterogeneity by Y02 Technology Domain: Forward Citations

Dependent Variable:	log_fwd_citations						
Model:	Buildings (1)	CCS (2)	Energy (3)	Other (4)	Production (5)	Transport (6)	Waste (7)
<i>Variables</i>							
Grant Rate (1 SD)	-0.0038 (0.0096)	0.0795** (0.0325)	0.0140*** (0.0046)	0.0307*** (0.0086)	0.0196*** (0.0068)	0.0476*** (0.0106)	-0.0009 (0.0198)
<i>Fixed-effects</i>							
Art-Unit × Filing-Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>							
Observations	72,906	6,328	241,847	58,693	120,123	132,687	7,462
R ²	0.49888	0.44259	0.48951	0.41777	0.44540	0.48221	0.49357

Clustered (*examiner_id*) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Dep. var: $\log(1 + \text{forward citations})$. CCS = Carbon Capture Storage. All columns include AU × FY FE and controls ($\log(1 + \text{claims})$, $\log(1 + \text{backward citations})$). SEs clustered at examiner level.

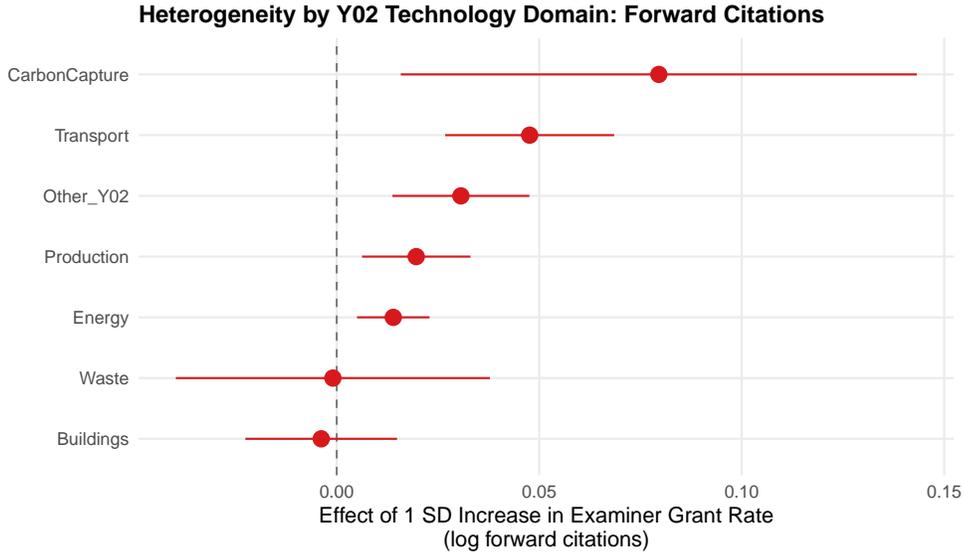


Figure 4: Heterogeneity by Y02 Technology Domain: Forward Citations

Notes: Coefficient on standardized examiner grant rate from domain-specific regressions of $\log(1 + \text{fwd. cit.})$ on grant rate with AU×FY FE and controls. SEs clustered at examiner level. 95% CIs shown. See Table 5.

Table 6: Robustness Checks, Falsification, and Aggregation

Specification	Coefficient	Std. Error	N	FE
<i>Panel A: Robustness</i>				
Poisson	-0.0025*	(0.0014)	640,845	AU×Y
Large cells	-0.0040**	(0.0017)	640,017	AU×Y
Winsorized	-0.0040**	(0.0017)	640,845	AU×Y
Experienced examiners	-0.0041**	(0.0017)	639,664	AU×Y
3yr horizon	-0.0040**	(0.0017)	640,845	AU×Y
10yr horizon	-0.0038**	(0.0017)	640,845	AU×Y
Grants only	-0.0075***	(0.0025)	442,287	AU×Y
<i>Panel B: Falsification</i>				
Placebo other subclass	0.001025***	(0.000310)	640,845	AU×Y
<i>Panel C: Collapsed Aggregation</i>				
Collapsed subclass×year	-0.1930**	(0.0847)	96	Year
Collapsed subclass×year	-0.1111***	(0.0352)	96	Sub+Year
Collapsed subclass×year	-0.2052***	(0.0597)	96	Sub+Year+Ctrl
Collapsed AU no controls	-0.0110	(0.0142)	980	AU+Y
Collapsed AU controls	0.0091	(0.0207)	980	AU+Y+Ctrl

Notes: Panel A reports the coefficient on examiner grant rate (1 SD) from variants of the main specification with application controls ($\log(1 + \text{claims})$, $\log(1 + \text{backward citations})$) and art-unit \times filing-year fixed effects. Panel B: placebo outcome is $\log(1 + \text{follow-on in other CPC subclasses})$, baseline specification without controls. The placebo outcome varies at the subclass \times year level; the small SE reflects limited residual variation after FE. Panel C collapses the data to the level at which the follow-on outcome varies. “Subclass×year” uses 96 CPC subclass-by-filing-year cells, weighted by cell size. “AU” uses art-unit-by-filing-year cells with AU + year FE. Dependent variable (except Panel B): $\log(1 + \text{follow-on Y02 patents, 5yr})$. Standard errors clustered at the examiner level (Panels A–B) or robust (Panel C). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

7. Robustness and Falsification

7.1 Alternative Specifications

Table 6 presents robustness checks (Panel A) and falsification tests (Panel B). The negative reduced-form effect on five-year follow-on is robust to:

- **Poisson model:** accommodating the count nature of the outcome
- **Large cells only** ($N \geq 20$): ensuring adequate within-cell variation
- **Winsorized at 99th percentile:** limiting the influence of extreme values
- **Experienced examiners** (≥ 10 applications): reducing measurement error in the instrument
- **Alternative horizons:** 3-year and 10-year follow-on windows
- **Grants-only subsample:** restricting to granted applications only

7.2 Inference Under Alternative Clustering

Table 7: Inference Under Alternative Clustering

Clustering Level	Coefficient	Std. Error	N
Examiner	-0.0067***	(0.0018)	640,845
Two-way (Exam x CPC)	-0.0067	(0.0057)	640,845
Art Unit	-0.0067**	(0.0027)	640,845
CPC Subclass x Year	-0.0067***	(0.0020)	640,845

Notes: Outcome is $\log(1 + \text{follow-on Y02 patents, 5yr})$. All rows use the baseline specification without application controls, with art-unit \times filing-year FE. CPC Subclass \times Year clustering addresses the shared-outcome concern that many applications share the same subclass-level follow-on count.

Table 7 reports standard errors under four clustering schemes for the baseline specification without application controls. Examiner-level clustering (SE = 0.0018) is the standard choice for examiner-IV designs, yielding $t = 3.7$. CPC subclass-by-filing-year clustering—which directly addresses the shared-outcome concern—yields SE = 0.0020 ($t = 3.4$). Art-unit clustering gives SE = 0.0027 ($t = 2.5$). Two-way clustering at the examiner-by-CPC-subclass level yields the most conservative standard error (SE = 0.0057, $t = 1.2$), rendering the coefficient statistically insignificant. The baseline coefficient is negative under all schemes, but its statistical significance depends substantially on the clustering choice.

7.3 Falsification: Placebo Outcomes

I test whether examiner grant rate predicts follow-on patenting in *different* CPC subclasses. For each application, other-subclass follow-on is the total Y02 follow-on count across all other CPC subclasses in the same filing year. The coefficient is small and positive (0.0010, SE = 0.0003; Panel B of Table 6)—opposite in sign to the main negative effect and roughly one-seventh its magnitude. This pattern is mechanically expected: within a filing year, own-subclass and other-subclass follow-on are complements (they sum to the year total, which is absorbed by FE), so a negative own-subclass effect implies a positive other-subclass coefficient. The small magnitude relative to the main effect is reassuring, but a strict null was never the appropriate benchmark for this test.

7.4 Permutation Inference

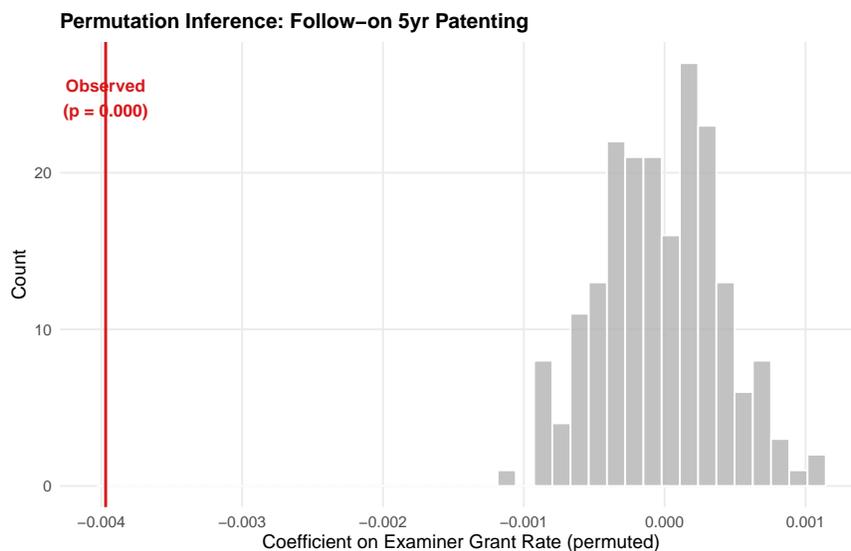


Figure 5: Permutation Inference: Follow-on 5yr Patenting

Notes: Distribution of the reduced-form coefficient under 200 random permutations of examiner grant rate within art-unit \times filing-year cells. The red line marks the observed coefficient. The permutation p-value confirms the regression inference.

Figure 5 shows the permutation distribution for the five-year follow-on coefficient. Randomly shuffling examiner grant rates within art-unit-by-filing-year cells produces a distribution centered near zero with a standard deviation of 0.0004. The observed coefficient of -0.004 falls far outside this distribution (permutation $p < 0.005$). This confirms that the negative reduced-form association is not an artifact of the regression specification or the shared-outcome structure at the application level. However, the permutation preserves the application-level

data structure rather than aggregating to the outcome level, so it does not fully address the pseudo-replication concern.

7.5 Aggregation Analysis

The most important robustness concern is that the 640,845 application-level observations share only 96 unique follow-on outcome values (8 CPC subclasses \times 12 filing years). Collapsing to the level at which the outcome genuinely varies is essential.

CPC subclass \times filing year (96 cells). Collapsing to this level—using the same $\log(1 + \text{follow-on count})$ dependent variable and weighting by cell size—the coefficient on mean examiner grant rate is -0.193 (SE = 0.085, $p = 0.025$) with year fixed effects, and -0.111 (SE = 0.035, $p = 0.002$) with subclass and year fixed effects (Panel C of [Table 6](#)). The negative sign is preserved and statistically significant. However, this collapsed design does not automatically inherit the application-level random-assignment guarantee. Examiner assignment is quasi-random within art-unit-by-year cells, not within subclass-by-year cells. The collapsed regression identifies the effect of subclass-year average examiner permissiveness, which could be confounded by compositional differences in application mix across subclass-year cells, time-varying art-unit composition, and the subclass imputation for abandoned applications.

Art-unit \times filing year (980 cells). Collapsing to this level—which better matches the assignment structure, since examiner assignment is randomized within art-unit-by-year—yields a null. The coefficient is -0.011 (SE = 0.014) without controls and 0.009 (SE = 0.021) with controls (Panel C of [Table 6](#)). Art units contain an average of 5.0 CPC subclasses, so the art-unit-level collapse mixes across subclasses, potentially diluting technology-specific effects. But the null at this level—which most closely matches the assignment mechanism—is also consistent with the interpretation that the subclass-level result is spurious.

The discrepancy between the two collapses is the paper’s central empirical finding. The subclass-level collapse produces a negative result where the outcome genuinely varies but where identification is less clean; the art-unit-level collapse preserves the assignment structure but yields a null. Neither collapse resolves the fundamental tension between outcome-level and assignment-level aggregation. I present both transparently and conclude that the evidence does not support a firm causal claim about blocking.

7.6 Temporal Evolution

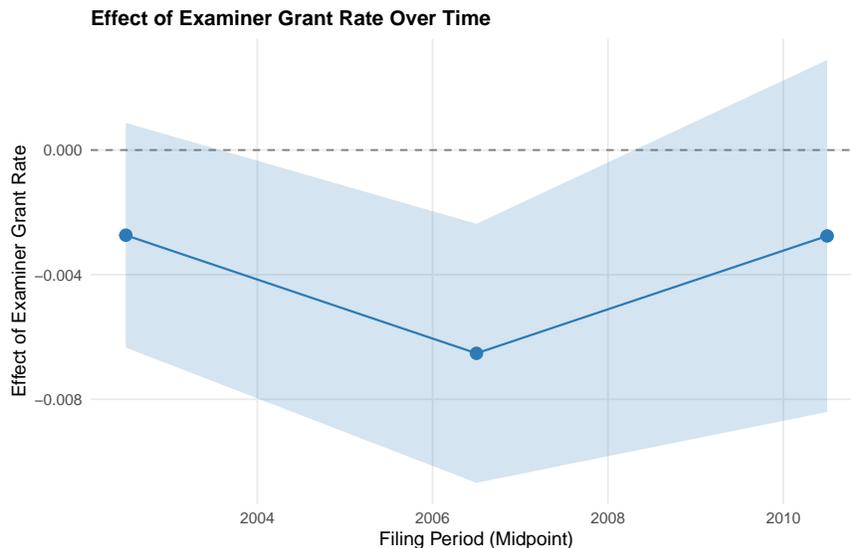


Figure 6: Effect of Examiner Grant Rate Over Time

Notes: Each point represents the coefficient on standardized examiner grant rate from a separate regression estimated on the indicated filing-era subsample, with application controls and art-unit-by-filing-year FE. Three eras: 2001–2004, 2005–2008, 2009–2012. Points are plotted at era midpoints and connected for visual continuity. Shaded area shows 95% confidence band.

Figure 6 plots the coefficient by filing era. The coefficient is -0.003 in 2001–2004, -0.007 in 2005–2008, and -0.003 in 2009–2012—all negative with no clear temporal trend.

8. Discussion

8.1 What the Evidence Shows

The paper cleanly establishes the first stage: application-level PatEx data with both grants and abandonments enables a proper examiner IV design for green technology ($F > 13,000$). The downstream evidence is mixed. The collapsed subclass-by-year analysis (96 cells) shows a negative association between average examiner permissiveness and follow-on Y02 patenting (-0.193 , $p = 0.025$), but this collapsed design does not inherit the application-level random-assignment guarantee. The art-unit-by-year collapse (980 cells)—which better matches the assignment structure—yields a null. Application-level estimates, while significant under examiner clustering, are based on pseudo-replicated outcomes and overstate precision.

Several mechanisms could produce a negative reduced-form effect.

Blocking. A marginal grant creates a property right that subsequent inventors must design around, license, or risk infringing. In green technology, where patent thickets are a documented concern, the marginal patent may deter entry into the same CPC subclass. The negative point estimate at the subclass-year level is consistent with this channel, but the null at the art-unit-year level makes this interpretation uncertain.

Cross-subclass redirection. The negative within-subclass effect could coexist with positive effects in adjacent subclasses if the marginal grant redirects inventive effort rather than reducing it. Green technologies are deeply interconnected—a battery patent (Y02E) may stimulate work in electric vehicles (Y02T) or grid management. The within-subclass measure would capture redirection as a negative effect.

Irrelevance at the margin. The effect, even where statistically detectable, is small. Green innovation may be driven primarily by market forces (energy prices, deployment subsidies), regulatory requirements, and public R&D funding ([Acemoglu et al. 2012](#); [Popp 2002](#))—none of which depend on whether a specific application was granted.

8.2 What the Evidence Does Not Show

The forward-citation result is substantially contaminated by the mechanical channel: abandoned applications cannot be cited because they do not exist as citable documents. I cannot separate genuine knowledge transfer from database visibility with the available data. Decomposing applicant-initiated versus examiner-initiated citations, or examining academic citations to patent disclosures, could shed light on this channel but is beyond the scope of the current study.

The IV estimates (-0.22 log points on follow-on, $+1.13$ on citations) require the strong exclusion restriction that examiner permissiveness affects outcomes only through the binary grant decision. Given that examiners who grant more frequently likely also differ in prosecution duration, claim scope, and continuation strategies ([Frakes and Wasserman, 2017](#)), these estimates should not be interpreted as causal effects of the grant decision and are not policy-relevant in their current form.

8.3 Comparison with Prior Literature

The mixed downstream evidence is most consistent with [Sampat and Williams \(2019\)](#), who found no effect of gene patent grants on follow-on research using the same examiner IV design. In both contexts—genomics and green technology—cumulative innovation appears largely independent of the marginal patent decision. The result contrasts with [Williams \(2013\)](#), who found 20–30 percent reductions from Celera Genomics’ IP claims—a regime-level

intervention qualitatively different from the marginal examiner effect identified here. Galasso and Schankerman (2015) found that court invalidation increased subsequent citations, but court invalidation publicly removes an existing IP barrier, while examiner denial simply prevents one from being created. Unlike these prior studies, which used application- or patent-specific outcomes, this paper’s follow-on measure is a coarse subclass-level aggregate that may not capture the technologically proximate innovation channel central to the blocking hypothesis.

8.4 Policy Implications

The results are informative but circumscribed. The downstream evidence is too sensitive to aggregation level, clustering, and outcome definition to support firm causal conclusions about blocking. The binding constraints on green innovation appear to lie in research funding, regulatory certainty, and market demand rather than in individual patent decisions. The IPCC’s call for tripling clean energy investment (IPCC, 2023) targets resources and incentives, not IP reform.

This design identifies a local margin, not broad IP regime changes. The infra-marginal incentive effect—the expectation of patent protection motivating R&D investment—is not identified. A world without green patents might feature dramatically less green R&D, even though the marginal grant decision has small effects on within-subclass follow-on activity.

8.5 Limitations

Several limitations warrant emphasis.

First, **pseudo-replication and aggregation sensitivity**. The primary follow-on outcome varies at the CPC subclass-by-filing-year level (96 unique values), yet the regression is estimated on 640,845 application-level observations. The subclass-level collapse preserves the negative sign; the art-unit-level collapse does not. This sensitivity is the paper’s most important limitation and prevents a definitive conclusion about the sign or magnitude of the effect.

Second, **limited balance evidence**. The quasi-random assignment assumption rests primarily on institutional arguments rather than comprehensive empirical validation. Small entity status—the only pre-treatment covariate observable for all applications—has no within-cell variation. The grants-only balance test conditions on a post-treatment outcome and is therefore not fully informative. Richer application-level covariates from PatEx (continuation status, applicant origin, inventor count) could strengthen the design but are not in the current extract.

Third, **Y02 classification for abandonments**. CPC subclass is imputed for abandoned applications using the art-unit modal subclass among grants. Art units contain an average of 5.0 subclasses, so some imputation error is inevitable. This may be classical measurement error (biasing toward zero) or systematic if the technology composition of abandonments differs from grants within an art unit.

Fourth, **treatment timing**. The outcome window begins at filing date, not at disposition, so some measured follow-on activity predates the grant decision. The 5-year and 10-year windows extend well beyond the typical 2–3 year examination period, but the reduced form is more precisely an ITT effect of examiner assignment from filing onward than a post-grant effect.

Fifth, **local estimand and sample period**. The design identifies the effect of assignment to a more permissive examiner for marginal applications in 2001–2012. The green innovation landscape has evolved substantially since then (Paris Agreement 2015, IRA 2022), and the marginal effect of patent decisions may differ in a more mature technology environment.

9. Conclusion

Does assignment to a more permissive patent examiner affect the pace of cumulative green innovation? This paper’s primary contribution is methodological: application-level PatEx data with both grants and abandonments enables a proper first stage in the examiner IV design for green technology ($F > 13,000$), resolving the selection problem inherent in grants-only analyses.

The downstream evidence is mixed and does not support a definitive causal conclusion. The available follow-on patenting outcome varies at the CPC subclass-by-filing-year level (96 cells), creating a mismatch with the application-level assignment design. Collapsing to the subclass-year level produces a significant negative association ($p = 0.025$); collapsing to the art-unit-year level—which better preserves the assignment structure—yields a null. Forward citations respond positively, but this is largely mechanical (abandoned applications cannot be cited).

The findings parallel [Sampat and Williams \(2019\)](#)’s result for gene patents: cumulative innovation in green technology appears largely independent of the marginal patent decision. A definitive answer to the blocking question will require downstream outcomes measured at the application level—citation-based follow-on, text-similarity-weighted subsequent inventions, or applicant-level innovation—rather than the coarse subclass-level aggregates available here. The forces shaping green innovation trajectories—carbon prices, R&D subsidies, deployment mandates, and market demand—appear to dominate the margin identified here.

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

- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous**, “The Environment and Directed Technical Change,” *American Economic Review*, 2012, *102* (1), 131–166.
- Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen**, “Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry,” *Journal of Political Economy*, 2016, *124* (1), 1–51.
- Budish, Eric, Benjamin N. Roin, and Heidi Williams**, “Do Firms Underinvest in Long-Run Research? Evidence from Cancer Clinical Trials,” *American Economic Review*, 2015, *105* (7), 2044–2085.
- Calel, Raphael and Antoine Dechezleprêtre**, “Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market,” *Review of Economics and Statistics*, 2016, *98* (1), 173–191.
- Chyn, Eric, Brigham Frandsen, and Emily Leslie**, “Examiner and Judge Designs,” *Working Paper*, 2024.
- Cockburn, Iain M., Samuel Kortum, and Scott Stern**, “Are All Patent Examiners Equal? The Impact of Characteristics on Patent Statistics and Litigation Outcomes,” *NBER Working Paper*, 2003.
- Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist**, “What Is a Patent Worth? Evidence from the US 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.
- Green, Jerry R. and Suzanne Scotchmer**, “On the Division of Profit in Sequential Innovation,” *RAND Journal of Economics*, 1995, *26* (1), 20–33.
- Heller, Michael A. and Rebecca S. Eisenberg**, “Can Patents Deter Innovation? The Anticommons in Biomedical Research,” *Science*, 1998, *280* (5364), 698–701.

- IPCC**, “Climate Change 2023: Synthesis Report,” Technical Report, Intergovernmental Panel on Climate Change 2023.
- Johnstone, Nick, Ivan Haščič, and David Popp**, “Renewable Energy Policies and Technological Innovation: Evidence Based on Patent Counts,” *Environmental and Resource Economics*, 2010, *45* (1), 133–155.
- Lemley, Mark A. and Bhaven Sampat**, “Examiner Characteristics and Patent Office Outcomes,” *Review of Economics and Statistics*, 2012, *94* (3), 817–827.
- Murray, Fiona, Philippe Aghion, Mathias Dewatripont, Julian Kolev, and Scott Stern**, “Of Mice and Academics: Examining the Effect of Openness on Innovation,” *Working Paper, NBER*, 2007.
- Nordhaus, William D.**, “An Economic Theory of Technological Change,” *American Economic Review*, 1969, *59* (2), 18–28.
- Popp, David**, “Induced Innovation and Energy Prices,” *American Economic Review*, 2002, *92* (1), 160–180.
- 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.
- Williams, Heidi L.**, “Intellectual Property Rights and Innovation: Evidence from the Human Genome,” *Journal of Political Economy*, 2013, *121* (1), 1–27.

A. Data Appendix

A.1 Data Sources

- **Patent Examination Research Dataset (PatEx):** Application-level data from Google BigQuery (`patents-public-data.uspto_oce_pair.application_data`). Contains 9.8 million utility patent applications with examiner IDs, filing dates, USPC classes, and disposition outcomes (granted/ISS, abandoned/ABN, pending/PEND). Filtered to resolved applications filed 2001–2012 with valid examiner assignments.
- **PatentsView CPC data** (`g_cpc_current.tsv`): CPC classification codes including Y02 green technology tags for all granted US patents.
- **PatentsView citations** (`g_us_patent_citation.tsv`): Complete forward and backward citation network among granted patents.
- **PatentsView patent data** (`g_patent.tsv`): Patent metadata including grant dates, claim counts, and patent type.

A.2 Sample Construction

Starting from 3.6 million resolved utility patent applications filed 2001–2012 with valid examiner assignments:

1. Identify Y02 art units: art units where $>10\%$ of granted applications carry Y02 CPC codes (minimum 50 Y02 grants)
2. Select all applications (granted + abandoned) in Y02 art units
3. Require at least 3 examiners per art-unit \times filing-year cell
4. Drop applications without valid CPC subclass (after imputation from art-unit modal subclass)

A.3 Variable Definitions

Examiner Grant Rate (LOO): Leave-one-out proportion of granted applications among the assigned examiner’s caseload in the art-unit \times filing-year cell: $(G_j - D_i)/(N_j - 1)$, where G_j is examiner j ’s total grants, D_i is the focal application’s grant status, and N_j is the examiner’s total applications in the cell. Standardized to mean zero, unit standard deviation.

Application Granted: Binary indicator: 1 if disposition is ISS (granted), 0 if ABN (abandoned).

Follow-on Y02 Patents (K-year): Count of Y02 patents filed within K years after the focal application's filing date in the same CPC subclass.

Forward Citations: Count of subsequent patents citing the focal patent (zero for abandoned applications).

B. Identification Appendix

B.1 Examiner Grant Rate Distribution

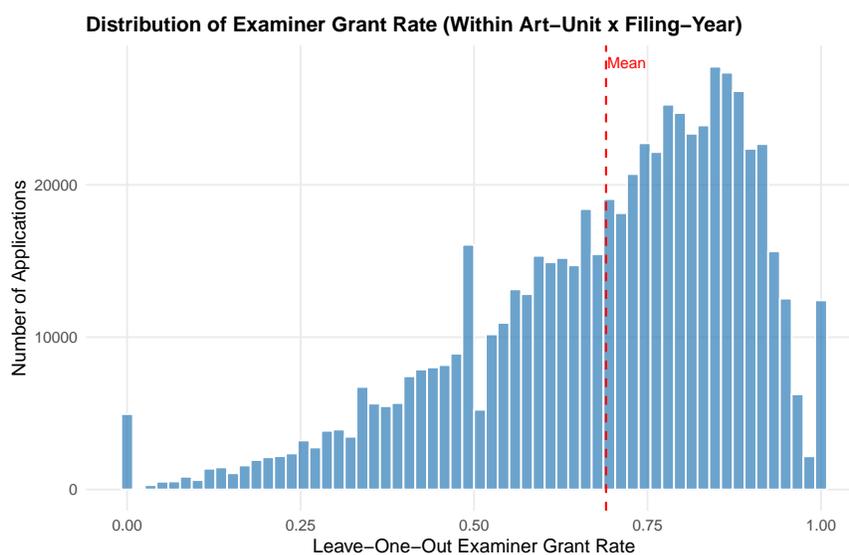


Figure 7: Distribution of Leave-One-Out Examiner Grant Rate

Notes: Distribution of leave-one-out examiner grant rate across all applications in Y02 art units. The dashed red line marks the mean.

B.2 Balance Test Detail

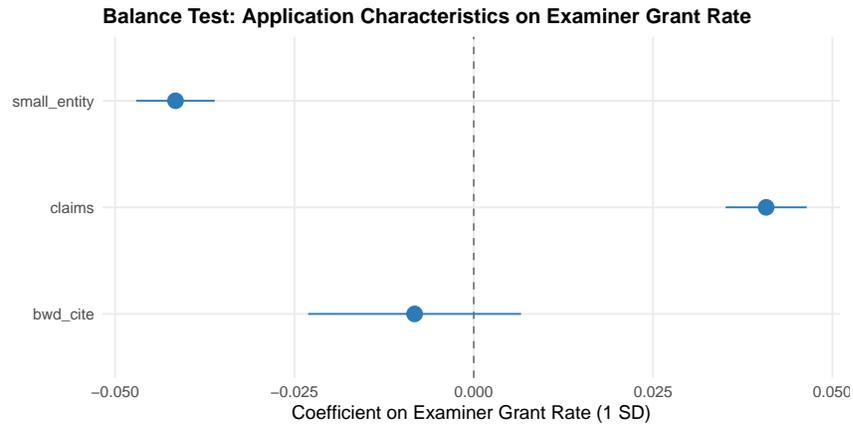


Figure 8: Balance Test: Application Characteristics on Examiner Grant Rate

Notes: Coefficients from regressions of application characteristics on standardized examiner grant rate with art-unit-by-filing-year fixed effects. Lines show 95% CIs.