

The Examination Lottery: Measuring Regulatory Inconsistency with Patent Continuation Twin Studies

APEP Autonomous Research* @olafdrw

March 29, 2026

Abstract

Patent examination is supposed to be consistent: the same invention should receive the same verdict regardless of which examiner reviews it. I test this assumption using 97,922 continuation-application twin pairs where the same invention family was reassigned to a different examiner within the same art unit at the USPTO. Reassignment increases outcome discordance—one application granted, the other abandoned—by 7.7 percentage points relative to same-examiner pairs (11.0 pp raw differential). A one-standard-deviation increase in examiner leniency raises the grant probability by 11.4 percentage points ($t = 33.8$). Small entities face 3.2 percentage points greater excess discordance than large entities. These results reveal a regulatory lottery in which examiner identity causally determines patent rights, with disproportionate consequences for resource-constrained inventors.

JEL Codes: O34, L51, K11

Keywords: patent examination, regulatory inconsistency, examiner leniency, continuation applications, small entities

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 22m).

1. Introduction

Nearly 30 percent of patent continuation applications receive a different verdict when reviewed by a different examiner in the same technology group. The same claims, the same prior art, the same legal standards—yet whether an invention becomes a patent depends substantially on which examiner happens to be available when the application arrives. This paper provides the first direct measurement of this examination lottery using a within-family twin study design that holds the invention family fixed.

The consistency of patent examination is a first-order policy question. Patents confer exclusive rights worth billions of dollars in aggregate, shaping innovation incentives, market structure, and technology diffusion (Jaffe and Lerner, 2004). If examination outcomes vary arbitrarily with examiner identity, then the patent system functions partly as a lottery—randomly allocating property rights in ways that distort resource allocation. Prior work has documented examiner heterogeneity using cross-sectional variation across different inventions (Lemley, 2001; Cockburn et al., 2003; Lichtman and Lemley, 2004), but comparing different inventions across examiners confounds examiner effects with invention quality. The central challenge is isolating the causal contribution of examiner identity while holding the underlying invention constant.

I exploit the institutional structure of continuation applications at the United States Patent and Trademark Office (USPTO). When inventors file continuation or divisional applications, the new application may be assigned to a different examiner than the parent—even within the same art unit. This reassignment is driven by workload balancing and examiner availability, creating quasi-random variation in examiner identity conditional on art-unit-by-filing-year fixed effects. Crucially, because parent and child applications belong to the same invention family, comparing their outcomes holds invention quality approximately fixed, yielding a twin study design with 702,490 same-art-unit pairs.

The results are striking. Among 97,922 pairs where the child was reassigned to a different examiner, 29.8 percent produced discordant outcomes (one granted, the other abandoned), compared to 18.7 percent for same-examiner pairs—an excess discordance of 11.0 percentage points. In regressions with art-unit-by-filing-year fixed effects, reassignment increases discordance by 7.7 percentage points ($t = 21.1$). The examiner’s leave-one-out grant rate is a powerful predictor of the child’s outcome: a one-standard-deviation increase in leniency raises the grant probability by 11.4 percentage points, with a first-stage F -statistic of 1,158. These magnitudes imply that examiner identity explains a substantial share of outcome variation beyond what the invention itself warrants.

The burden of this regulatory lottery falls disproportionately on small entities. Among

reassigned pairs, small-entity applications exhibit 38.4 percent discordance compared to 27.2 percent for large entities—a gap of 11.2 percentage points. In a triple-difference specification, the small-entity-by-reassignment interaction is 3.2 percentage points ($t = 4.3$). The leniency effect is also larger for small entities (0.692 versus 0.626), consistent with small inventors having fewer resources to appeal unfavorable decisions or navigate examiner-specific preferences.

This paper contributes to three literatures. First, it advances the measurement of examiner effects in patent examination. [Sampat and Williams \(2019\)](#) show that examiner leniency affects downstream innovation outcomes, and [Farre-Mensa et al. \(2020\)](#) estimate that patents cause 36 percent higher employment for startups, using examiner leniency as an instrument. Both compare different inventions across examiners; I hold invention quality fixed by comparing applications within the same family. The 29.8 percent discordance rate among reassigned twins provides a direct measure of the examination lottery that these cross-sectional designs cannot isolate.

Second, this paper relates to the growing literature on regulatory discretion and its distributional consequences. [Kling \(2006\)](#) and [Dobbie et al. \(2018\)](#) show that judge identity causally affects sentencing and downstream outcomes in criminal justice. [Maestas et al. \(2013\)](#) demonstrate similar leniency effects in disability determination. My results extend the examiner-leniency paradigm to intellectual property, where the stakes—exclusive market rights rather than years of incarceration—are economic rather than punitive, and the small-entity penalty reveals how regulatory inconsistency compounds existing inequalities in the innovation system.

Third, I contribute to debates about patent quality and patent reform. [Frakes and Wasserman \(2017\)](#) show that examiner time pressure reduces examination thoroughness, and [Galasso and Schankerman \(2015\)](#) find that patent invalidation spurs follow-on innovation. My finding that 29.8 percent of reassigned continuations receive different verdicts suggests that the examination process itself introduces noise into the patent system at a scale comparable to the effects of intentional policy changes. The portability of this “examination lottery” concept extends beyond patents: any administrative adjudication system with examiner discretion—immigration, tax audits, disability reviews—faces the same consistency challenge.

The remainder of the paper proceeds as follows. [Section 2](#) describes the institutional setting. [Section 3](#) presents the data. [Section 4](#) details the empirical strategy. [Section 5](#) reports results. [Section 6](#) discusses implications.

2. Institutional Background

Patent examination at the USPTO. The USPTO receives over 600,000 utility patent applications annually. Each application is assigned to one of approximately 8,000 examiners organized into art units—groups of 10–20 examiners specializing in a narrow technology area. The examiner reviews the application against prior art and statutory requirements (novelty under 35 U.S.C. §102, non-obviousness under §103, written description under §112), issuing office actions that may accept or reject claims. Applications ultimately receive a “disposal”—issuance as a patent (ISS) or abandonment (ABN)—typically within 3–5 years of filing (Frakes and Wasserman, 2017).

Continuation and divisional applications. Inventors may file continuation applications that claim priority from a parent application, maintaining the same effective filing date while presenting different or refined claims. Divisional applications arise when the examiner determines that the parent contains multiple distinct inventions. Both continuation (CON) and divisional (DIV) applications enter the same art unit as the parent but may be assigned to a different examiner based on workload and availability. This reassignment is the source of identifying variation in this paper.

Examiner assignment. Within an art unit, applications are assigned to examiners through a combination of specialty matching and workload balancing. The assignment process does not condition on application quality or the identity of the inventor. While examiners may develop expertise in narrower subfields within an art unit, the decision of which examiner receives a particular continuation is primarily administrative—an “as-good-as-random” assignment conditional on art unit and filing year (Sampat and Williams, 2019).

Examiner heterogeneity. Prior work documents substantial variation in examiner grant rates even within art units. Cockburn et al. (2003) find that allowance rates range from 35 to 97 percent across examiners in the same technology center. Lemley (2001) argues that this heterogeneity reflects differences in effort, risk tolerance, and interpretation of legal standards rather than case mix. My data confirm this: among examiners with 50 or more decided applications, the within-art-unit standard deviation of grant rates is 0.172.

3. Data

I construct the dataset from the USPTO Patent Examination Research Dataset (PatEx) (Marco et al., 2015), accessed via Google BigQuery. PatEx contains the universe of 9.8 million

Table 1: Summary Statistics: Patent Continuation Twin Pairs

	N	Mean	SD
Child granted	702,490	0.781	0.414
Parent granted	702,490	0.841	0.366
Discordant outcome	702,490	0.203	0.402
Reassigned (diff. examiner)	702,490	0.139	0.346
Small entity	702,490	0.219	0.414
Examiner grant rate (LOO)	702,490	0.749	0.172
Filing year	702,490	2007.792	4.519

Notes: Same-art-unit continuation and divisional application pairs filed 1998–2015 with resolved outcomes (issued or abandoned). Examiner grant rate is leave-one-out.

patent applications with filing dates, examiner identifiers, art unit assignments, disposal types, and entity status. I link parent-child continuation pairs using the continuity data and merge office action records from the Office Action Research Dataset (Lu et al., 2017).

Sample construction. I begin with all continuation (CON) and divisional (DIV) parent-child pairs filed between 1998 and 2015 where both applications have resolved outcomes (issued or abandoned) and non-missing examiner identifiers. I restrict to pairs where the child is in the same art unit as the parent, yielding 702,490 same-art-unit pairs. Of these, 97,922 (13.9 percent) were reassigned to a different examiner. The filing-year restriction ensures at least five years for disposition by the time of data collection.

Key variables. The outcome is a binary indicator for whether the child application was granted (disposal type ISS). The treatment variable in the twin comparison is an indicator for reassignment (different examiner than the parent). For the leniency analysis, I construct each examiner’s leave-one-out (LOO) grant rate: the share of all other applications decided by that examiner that were granted. The LOO rate averages 0.749 with a standard deviation of 0.172, indicating substantial examiner-level variation in grant propensity.

Table 1 presents summary statistics. The overall grant rate for child applications is 78.1 percent, slightly below the parent grant rate of 84.0 percent—consistent with continuations facing somewhat tougher examination since the parent’s claims have already been partially adjudicated. Small entities (those qualifying for reduced USPTO fees) account for 21.9 percent of the sample.

4. Empirical Strategy

Twin study design. The core comparison exploits the within-family structure of continuation applications. For a given parent-child pair (p, c) in art unit a filed in year t , I estimate:

$$\text{Discordant}_{pc} = \alpha + \beta \cdot \text{Reassigned}_{pc} + \gamma \cdot \text{ParentGranted}_{pc} + \delta_{at} + \varepsilon_{pc} \quad (1)$$

where $\text{Discordant}_{pc} = \mathbb{I}[\text{Disposal}_c \neq \text{Disposal}_p]$, $\text{Reassigned}_{pc} = \mathbb{I}[\text{Examiner}_c \neq \text{Examiner}_p]$, and δ_{at} are art-unit-by-filing-year fixed effects. The coefficient β measures the causal effect of examiner reassignment on outcome inconsistency, identified under the assumption that reassignment is orthogonal to invention quality within art-unit-year cells.

Examiner leniency design. To measure the intensity of the examination lottery, I estimate:

$$\text{Granted}_c = \alpha + \beta \cdot \text{Leniency}_c^{-i} + \gamma \cdot \text{ParentGranted}_{pc} + \delta_{at} + \varepsilon_c \quad (2)$$

where Leniency_c^{-i} is the leave-one-out grant rate of the child’s examiner. This specification is estimated on the subsample of reassigned pairs, where the child examiner is plausibly as-good-as-random conditional on δ_{at} . The coefficient β captures the partial derivative of the grant probability with respect to examiner leniency, holding the invention family fixed through parent outcome and art-unit-year fixed effects.

Identification assumptions. The identifying assumption is that examiner assignment for continuation applications is independent of application quality conditional on art-unit-by-filing-year fixed effects. This follows from the administrative assignment mechanism described in [Section 2](#): within an art unit and year, which examiner receives a continuation depends on workload and availability rather than case characteristics.

I probe this assumption with balance tests. [Table 3](#) shows that conditional on art-unit-year fixed effects, the correlation between examiner leniency and small-entity status is -0.009 ($t = -5.5$), and the correlation with parent grant status is 0.007 ($t = 4.7$). While both are statistically significant—as expected with nearly 100,000 observations—they are economically negligible, explaining less than 0.1 percent of leniency variation (within $R^2 < 0.001$). The identifying assumption requires conditional independence, not zero correlation in finite samples; these magnitudes are consistent with as-good-as-random assignment.

Scope of the twin comparison. An important caveat is that continuation and divisional applications are not identical to their parents. Children may present refined claims, different

prosecution strategies, or narrower invention scope. The within-family design holds the invention *family* fixed—not the precise claim set. The parent outcome control absorbs family-level quality differences, but residual within-family heterogeneity remains. The results should therefore be interpreted as measuring the contribution of examiner identity to outcome variation *within families*, not as a pure test of whether identical applications receive different verdicts. Continuations (CON) provide the cleaner comparison because they pursue substantially similar claims; divisionals (DIV) may involve distinct inventions within the same specification. Robustness results in [Table 5](#) confirm that the leniency effect is similar across both types.

Inference. Standard errors are clustered at the examiner level for specifications using examiner leniency and at the art-unit level for discordance specifications, reflecting the levels at which treatment varies. With 10,575 unique examiners and 671 art units, cluster counts are sufficiently large for reliable asymptotic inference.

5. Results

5.1 The Examination Lottery

[Table 2](#) documents the core fact. Among same-art-unit pairs where the child was reassigned to a different examiner, 29.8 percent produced discordant outcomes—one application granted, the other abandoned. For same-examiner pairs, discordance is 18.7 percent. The raw excess discordance of 11.0 percentage points captures the additional inconsistency introduced by examiner reassignment.

The regression estimate in Panel B of [Table 2](#) confirms this pattern with art-unit-by-filing-year fixed effects: reassignment increases discordance by 7.7 percentage points (SE = 0.004). The difference from the raw gap reflects composition: reassigned pairs are somewhat more concentrated in technology areas with higher baseline discordance.

5.2 Balance

[Table 3](#) reports balance tests for the examiner leniency instrument. Conditional on art-unit-year fixed effects, examiner leniency is nearly orthogonal to observable characteristics of the application: the coefficient on small-entity status is -0.009 and on parent-grant status is 0.007 . Both are statistically significant but explain less than 0.1 percent of leniency variation, consistent with as-good-as-random assignment within art-unit-year cells.

Table 2: Discordance Rates by Examiner Reassignment and Entity Size

	All Pairs		By Entity Size	
	N	Discordance	Small	Large
<i>Panel A: Levels</i>				
Different examiner	97,922	0.298	0.384	0.272
Same examiner	604,568	0.187	0.252	0.170
<i>Panel B: Difference</i>				
Excess discordance		0.110 (0.003)		
Art unit \times year FE		Yes		
Clustered SE (art unit)		Yes		

Notes: Discordance means the child and parent applications received different dispositions (one issued, one abandoned). Panel B reports the coefficient on a reassignment indicator from a regression of discordance on reassignment with art-unit \times filing-year fixed effects. Standard errors clustered at the art-unit level in parentheses.

5.3 Examiner Leniency and Grant Outcomes

Table 4 presents the leniency results. Column (1) shows the preferred specification among reassigned pairs: a one-unit increase in examiner LOO grant rate raises the child’s grant probability by 0.661 (SE = 0.020), controlling for parent outcome and art-unit-year fixed effects. At the mean leniency standard deviation of 0.172, this implies that a one-standard-deviation increase in examiner leniency raises the grant probability by 11.4 percentage points—a large effect relative to the 68.5 percent baseline grant rate for reassigned children. The t -statistic on leniency exceeds 33, confirming that examiner grant propensity is a powerful predictor of individual outcomes.

Column (3) extends to the full sample and interacts reassignment with leniency. The leniency effect for same-examiner pairs is 0.491, reflecting the mechanical correlation between an examiner’s leniency and their own decisions. For reassigned pairs, the additional leniency effect is 0.112 ($t = 8.7$), consistent with reassigned children being more sensitive to the lottery of examiner assignment since they face a new decision-maker who has no memory of the parent case.

5.4 The Small-Entity Penalty

The examination lottery falls disproportionately on small entities. In the raw data, reassigned small-entity pairs exhibit 38.4 percent discordance versus 27.2 percent for large entities—a gap of 11.2 percentage points (Table 2). The triple-difference estimate in the regression framework is 3.2 percentage points ($t = 4.3$): small entities face significantly greater excess

Table 3: Balance Tests: Examiner Leniency and Observable Characteristics

	Dep. var.: Examiner leniency (LOO)	
	(1)	(2)
Small entity	-0.0092 (0.0017)	
Parent granted		0.0072 (0.0015)
Art unit \times year FE	Yes	Yes
N	97,922	97,922
R^2	0.643	0.643

Notes: Dependent variable is the child examiner’s leave-one-out grant rate. Sample restricted to reassigned same-art-unit continuation pairs. Standard errors clustered at the examiner level in parentheses.

discordance from reassignment even after absorbing art-unit-year fixed effects and controlling for parent outcome.

The leniency analysis confirms this pattern. The coefficient on examiner LOO rate is 0.692 for small entities versus 0.626 for large entities. Small inventors’ grant outcomes are more sensitive to examiner identity, consistent with large entities having greater capacity to appeal adverse decisions, hire experienced patent attorneys, or restructure claims in response to examiner-specific feedback ([Lanjouw and Schankerman, 2001](#)).

5.5 Robustness

[Table 5](#) demonstrates that the leniency effect is stable across multiple dimensions. The coefficient is similar for continuations (0.639) and divisionals (0.694), for early filings (0.650, 1998–2006) and late filings (0.675, 2007–2015), and regardless of whether the parent was granted (0.647) or abandoned (0.666). Using within-art-unit-year standardized leniency yields a coefficient of 0.065 per standard deviation ($t = 31.0$), confirming that the result is not driven by cross-art-unit variation.

6. Discussion

These results reveal a regulatory lottery at the heart of the U.S. patent system. When the same invention family faces a different examiner, 29.8 percent of applications receive a different verdict. This is not a claim that examiners are incompetent or biased—the examination process involves genuine judgment under uncertainty, and reasonable experts may disagree. The point is that the magnitude of this disagreement is large enough to constitute a first-order

Table 4: Examiner Leniency and Patent Grant Outcomes

	Dep. var.: Child application granted		
	(1)	(2)	(3)
	Reassigned	Reassigned	Full sample
Examiner leniency (LOO)	0.661 (0.020)	0.661 (0.020)	0.491 (0.009)
Parent granted	0.166 (0.006)		0.277 (0.003)
Reassigned			-0.120 (0.010)
Art unit \times year FE	Yes	Yes	Yes
Parent outcome control	Yes	No	Yes
N	97,922	97,922	702,490
R^2	0.282	0.282	0.240

Notes: Examiner leniency is the leave-one-out grant rate of the child application’s examiner. Column (1) controls for parent outcome; column (2) omits parent control; column (3) includes the full sample with a reassignment indicator and its interaction with leniency. Standard errors clustered at the examiner level in parentheses.

distortion in the allocation of patent rights.

Economic magnitude. The examination lottery operates at considerable scale. Among 97,922 reassigned same-AU pairs, the 7.7-percentage-point excess discordance implies that roughly 7,500 applications received outcomes that would have been different under their parent’s examiner. While this paper does not directly estimate downstream innovation effects, prior work suggests the stakes are material: [Farre-Mensa et al. \(2020\)](#) find that patent grants cause 36 percent higher employment for startups, implying that even a fraction of these discordant outcomes could meaningfully affect resource allocation in the innovation system.

The small-entity penalty. The disproportionate burden on small entities has implications for innovation policy. If the patent system functions as a lottery, resource-constrained inventors—who cannot afford repeated filings, experienced attorneys, or examiner-specific claim strategies—are most vulnerable to adverse draws. This compounds existing disparities documented by [Jensen et al. \(2018\)](#) in patent prosecution success rates.

Limitations. The balance tests reveal small but statistically significant correlations between examiner leniency and application observables, suggesting that assignment is not perfectly random. While these correlations are economically negligible (within $R^2 < 0.001$), they warrant caution in interpreting the leniency estimates as strictly causal. The twin design

Table 5: Robustness: Examiner Leniency Effect Across Subsamples

	Coefficient	SE	N
<i>Panel A: Application type</i>			
Continuations (CON)	0.639	(0.024)	60,032
Divisionals (DIV)	0.694	(0.025)	37,890
<i>Panel B: Time period</i>			
Early (1998–2006)	0.650	(0.025)	55,565
Late (2007–2015)	0.675	(0.028)	42,357
<i>Panel C: Parent outcome</i>			
Parent granted	0.647	(0.021)	79,398
Parent abandoned	0.666	(0.037)	18,524
<i>Panel D: Alternative leniency</i>			
Standardized (within AU×year)	0.065	(0.002)	97,834

Notes: Each row reports the coefficient on examiner leniency from a regression of child grant outcome on examiner leniency with art-unit \times filing-year fixed effects. Sample restricted to reassigned same-art-unit pairs unless noted. Standard errors clustered at the examiner level in parentheses. Panel D reports the coefficient on within-AU \times year standardized leniency.

addresses this concern by holding the invention family fixed within families, but some residual selection—for instance, if certain types of continuations are systematically routed to particular examiners—cannot be fully ruled out. Additionally, the LOO leniency measure uses an examiner’s full career history rather than contemporaneous rates, potentially introducing measurement error that attenuates estimates.

Policy implications. Reducing examination inconsistency does not require eliminating examiner discretion. Targeted interventions—standardized examination guidelines with calibration exercises, blind second opinions on borderline cases, or statistical monitoring of examiner-level discordance rates—could reduce arbitrary variation while preserving the benefits of specialized human judgment. The twin study methodology developed here could serve as an ongoing audit tool for the patent office.

7. Conclusion

Patent examination operates partly as a lottery. Using 97,922 within-family twin comparisons at the USPTO, I show that examiner reassignment raises outcome discordance by 7.7 percentage points, with small entities bearing a disproportionate 3.2-percentage-point additional burden. The portable lesson is that any administrative adjudication system with examiner discretion faces a consistency challenge that can be measured—and potentially mitigated—using the twin study design introduced here. The examination lottery is not

unique to patents; it is a structural feature of regulatory discretion that deserves systematic attention.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

- 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.
- Dobbie, Will, Jacob Goldin, and Crystal S Yang**, “The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges,” *American Economic Review*, 2018, *108* (2), 201–240.
- 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**, “Does Patent Invalidation Reduce Strategic Patenting?,” *American Economic Review*, 2015, *105* (10), 3145–3179.
- Jaffe, Adam B and Josh Lerner**, “Innovation and Its Discontents,” *Innovation Policy and the Economy*, 2004, *4*, 27–73.
- Jensen, Kyle, Fiona Murray, and Olav Sorenson**, “Gender Differences in Obtaining and Maintaining Patent Rights,” *Nature Biotechnology*, 2018, *36*, 307–309.
- Kling, Jeffrey R**, “Incarceration Length, Employment, and Earnings,” *American Economic Review*, 2006, *96* (3), 863–876.
- Lanjouw, Jean O and Mark Schankerman**, “Characteristics of Patent Litigation: A Window on Competition,” *RAND Journal of Economics*, 2001, *32* (1), 129–151.
- Lemley, Mark A**, “Rational Ignorance at the Patent Office,” *Northwestern University Law Review*, 2001, *95* (4), 1495–1532.
- Lichtman, Doug and Mark A Lemley**, “Rethinking Patent Law’s Presumption of Validity,” *Stanford Law Review*, 2004, *60*, 45–72.
- Lu, Guan-Cheng, Anthony Cardillo, Josh Lerner, and Alan C Marco**, “The USPTO Patent Prosecution Research Data: Unlocking Office Action Traits,” *USPTO Economic Working Paper*, 2017, (2017-10).

Maestas, Nicole, Kathleen J Mullen, and Alexander Strand, “Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt,” *American Economic Review*, 2013, *103* (5), 1797–1829.

Marco, Alan C, Joshua D Sarnoff, and Charles A deGrazia, “The USPTO Patent Examination Research Dataset: A Window on Patent Processing,” Technical Report 2015-4, USPTO Economic Working Paper 2015.

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.

A. Data Appendix

Data sources. The primary data come from the USPTO Patent Examination Research Dataset (PatEx), accessed via Google BigQuery (`patents-public-data.uspto_oce_pair`). This dataset contains the universe of 9.8 million utility patent applications filed at the USPTO, including examiner identifiers, art unit assignments, filing dates, disposal types, and entity status indicators. Continuation and divisional parent-child links come from the continuity tables within the same dataset. Office action data (rejection types and dates) come from the Office Action Research Dataset (`patents-public-data.uspto_oce_office_actions`).

Sample construction. Starting from all continuation (CON) and divisional (DIV) parent-child pairs with child filing dates between January 1, 1998 and December 31, 2015:

1. Require both parent and child to have resolved outcomes (disposal type ISS or ABN): 1,160,090 pairs
2. Require non-missing examiner identifiers and art units for both: 1,160,090 pairs
3. Restrict to same-art-unit pairs (child art unit equals parent art unit): 702,490 pairs
4. Of these, 97,922 (13.9%) have different examiners (reassigned)

Variable definitions.

- **Granted:** Binary indicator equal to 1 if disposal type is ISS (issued).
- **Discordant:** Binary indicator equal to 1 if child and parent disposal types differ.
- **Reassigned:** Binary indicator equal to 1 if child examiner ID differs from parent examiner ID.
- **Examiner leniency (LOO):** Leave-one-out grant rate: for child application i assigned to examiner j , $\text{Leniency}_{j(-i)} = \frac{\sum_{k \neq i} \text{Granted}_k}{\sum_{k \neq i} 1}$ over all applications decided by examiner j .
- **Small entity:** Binary indicator equal to 1 if the child application’s entity status is “small” (qualifying for reduced USPTO fees under 37 CFR §1.27).
- **Technology center:** First two digits of the art unit code, grouping related technology areas.

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Child application granted	0.661	0.020	0.414	0.300	0.009	Large positive
Discordant outcome	0.077	0.004	0.402	0.192	0.009	Large positive
<i>Panel B: Heterogeneous (entity size)</i>						
Child granted (small entity)	0.692	0.036	0.470	0.290	0.015	Large positive
Child granted (large entity)	0.626	0.022	0.391	0.292	0.010	Large positive

Notes: **Country:** United States. **Research question:** Does quasi-random reassignment of continuation patent applications to different examiners within the same art unit cause inconsistent grant outcomes, and do small entities bear disproportionate costs? **Policy mechanism:** When inventors file continuation or divisional applications at the USPTO, workload-balancing reassignment exposes the same invention family to different examiners with heterogeneous grant propensities, creating a regulatory lottery that randomly reallocates patent rights. **Outcome definition:** Child application granted is a binary indicator equal to one if the continuation application was issued as a patent (disposal type ISS). Discordant outcome equals one if parent and child received different dispositions. **Treatment:** Continuous examiner leave-one-out grant rate for rows 1, 3, 4; binary reassignment indicator for row 2. **Data:** USPTO PAIR via Google BigQuery, 1998–2015 filing years, application-pair level, 702,490 same-art-unit pairs. **Method:** OLS with art-unit \times filing-year fixed effects, standard errors clustered at examiner level (rows 1, 3, 4) or art-unit level (row 2). **Sample:** Continuation (CON) and divisional (DIV) applications with resolved outcomes (issued or abandoned) where both parent and child are in the same art unit. $SDE = \hat{\beta} \times SD(X)/SD(Y)$ for continuous treatment, $\hat{\beta}/SD(Y)$ for binary treatment. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

B. Robustness Appendix

The leniency effect is stable across all subsamples tested (Table 5). Continuations (CON) and divisionals (DIV) yield similar coefficients (0.639 vs. 0.694), ruling out concerns that the results are driven by one application type. The temporal stability across early (1998–2006: 0.650) and late (2007–2015: 0.675) periods indicates that the examination lottery is not an artifact of a particular policy regime or examiner cohort. Conditioning on parent outcome (granted vs. abandoned) produces nearly identical estimates (0.647 vs. 0.666), confirming that the leniency effect does not reflect differential selection into continuation filing based on parent success.

C. Standardized Effect Sizes