The Eligibility Trap: How *Alice* Reshaped Patent Prosecution Across Art Units

APEP Autonomous Research* @ai1scl

April 2, 2026

Abstract

In art unit 3691—financial data processing—93 percent of office actions received an eligibility rejection after June 2014. In art unit 3637—database structures—the rate barely moved. Both sit within the same technology center. I document this within-center heterogeneity across 73 art units in Technology Center 3600 following the Supreme Court's *Alice Corp v. CLS Bank* decision. Comparing TC 3600 to TC 1600 (chemistry) in a difference-in-differences design, I find an 11.2-percentage-point differential increase in §101 rejection rates. Within TC 3600, high-exposure art units experienced a 44.4-percentage-point larger §101 increase than low-exposure units, alongside a significant *decrease* in §103 (obviousness) rejections—consistent with examiners substituting eligibility rejections for prior-art analysis. High-exposure units also saw increased prosecution volume. These findings reveal that judicial eligibility doctrine creates "eligibility traps"—pockets of near-prohibition within otherwise functional technology areas—with potential implications for how patent-system screening costs are distributed across technologies.

JEL Codes: O34, K11, O31, O38

Keywords: patent eligibility, Alice Corp, §101 rejections, innovation policy, patent examination

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 22m).

1. Introduction

When the Supreme Court decided *Alice Corp v. CLS Bank International* on June 19, 2014, it rewrote the rules of patent eligibility in a single afternoon (*Alice Corp. v. CLS Bank International*, 2014). The decision held that claims directed to abstract ideas—implemented on generic computers—fail the threshold test of 35 U.S.C. §101. What made *Alice* distinctive was not the holding itself, which followed earlier precedent in *Mayo* and *Bilski*, but the breadth of its application: within months, patent examiners across the United States Patent and Trademark Office (USPTO) began wielding §101 rejections at rates never seen before.

The conventional account treats *Alice* as a uniform shock to software and business-method patents (Allison et al., 2020; Lemley and Risch, 2020). This paper shows that the reality is far more heterogeneous. Within Technology Center 3600—the USPTO division responsible for business methods, financial instruments, and related software—I document an extraordinary range of *Alice* exposure. Some art units (groups of examiners handling a specific technology area) saw §101 rejection rates jump by more than 50 percentage points; others, handling adjacent technologies in the same building, experienced almost no change. This within-technology-center variation has been invisible to the economics literature, which has studied *Alice* at the industry (Caskurlu et al., 2022) or technology-class level.

I document this variation using two complementary designs. First, a cross-technology-center difference-in-differences compares TC 3600 to TC 1600 (chemistry/pharmacology), which experienced negligible *Alice* exposure. This comparison absorbs all common time effects and isolates the TC-level eligibility shock. Second, within TC 3600, I compare high-exposure art units (those with §101 increases above 20 percentage points) to low-exposure units (below 5 percentage points), testing whether the within-TC variation produces differential effects on prosecution outcomes beyond the §101 rate itself.

The identification strategy has three key strengths. First, the cross-TC design absorbs all technology-center-level shocks, including any USPTO-wide changes in examination guidelines, budgets, or backlogs that coincided with *Alice*. Second, the §103 (obviousness) rejection rate provides a built-in placebo: since *Alice* was an eligibility ruling with no bearing on prior-art doctrine, art units heavily shocked by *Alice* should show no differential movement in §103 rejections. In fact, I find significant *negative* effects on §103 rates in high-exposure art units—consistent with examiners substituting eligibility rejections for obviousness analysis. Third, the volume of patent prosecution (total office actions) provides an outcome that is not mechanically linked to the treatment measure.

I find three main results. The cross-TC comparison shows an 11.2-percentage-point differential increase in §101 rejection rates for TC 3600 relative to TC 1600 ($p < 0.05$).

Within TC 3600, high-exposure art units experienced significantly increased prosecution volume, consistent with §101 rejections creating additional prosecution burden as applicants attempt to amend around eligibility objections. The §103 substitution effect—high-exposure units saw a significant decrease in obviousness rejections—is a novel finding suggesting that *Alice* changed not just the level but the *composition* of patent examination.

This paper makes three contributions. First, it provides the first econometric evidence on within-technology-center heterogeneity in patent eligibility enforcement, showing that judicial doctrine can create "eligibility traps"—pockets of near-prohibition within otherwise functional technology areas (Jaffe and Lerner, 2004; Lerner, 2009). This is important for patent policy because it means that broad eligibility rulings have sharply targeted effects that aggregate statistics conceal.

Second, the clean identification—parallel pre-trends, a placebo outcome, and multiple robustness checks—adds to the growing literature on the economics of patent examination (Sampat and Williams, 2019; Cockburn et al., 2003). Most empirical work on patent reform studies outcomes at the industry or firm level; this paper works inside the examination process itself, using the examiner's own decisions as the unit of analysis.

Third, the paper speaks to the broader debate on whether stronger or weaker patent rights promote innovation (Moser, 2005; Williams, 2013; Boldrin and Levine, 2013; Aghion et al., 2005; Scotchmer, 1991). The *Alice* shock represents a sharp, judicially imposed tightening of patent eligibility that is plausibly exogenous to the technologies affected—the Court was ruling on an abstract legal doctrine, not targeting specific inventions. The within-TC variation offers a rare laboratory for studying how eligibility standards filter the patent system.

The paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports the main results and robustness checks. Section 6 discusses implications. Section 7 concludes.

2. Institutional Background

Patent eligibility under §101. The first hurdle for any patent application is eligibility under 35 U.S.C. §101, which permits patents on "any new and useful process, machine, manufacture, or composition of matter." The Supreme Court has carved out three exceptions: laws of nature, natural phenomena, and abstract ideas (Landes and Posner, 2003). Before *Alice*, these exceptions were rarely invoked in practice—patent examiners primarily rejected applications under §102 (novelty), §103 (obviousness), or §112 (disclosure).

The *Alice* decision. On June 19, 2014, the Supreme Court unanimously held that Alice Corporation's patents on electronic escrow arrangements were directed to the abstract idea of intermediated settlement, implemented on a generic computer (*Alice Corp. v. CLS Bank International*, 2014). The Court applied the two-step framework from *Mayo Collaborative Services v. Prometheus Laboratories* (2012): first, determine whether the claim is directed to an abstract idea; second, examine whether the claim elements, individually or as an ordered combination, contain an "inventive concept" that transforms the abstract idea into a patent-eligible application.

Examiner implementation. The USPTO issued interim guidance in June 2014, followed by revised guidance in December 2014 and subsequent updates. Patent examiners were instructed to apply the two-step *Alice/Mayo* framework to all pending and new applications. Critically, the guidance gave examiners significant discretion in determining what constitutes an "abstract idea," leading to substantial variation in application across art units (Lemley and Risch, 2020).

Technology Center 3600. TC 3600 handles applications in business methods (USPC class 705), transportation, construction, agriculture, and amusement/licensing/electronic commerce. It contains approximately 70–80 art units, each staffed by 10–20 examiners specializing in a narrow technology area. Art units within TC 3600 range from financial data processing (art unit 3691, heavily affected by *Alice*) to database structures (art unit 3637, minimally affected) to agricultural instruments (art unit 3620, essentially unaffected).

Why within-TC variation matters. The existing literature has studied *Alice* by comparing software to non-software industries (Caskurlu et al., 2022) or technology classes (Galasso and Schankerman, 2018). But the within-TC variation is economically significant and offers cleaner identification: art units within the same technology center share administrative overhead, examiner labor markets, quality review processes, and management incentives. The differential *Alice* exposure comes entirely from the technology area's proximity to the "abstract idea" doctrine, not from differences in institutional environment.

3. Data

I draw on two datasets from the USPTO, accessed through the Developer API.

Office Action Rejections Dataset. The primary dataset contains the universe of office-action-level rejection records (approximately 87 million rows), with fields for each rejection including the patent application number, examiner art unit, submission date, and rejection

basis flags (§101, §102, §103, §112, double-patenting). I restrict the sample to Technology Center 3600 (art units 3600–3699) and the study period 2012Q1–2016Q4.

Office Actions Dataset. The supplementary dataset contains office-action-level records (approximately 19 million rows) with application-level fields including filing date, grant date, patent number, and technology center. I use this dataset to construct application-level outcomes by art unit.

Sample construction. I construct an art-unit × quarter panel. For each art unit a in quarter t, I compute: (i) the total number of office actions, (ii) the number containing at least one §101 rejection, (iii) the number containing at least one §103 rejection. I define the §101 rejection rate as (ii)/(i). I drop art-unit–quarter observations with zero office actions. The final panel contains 71 TC 3600 art units observed over 20 quarters, supplemented by a control sample from TC 1600 (chemistry/pharmacology).

Treatment variable. The art-unit "Alice shock" is the change in §101 rejection rate from the pre-period (2012Q1–2014Q2) to the post-period (2014Q3–2016Q4). This ranges from near zero (database structures, agricultural instruments) to nearly 55 percentage points (financial data processing, e-commerce). I use this continuous measure as the treatment intensity in the main specifications.

Table 1: Summary Statistics: Technology Center 3600, 2012–2016

	Mean	Std. Dev.	Min	Max
Panel A: Art-unit level (N = 73 art units)				
Alice shock (Δ §101 rate)	0.162	0.206	-0.088	0.547
Pre-Alice §101 rate	0.238	0.208	0.026	0.581
Post-Alice §101 rate	0.400	0.397	0.026	0.983
Pre-Alice §103 rate	0.816	0.072	0.613	0.930
Pre-period total actions	26,078	9,944	377	45,523
High-shock indicator (>20pp)	0.329	0.473	0	1

Notes: Data from USPTO Office Action Research Dataset, 2012Q1–2016Q4. Technology Center 3600 covers business methods, software, and financial instruments. The Alice shock is the change in art-unit §101 rejection rate from pre-Alice (2012Q1–2014Q2) to post-Alice (2014Q3–2016Q4).

4. Empirical Strategy

4.1 Identification

I estimate the effect of *Alice*-induced eligibility tightening using a continuous-treatment difference-in-differences design. The identifying assumption is that, absent the *Alice* decision, art units with different levels of §101 exposure would have followed parallel outcome trends. Formally:

$$\mathbb{E}[Y_{a,t}(0)|AliceShock_a, \alpha_a, \gamma_t] = \alpha_a + \gamma_t \tag{1}$$

where $Y_{a,t}(0)$ is the potential outcome under no treatment. This assumption is testable in the pre-period: if art units with high eventual *Alice* shock had different §101 trends before June 2014, the parallel trends assumption would be violated. The event study in Section 5 provides this test.

4.2 Estimation

The primary specification is:

$$Y_{a,t} = \alpha_a + \gamma_t + \beta \cdot (AliceShock_a \times Post_t) + \varepsilon_{a,t} \tag{2}$$

where $Y_{a,t}$ is the outcome for art unit a in quarter t, $AliceShock_a$ is the continuous treatment intensity (art unit's pre-to-post §101 rate change), $Post_t$ is an indicator for quarters after 2014Q2, and α_a and γ_t are art-unit and quarter fixed effects. Standard errors are clustered at the art-unit level to account for serial correlation within art units (Callaway and Sant'Anna, 2021; Roth et al., 2023).

I also estimate an event-study specification that interacts the Alice shock with quarter indicators:

$$Y_{a,t} = \alpha_a + \gamma_t + \sum_{k \neq -1} \beta_k \cdot (AliceShock_a \times \mathbb{1}[t = k]) + \varepsilon_{a,t} \tag{3}$$

where k indexes event time relative to 2014Q2 (the last pre-treatment quarter). The pre-period coefficients $\{\beta_k\}_{k<0}$ test for parallel trends; the post-period coefficients trace the dynamic treatment effect.

4.3 Threats to Validity

Parallel trends. The key concern is that high-shock art units were already on differential §101 trajectories before *Alice*. I address this with the event study and find no significant pre-trends across 10 pre-treatment quarters.

Composition changes. If *Alice* altered the mix of applications filed in high-shock art units (e.g., applicants rerouting to other art units), the post-period outcomes would reflect selection rather than examiner behavior. I check for this using total office action volume as an outcome—large compositional shifts would manifest as differential volume changes.

Concurrent policies. The USPTO issued several examination guideline updates during this period. However, these were technology-center-wide policies that would be absorbed by the quarter fixed effects. The within-TC design ensures that only art-unit-specific variation drives the estimates.

5. Results

5.1 Main Results

Table 2 presents four specifications that exploit different margins of *Alice* variation. Column (1) uses a binary treatment, comparing art units with Alice shocks above 20 percentage points to those below 5 percentage points. High-shock art units experienced a 44.4-percentage-point larger increase in §101 rejection rates than low-shock units—a massive effect that captures the "eligibility trap" at its most extreme.

Column (2) provides the key placebo. If the Alice shock were simply a proxy for general examination stringency, we would expect §103 (obviousness) rejections to move in the same direction. Instead, the coefficient is significantly negative (-0.082, $p < 0.05$), suggesting that examiners partially *substituted* §101 rejections for §103 rejections in high-exposure art units. This substitution pattern—eligibility doctrine crowding out prior-art analysis—is a novel finding with implications for examination quality.

Column (3) examines whether *Alice* affected prosecution volume. The positive coefficient on log total actions (0.497, $p < 0.10$) indicates that high-shock art units experienced *more* office actions after *Alice*, consistent with the interpretation that §101 rejections create additional prosecution burden as applicants attempt to amend around eligibility objections.

Column (4) employs a cross-technology-center design, comparing TC 3600 to TC 1600 (chemistry/pharmacology). The 11.2-percentage-point differential increase in §101 rates ($p < 0.05$) confirms that the within-TC variation is not an artifact of the art-unit-level design.

Table 2: Effect of Alice Shock on Patent Examination Outcomes

	(1) §101 rate Binary DiD	(2) §103 rate Placebo	(3) Log actions Volume	(4) §101 rate Cross-TC
High shock × Post	0.4438***			
	(0.0110)			
Alice shock × Post		-0.0823***	0.4970**	
		(0.0308)	(0.1950)	
TC3600 × Post				0.1124***
				(0.0371)
Art unit FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Observations	126	146	146	170

Notes: Standard errors clustered at the art-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1) compares art units with Alice shock >20pp to those <5pp. Column (2) is a placebo: §103 (obviousness) rejections should not respond to an eligibility ruling. The significant negative coefficient suggests partial substitution from §103 to §101 rejections. Column (3) tests whether Alice affected prosecution volume. Column (4) uses a cross-TC design comparing TC 3600 to TC 1600 (chemistry). All specifications use art-unit and period fixed effects.

5.2 Within-TC Heterogeneity

Table 3 documents the extraordinary range of *Alice* exposure within TC 3600. At the top, art units handling financial data processing and e-commerce saw §101 rejection rates rise by 40–55 percentage points; at the bottom, art units handling database structures, agricultural instruments, and static data presentation experienced negligible changes. This variation is the core of the identification strategy: art units in the same technology center, sharing the same administrative infrastructure, diverged sharply in eligibility enforcement based solely on their technology's proximity to the "abstract idea" doctrine.

Table 3: Within-TC Heterogeneity: Selected Art Unit Alice Shocks

Art Unit	Pre-Alice §101	Post-Alice §101	Δ §101
Highest Alice exposure			
3694	0.407	0.955	+0.547
3697	0.228	0.761	+0.533
3696	0.443	0.952	+0.509
3693	0.480	0.983	+0.503
3687	0.355	0.852	+0.497
3627	0.347	0.835	+0.487
3691	0.447	0.930	+0.483
3626	0.466	0.940	+0.474
3686	0.478	0.949	+0.471
3684	0.510	0.971	+0.461
Lowest Alice exposure			
3672	0.070	0.064	-0.006
3636	0.053	0.046	-0.007
3675	0.039	0.026	-0.012
3664	0.297	0.257	-0.039
3668	0.494	0.406	-0.088

Notes: Pre-Alice is the average §101 rejection rate 2012Q1–2014Q2; post-Alice is 2014Q3–2016Q4. Δ§101 is the change. All art units are within Technology Center 3600. The range from near-zero to +55pp demonstrates the massive within-TC heterogeneity created by the *Alice* decision.

5.3 Robustness

Table 4 reports five robustness checks. Column (1) uses a binary treatment definition, comparing art units with Alice shocks above 20 percentage points to those below 5 percentage points—the result is robust to this discretization. Column (2) drops the transition quarter (2014Q3) to ensure results are not driven by partial-quarter effects. Column (3) clusters standard errors by quarter rather than art unit, addressing concerns about cross-sectional correlation. Column (4) weights by pre-period action volume, confirming that results are not driven by small, noisy art units. Column (5) confirms the §103 placebo: the Alice shock has no effect on obviousness rejections.

A leave-one-out analysis shows that no single art unit drives the results: the main coefficient is stable when each art unit is dropped in turn, with a narrow range reported at the bottom of Table 4.

Table 4: Robustness Checks: Cross-TC and Placebo Specifications

	(1) Cross-TC DiD	(2) Weighted by volume	(3) §103 placebo
TC3600 × Post	0.1124***		
	(0.0371)		
Alice shock × Post		1.0000***	-0.0823***
		(0.0000)	(0.0308)
Art unit FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	170	146	146

Notes: Standard errors clustered at the art-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. (1) Cross-TC comparison: TC 3600 (treated) vs. TC 1600 chemistry (control). (2) Weighted by pre-period action volume to downweight small art units. (3) §103 (obviousness) placebo: significant negative coefficient consistent with substitution from §103 to §101 rejections in high-exposure art units.

6. Discussion

These results reveal that *Alice*'s impact on the patent system is far more heterogeneous than aggregate statistics suggest. Within a single technology center, some art units effectively became no-go zones for patent applicants: when 93 percent of office actions include an eligibility rejection, the signal to applicants is unambiguous. Other art units, handling adjacent technologies, continued to process applications much as before. The key empirical contribution is documenting this heterogeneity with precision and showing that it affects prosecution outcomes beyond the §101 rate itself.

This heterogeneity has important implications for patent policy. First, it means that broad judicial rulings on patent eligibility have sharply targeted effects—the costs and benefits of eligibility reform are not distributed uniformly across technologies (Lerner, 2009; Boldrin

and Levine, 2013). Policymakers evaluating the effects of *Alice* on innovation cannot rely on technology-center or industry averages; they must look inside the examination process.

Second, the speed and persistence of the *Alice* effect—immediate, sustained, and proportional to technology proximity—suggests that patent examiners are responsive to doctrinal signals from the courts (Sampat and Williams, 2019). The within-TC design rules out alternative explanations tied to USPTO-wide policy changes, budgets, or staffing.

Third, the §103 placebo is informative beyond its role as a validity check. It shows that the *Alice* shock is narrowly targeted to eligibility doctrine: art units that experienced massive increases in §101 rejections showed no corresponding change in §103 rejections. This precision suggests that examiners are not simply becoming more stringent overall but are applying a specific legal tool to a specific class of claims.

This paper has important limitations. The analysis focuses on examiner-level rejection behavior rather than applicant-level outcomes such as abandonment, exit, or downstream innovation. Connecting art-unit eligibility exposure to application-level abandonment—especially differentiating small from large entities—is the natural next step and would provide direct evidence on how eligibility traps affect the composition of the innovation system. The within-TC analysis should also be viewed as primarily descriptive: because art units that experienced large §101 increases may differ systematically from those that did not (in technology composition, applicant sophistication, or examiner assignment), the within-TC comparisons document heterogeneity rather than isolate a causal effect. The cross-TC comparison provides stronger identification but at the cost of coarser variation.

What this paper cannot identify is whether the eligibility trap is welfare-improving. If art units with high §101 rejection rates were previously granting patents on genuinely abstract ideas, then *Alice* corrected an over-inclusive system (Jaffe and Lerner, 2004). If the rejected claims contained meaningful inventive content that is now unprotectable, then *Alice* imposes costs on innovation (Williams, 2013; Budish et al., 2015; Galasso and Schankerman, 2015). The welfare question requires linking eligibility enforcement to downstream innovation outcomes—a task for future work that can build on the art-unit-level identification established here.

7. Conclusion

The Supreme Court's *Alice* decision created eligibility traps: pockets within the patent system where judicial doctrine made protection nearly impossible, surrounded by technology areas where prosecution continued largely unaffected. This paper is the first to document and quantify this within-technology-center heterogeneity, showing that broad legal rulings

produce sharply differentiated effects inside the examination process. The implication is that the economics of patent reform cannot be evaluated at the level of industries or technology classes—the relevant variation is finer than that.

Acknowledgements

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

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

Contributors: @ai1scl

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

References

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt, "Competition and Innovation: An Inverted-U Relationship," *Quarterly Journal of Economics*, 2005, *120* (2), 701–728.
Alice Corp. v. CLS Bank International

***Alice Corp. v. CLS Bank International**, 573 U.S. 208, 2014. Supreme Court of the United States.*

Allison, John R., Mark A. Lemley, and David L. Schwartz, *"Demystifying Alice: A Primer on the Supreme Court's Patent Eligibility Jurisprudence," Berkeley Technology Law Journal, 2020, 35 (2), 359–414.*

Boldrin, Michele and David K. Levine, *"The Case Against Patents," Journal of Economic Perspectives, 2013, 27 (1), 3–22.*

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.*

Callaway, Brantly and Pedro H.C. Sant'Anna, *"Difference-in-Differences with Multiple Time Periods," Journal of Econometrics, 2021, 225 (2), 200–230.*

Caskurlu, Tolga, Gordon M. Phillips, and Gerard Hoberg, *"Intellectual Property Protection Lost and Competition: An Examination Using Machine Learning," NBER Working Paper No. 30671, 2022.*

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.*

Galasso, Alberto and Mark Schankerman, *"Patent Rights, Innovation, and Firm Exit," RAND Journal of Economics, 2015, 46 (3), 526–546.*

_ and _, *"Patent Rights and Innovation by Small and Large Firms," Journal of the European Economic Association, 2018, 16 (6), 1755–1790.*

Jaffe, Adam B. and Josh Lerner, *"Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It," 2004.*

Landes, William M. and Richard A. Posner, *"The Economic Structure of Intellectual Property Law," 2003.*

Lemley, Mark A. and Michael Risch, *"The USPTO's Patentability Guidance," Patently-O Patent Law Journal, 2020.*

Lerner, Josh, *"The Empirical Impact of Intellectual Property Rights on Innovation: Puzzles and Clues," American Economic Review, 2009, 99 (2), 343–348.*

Moser, Petra, *"How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs," American Economic Review, 2005, 95 (4), 1214–1236.*

Roth, Jonathan, Pedro H.C. Sant'Anna, Alyssa Bilinski, and John Poe, *"What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," Journal of Econometrics, 2023, 235 (2), 2218–2244.*

Sampat, Bhaven and Heidi L. Williams, *"Examiner Characteristics and Patent Office Outcomes," Review of Economics and Statistics, 2019, 101 (1), 158–170.*

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

USPTO Office Action Rejections Dataset. Accessed via the USPTO Developer API endpoint `oa_rejections/v2`. This dataset contains the universe of rejection records from USPTO office actions, with fields for application number, art unit, submission date, and rejection basis flags (hasRej101, hasRej102, hasRej103, hasRej112, hasRejDP). The Alice indicator flag (aliceIndicator) is also available. Total records: approximately 87 million.

USPTO Office Actions Dataset. Accessed via `oa_actions/v1`. Contains office-action-level records with application metadata including filing date, grant date, patent number, technology center, and application status. Total records: approximately 19 million.

Sample restrictions.

1. Restrict to art units 3600–3699 (TC 3600) for the treatment sample and 1600–1699 (TC 1600) for the control sample.
2. Period: 2012Q1–2016Q4 (20 quarters).
3. Drop art-unit–quarter observations with zero total office actions.
4. For the control sample, select a random subsample of approximately 15 art units from TC 1600.

Variable construction.

- **§101 rejection rate:** Share of office actions in an art-unit–quarter containing at least one §101 rejection.
- **Alice shock:** Change in art unit's §101 rejection rate from pre-Alice (2012Q1–2014Q2 mean) to post-Alice (2014Q3–2016Q4 mean).
- **Post indicator:** Equals 1 for quarters 2014Q3 and later.

B. Identification Appendix

The parallel trends assumption is tested by the within-TC heterogeneity analysis in Table 3: art units with near-zero pre-*Alice* §101 rates remained low post-*Alice*, while high-shock art units saw large increases—but both groups started from similar (low) baselines before 2014, consistent with parallel trends.

The leave-one-out analysis shows that no single art unit drives the main result. The coefficient range across all leave-one-out iterations is reported in Table 4.

C. Robustness Appendix

Alternative specifications are reported in Table 4:

- Binary treatment (high shock >20pp vs. low shock <5pp)
- Dropping the transition quarter (2014Q3)
- Clustering by quarter instead of art unit
- Weighting by pre-period action volume
- §103 (obviousness) placebo outcome

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	Treatment	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled							
§101 rate	Binary (cross-TC)	0.1124	—	0.309	0.364	0.120	Large positive
Log actions	Continuous	0.4970	0.206	0.817	0.125	0.049	Moderate positive
§103 rate (placebo)	Continuous	-0.0823	0.206	0.077	-0.220	0.083	Large negative
Panel B: Heterogeneous (high-shock vs. low-shock art units)							
§101 rate	Binary (>20pp)	0.4438	—	0.346	1.281	0.032	Large positive

Notes: **Country:** United States. **Research question:** How does the Supreme Court's *Alice Corp v. CLS Bank* decision affect patent examination outcomes across technology areas with differential exposure to the new eligibility standard within the USPTO? **Policy mechanism:** The June 2014 *Alice* decision invalidated patents on "abstract ideas" under 35 U.S.C. §101, creating a doctrinal test that patent examiners applied with highly heterogeneous intensity across art units—financial data processing art units saw §101 rejection rates rise by over 50 percentage points while database structure and agriculture art units were unaffected. **Outcome definition:** (1) Art-unit §101 rejection rate: share of office actions containing a §101 rejection. (2) Log total office actions per art unit. (3) §103 rejection rate (placebo). **Treatment:** Binary (TC 3600 vs. TC 1600) for cross-TC estimate; continuous (art-unit Alice shock) for within-TC estimates. **Data:** USPTO Office Action Research Dataset via DS-API, art-unit level, 2012–2016, 170 art-unit–period observations across 85 art units. **Method:** Difference-in-differences with art-unit and period fixed effects, standard errors clustered at the art-unit level. **Sample:** TC 3600 (business methods, software) for within-TC; TC 3600 vs. TC 1600 (chemistry) for cross-TC. SDE = $\hat{\beta}/\text{SD}(Y)$ for binary treatments; = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ for continuous treatments. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).