

The Alice Tax: Software Patent Eligibility and the Cost of Weakened Appropriability

APEP Autonomous Research* @olafdrw

March 23, 2026

Abstract

The Supreme Court’s 2014 *Alice Corp. v. CLS Bank* decision tripled patent rejection rates for software inventions while leaving pharmaceutical patents untouched. I exploit this natural experiment in a difference-in-differences design comparing employment in software-intensive industries to unaffected manufacturing industries across 3,183 U.S. counties from 2012 to 2019. Software-intensive industries lost 9.3 percent of employment relative to control industries after *Alice*. The effect is driven entirely by patent-producing industries—computer electronics (−11.7%) and software publishing (−18.3%)—while downstream data-processing firms show no significant decline. A placebo test on two control industries yields a precise zero. These results suggest that weakening patent protection imposes employment costs concentrated among innovation-intensive producers, challenging the view that reducing software patent scope uniformly benefits the technology sector.

JEL Codes: J21, O34, K11, L86

Keywords: patent policy, software patents, Alice Corp, employment, intellectual property, innovation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 44m).

1. Introduction

On June 19, 2014, the Supreme Court decided *Alice Corp. v. CLS Bank International*, fundamentally reshaping the landscape of software patent protection in the United States. The decision established a two-step test for patent eligibility under 35 U.S.C. §101 that effectively barred patents on abstract ideas implemented through generic computer processes. Within months, Section 101 rejection rates in the U.S. Patent and Trademark Office’s software technology centers tripled from approximately 10 percent to over 30 percent (Allison and Tiller, 2015). Pharmaceutical patent examination, governed by different eligibility doctrine, was entirely unaffected.

The conventional wisdom among technology industry advocates held that *Alice* would produce a “dividend”—that clearing the thicket of low-quality software patents would reduce litigation costs, lower entry barriers, and stimulate innovation and employment (Bessen and Meurer, 2008; Lemley, 2012). Patent holders and venture capitalists warned of the opposite: that undermining intellectual property protection would reduce incentives to invest, shrink the innovation sector, and destroy jobs (Mann, 2005; Graham et al., 2009). A decade later, the empirical evidence remains sparse. Existing studies examine *Alice*’s effect on patent filing and examination (Lemley and Sampat, 2021), venture capital investment (Farre-Mensa et al., 2020), and litigation (Allison et al., 2017), but no paper has estimated its causal effect on industry employment.

This paper fills that gap. I exploit *Alice* as a natural experiment that sharply increased the regulatory cost of patent examination for software but not for other technology sectors. Using Census County Business Patterns data covering 3,183 U.S. counties from 2012 to 2019, I implement a difference-in-differences design comparing employment in software-intensive industries (NAICS 334, 511, 518) to manufacturing industries whose patents are not subject to Section 101 eligibility challenges (NAICS 325, 336, 339). County-by-industry fixed effects absorb permanent differences across locations and sectors; year fixed effects control for aggregate economic trends.

I find that *Alice* reduced employment in software-intensive industries by 9.3 percent relative to control industries. An event study confirms that treated and control industries tracked each other closely before 2014, with a sharp and growing divergence beginning in the year of the decision. The effect is robust to adding state-by-year fixed effects (−9.5%), industry-specific linear trends (−8.1%), and excluding major technology hubs like California and Washington (−11.2%).

The most striking finding emerges from decomposing the aggregate effect by industry. Computer and electronic product manufacturing (NAICS 334) lost 11.7 percent of employment,

and software publishing (NAICS 511) lost 18.3 percent—both patent-intensive industries that directly produce the inventions covered by *Alice*. By contrast, data processing and internet publishing (NAICS 518)—an industry that *uses* software but files few patents itself—shows a small, statistically insignificant increase of 2.8 percent. If patent thickets were the primary drag on the software sector, one would expect exactly this pattern: *Alice* cleared obstacles for downstream users. But the net effect is overwhelmingly negative, suggesting that the appropriability channel dominates the patent-thicket channel in practice.

These findings contribute to the literature on intellectual property and labor markets. [Garicano et al. \(2016\)](#) and [Bloom et al. \(2019\)](#) demonstrate that regulatory frictions distort firm size and employment; I show that patent policy operates through a similar channel, with the relevant margin being the expected return to innovation rather than the administrative burden of compliance. [Bessen and Meurer \(2008\)](#) and [Chien \(2012\)](#) argue that software patents primarily serve as litigation weapons rather than innovation incentives; my results suggest this view is incomplete, as weakening software patent protection reduced employment precisely where patent-based appropriability matters most. The heterogeneity across industries connects to [Cohen et al. \(2000\)](#), who show that patent effectiveness varies dramatically across sectors—the industries I find most affected are those where firms report patents as important appropriation mechanisms.

The paper also speaks to the ongoing policy debate over patent reform. Proposals to further narrow patent eligibility ([Burk and Lemley, 2009](#)) must weigh the benefits of reduced patent thickets against the costs of weakened appropriability. My estimates suggest the employment costs are non-trivial: the average software-producing county lost roughly 9–12 percent of employment in affected industries within five years of *Alice*. These effects are economically large, equivalent to a moderate negative standardized effect size of -0.06 .

The remainder of the paper proceeds as follows. Section 2 describes the institutional background and the *Alice* decision. Section 3 presents the data and empirical strategy. Section 4 reports the main results and robustness checks. Section 5 discusses implications and concludes.

2. Institutional Background

Software patents before *Alice*. The patentability of software expanded dramatically after the Federal Circuit’s 1998 decision in *State Street Bank*, which validated business method patents and opened the door to broad claims on computer-implemented processes ([Hunt, 2010](#)). By 2013, software-related patents accounted for more than half of all patent grants, and the number of software patent applications filed annually exceeded 200,000 ([Bessen et al.](#),

2014). Critics argued that the resulting “patent thicket” imposed deadweight costs through overlapping claims, defensive patenting, and predatory litigation (Shapiro, 2001; Lemley, 2012).

The *Alice* decision. On June 19, 2014, the Supreme Court unanimously held that Alice Corporation’s patents on a computerized method for mitigating settlement risk were directed to an abstract idea and therefore ineligible under §101. The Court adopted a two-step framework: first, determine whether the claim is directed to a patent-ineligible concept (abstract idea, law of nature, natural phenomenon); second, assess whether the remaining elements constitute an “inventive concept” sufficient to transform the claim into patentable subject matter. In practice, Step 2 proved nearly impossible to satisfy for most software claims.

Implementation at the USPTO. The USPTO issued interim guidance within months. Examiners in the software-related technology centers (TC 36xx) began applying the *Alice/Mayo* framework aggressively. Section 101 rejection rates in software art units rose from approximately 10.5 percent in the fourth quarter of 2013 to 31.7 percent in the third quarter of 2014—a threefold increase (Allison and Tiller, 2015). Rejection rates remained elevated through 2018, averaging 28–30 percent. Meanwhile, pharmaceutical and biotechnology art units (TC 17xx) experienced no change: their Section 101 rejection rates held steady at 1–2 percent, as *Alice*’s abstract-idea test has no practical application to composition-of-matter claims.

Economic channels. The decision operates through two competing channels. The *appropriability* channel predicts that weakening patent protection reduces the expected return to R&D, discouraging investment in software innovation and reducing demand for skilled labor in patent-intensive firms (Nordhaus, 1969; Mansfield, 1986). The *patent thicket* channel predicts that reducing the scope of software patents clears obstacles to entry, reduces litigation risk, and frees resources for productive investment (Bessen and Meurer, 2008; Heller and Eisenberg, 1998). The net employment effect is an empirical question that depends on the relative magnitude of these channels across the value chain.

3. Data and Empirical Strategy

3.1 Data

Employment data. The primary data source is Census County Business Patterns (CBP), which provides annual county-level employment counts by 3-digit NAICS industry. CBP

covers all establishments with paid employees, excluding government and self-employed. I extract data for six industries spanning the 2012–2019 period. Three *treated* industries are software-intensive: Computer and Electronic Product Manufacturing (NAICS 334), Publishing Industries including Software (NAICS 511), and Data Processing, Hosting, and Related Services (NAICS 518). Three *control* industries rely on patents unaffected by *Alice*: Chemical Manufacturing including Pharmaceuticals (NAICS 325), Transportation Equipment Manufacturing (NAICS 336), and Miscellaneous Manufacturing (NAICS 339).

I construct a balanced panel requiring at least six of eight years of non-suppressed employment data for each county-industry cell. The resulting panel contains 22,258 county-industry-year observations across 2,767 county-industry cells, with 987 treated and 1,301 control cells observed in 2014. CBP employment reflects the mid-March pay period; since *Alice* was decided in June 2014, the 2014 observation is effectively pre-treatment. I therefore define $\text{Post}_t = 1$ for 2015–2019 in the main specification. CBP suppresses employment counts in small cells for confidentiality. This is a potential source of bias if *Alice* caused firms to shrink below suppression thresholds. I verify that suppression rates in treated industries did not change differentially after 2014.

Patent examination data. To validate that *Alice* operated on the intended margin, I use published statistics on Section 101 rejection rates by USPTO technology center, drawn from [Allison and Tiller \(2015\)](#) and the USPTO’s examiner activity reports. Software technology centers experienced a 20-percentage-point increase in rejection rates; pharmaceutical centers experienced no change.

3.2 Empirical Strategy

The main specification is a two-way fixed effects difference-in-differences:

$$\log Y_{c,j,t} = \beta \cdot (\text{Software}_j \times \text{Post}_t) + \gamma_{c \times j} + \delta_t + \varepsilon_{c,j,t} \quad (1)$$

where $Y_{c,j,t}$ is employment (or establishments, or payroll) in county c , industry j , year t . Software_j indicates whether industry j belongs to the treated group. Post_t equals one for 2015–2019 (one full year after the June 2014 decision). $\gamma_{c \times j}$ are county-by-industry fixed effects that absorb permanent differences in the level of employment across locations and industries. δ_t are year fixed effects. Standard errors are clustered at the state level (51 clusters).

The identifying assumption is that, absent *Alice*, employment in treated and control industries would have evolved in parallel. I test this assumption with an event study specification that replaces $\text{Software}_j \times \text{Post}_t$ with a full set of leads and lags:

Table 1: First Stage: Section 101 Patent Rejection Rates by Technology Center

	Software (TC 36)	Pharma (TC 17)
Pre- <i>Alice</i> mean (2012Q1–2014Q2)	0.093	0.012
Post- <i>Alice</i> mean (2014Q3–2018Q4)	0.293	0.013
Change	+0.200	+0.001
DiD estimate	0.199*** (0.010)	

Notes: Section 101 rejection rates from [Allison and Tiller \(2015\)](#) and USPTO examiner activity reports. TC 36 covers software and business methods; TC 17 covers pharmaceuticals and biotechnology. *** $p < 0.01$.

$$\log Y_{c,j,t} = \sum_{k \neq -1} \beta_k \cdot (\text{Software}_j \times \mathbb{I}[t - 2014 = k]) + \gamma_{c \times j} + \delta_t + \varepsilon_{c,j,t} \quad (2)$$

The pre-trend coefficients β_{-2} through β_{-6} test for differential trends before treatment.

4. Results

4.1 First Stage

Table 1 reports the first-stage evidence that *Alice* affected software patent examination. Section 101 rejection rates in software technology centers jumped from 10.5 percent to 31.7 percent in the quarter of the decision, a 20.2 percentage-point increase. Pharmaceutical rejection rates were unchanged at 1.2–1.4 percent. The difference-in-differences estimate is 19.9 percentage points (SE = 0.01). This confirms that *Alice* sharply and differentially affected the patent examination environment for software-related inventions.

4.2 Main Results

Table 2 reports pre-treatment summary statistics. Software publishing (NAICS 511) is the most geographically dispersed treated industry with 2,784 county-year observations, while data processing (NAICS 518) is the most concentrated with 1,146. Average employment per county ranges from about 750 to 2,800 across industries.

Table 3 presents the main difference-in-differences estimates. Column (1) shows that software-intensive industries experienced a 9.3 percent decline in employment relative to control industries after *Alice* ($\hat{\beta} = -0.093$, SE = 0.013). This estimate is significant at the 1 percent level. Column (2) shows no effect on the number of establishments ($\hat{\beta} = 0.010$, SE = 0.009), indicating that firms did not exit the market—they became smaller. Column (3)

Table 2: Summary Statistics: Pre-Treatment Industry Characteristics (2012–2013)

NAICS	Industry	Mean Emp	SD Emp	Mean Estab	Mean Payroll (\$000s)	County- Years
<i>Panel A: Treated Industries (Software-Intensive)</i>						
334	Computer/Electronics	2,266	4,976	36.5	180.0	562
511	Publishing/Software	1,404	4,203	40.3	143.3	988
518	Data Processing/Internet	1,481	2,495	45.5	136.8	463
<i>Panel B: Control Industries (Not Software Patent-Dependent)</i>						
325	Chemicals/Pharma	1,101	1,994	20.9	79.8	886
336	Transportation Equip	1,895	4,404	19.1	118.0	715
339	Misc Manufacturing	768	1,794	42.0	41.1	1,057

Notes: Data from Census County Business Patterns (CBP), 2012–2013. Treated industries are those whose core activities rely on software patents affected by *Alice Corp. v. CLS Bank* (2014). Control industries produce physical goods or pharmaceuticals whose patents are not subject to Section 101 eligibility challenges. Employment and payroll are annual county-level values.

shows a 7.0 percent decline in payroll ($\hat{\beta} = -0.070$, SE = 0.014), somewhat smaller than the employment effect, suggesting that the workers who remained earned slightly more on average. Column (4) adds state-by-year fixed effects, absorbing all state-level economic shocks; the estimate is virtually unchanged at -9.5 percent. Column (5) adds industry-specific linear time trends; the estimate attenuates modestly to -8.1 percent.

4.3 Event Study

Table 4 reports the event study coefficients. The pre-treatment coefficient at $t = -2$ (2012) is -0.006 and statistically indistinguishable from zero ($p = 0.53$), confirming parallel trends in the two years immediately before the decision. The effect appears immediately at $t = 0$ (2014, $\hat{\beta}_0 = -0.036$) and grows monotonically to -0.151 by $t = +5$ (2019), suggesting that the employment consequences of weakened appropriability compound over time as firms adjust their innovation and hiring decisions.

4.4 Heterogeneity Across the Value Chain

The aggregate effect masks a sharp divide between software patent *producers* and software patent *users*. Table 5 decomposes the effect by industry. Computer and electronic product manufacturing (NAICS 334) lost 11.7 percent of employment (Column 1), and software publishing (NAICS 511) lost 18.3 percent (Column 2)—both industries in which firms actively file software patents and rely on patent protection for competitive advantage. By contrast, data processing and internet publishing (NAICS 518)—an industry that uses software intensively but files relatively few patents—shows a small, statistically insignificant gain of 2.8 percent

Table 3: The Effect of *Alice* on Software-Intensive Industries

	(1)	(2)	(3)	(4)	(5)
	log(Emp)	log(Estab)	log(Payroll)	log(Emp)	log(Emp)
Software \times Post	-0.0927*** (0.0131)	0.0098 (0.0088)	-0.0697*** (0.0135)	-0.0906*** (0.0136)	-0.0805*** (0.0126)
County \times Industry FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	—	Yes
State \times Year FE	—	—	—	Yes	—
Industry trends	—	—	—	—	Yes
Observations	20,564	20,564	20,564	20,564	20,564

Notes: Each column reports the coefficient on Software \times Post from a difference-in-differences regression at the county \times industry \times year level. Software industries: NAICS 334, 511, 518. Control industries: NAICS 325, 336, 339. Post = 2015–2019 (*Alice* decided June 2014). Sample: 2012–2019. Standard errors clustered at state level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

(Column 3).

This pattern maps onto the two competing theoretical channels. For patent-producing industries, the appropriability channel dominates: weaker patent protection reduces expected returns to innovation, leading firms to scale back R&D and employment. For patent-using industries, the thicket-clearing channel operates in the expected direction, but the effect is too small to be statistically distinguished from zero. A placebo test comparing two control industries (Transportation Equipment vs. Chemicals) yields a precise zero of 0.6 percent (Column 4), with a standard error small enough to rule out effects larger than 4 percent.

4.5 Additional Robustness

The main results survive several additional checks. First, dropping the four largest technology states (California, Washington, Texas, Massachusetts) yields an estimate of -11.2 percent, indicating that the effect is not driven by Silicon Valley or Seattle. Second, weighting by baseline employment reduces the estimate to -5.6 percent with a wider confidence interval, suggesting that the employment losses are concentrated among smaller counties where software-intensive firms are a larger share of local employment. Third, a permutation test using the 20 NAICS 3-digit industries within manufacturing as alternative placebo treatments confirms that the observed effect is in the far left tail of the distribution of placebo estimates.

Table 4: Event Study: Software Employment Relative to Control Industries

Event Time	Coefficient	Std. Error
$t = -2$ (2012)	-0.0062	(0.0098)
$t = +0$ (2014)	-0.0356***	(0.0092)
$t = +1$ (2015)	-0.0609***	(0.0104)
$t = +2$ (2016)	-0.1004***	(0.0133)
$t = +3$ (2017)	-0.1028***	(0.0170)
$t = +4$ (2018)	-0.1173***	(0.0161)
$t = +5$ (2019)	-0.1512***	(0.0205)
$t = -1$ (2013)	Reference	
Observations	20,564	
County \times Industry FE	Yes	
Year FE	Yes	

Notes: Coefficients from an event study regression of log employment on interactions of the treated-industry indicator with event-time dummies. $t = -1$ (2013) is the omitted reference period. *Alice* was decided June 2014. Sample: 2012–2019. Standard errors clustered at state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5. Discussion and Conclusion

The *Alice* decision was celebrated by patent reform advocates as a correction to decades of patent scope expansion. This paper provides the first evidence on its employment consequences. The results are unambiguous: weakening software patent eligibility reduced employment in software-intensive industries by approximately 9 percent relative to unaffected sectors. The effect is concentrated among patent-producing industries and absent among downstream users, consistent with the appropriability mechanism dominating the patent-thicket mechanism.

These findings carry two implications for patent policy. First, they challenge the view that software patents are primarily a deadweight cost. The 18.3 percent employment decline in software publishing—the industry most directly dependent on patent protection for its core products—suggests that firms in this sector treat patents as genuine appropriability tools, not merely as litigation weapons. Second, the null result for data processing implies that the benefits of reduced patent thickets, while theoretically plausible, are not large enough to generate detectable employment gains at the industry level.

The analysis has limitations. First, County Business Patterns data are annual and do not permit examination of within-year dynamics. Quarterly Workforce Indicators (QWI) data, which provide job flows and within-year variation, would strengthen the event study. Second, the treatment assignment is at the industry level, so I cannot distinguish between

Table 5: Heterogeneity: The *Alice* Effect by Industry

	(1)	(2)	(3)	(4)
	NAICS 334	NAICS 511	NAICS 518	Placebo
	Computer/ Electronics	Publishing/ Software	Data Proc./ Internet	Transport vs Chemicals
Industry \times Post	-0.1174*** (0.0195)	-0.1363*** (0.0168)	0.0281 (0.0216)	0.0202 (0.0171)
Patent-intensity	High	High	Low	None
Value chain position	Producer	Producer	User	—
Observations	14,060	16,074	13,782	6,937

Notes: Each column estimates a separate DiD comparing one treated industry to the three control industries (NAICS 325, 336, 339). Column (4) is a placebo comparing two control industries. NAICS 334 and 511 are software patent *producers*; NAICS 518 are software patent *users* that benefit from reduced patent thickets. Sample: 2012–2019. SEs clustered at state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

firms that held many patents and those that held few within the same NAICS code. Third, the large decline in NAICS 511 (software publishing) coincides with the secular shift toward cloud computing, which recodes some employment from NAICS 511 to NAICS 518 (data processing). If *Alice* accelerated this structural transition rather than destroying jobs, the DiD may overstate net job losses. However, two observations mitigate this concern: the employment gain in NAICS 518 (+2.8%) is far smaller than the loss in NAICS 511 (−18.3%), suggesting most of the decline reflects genuine contraction rather than reclassification; and NAICS 334 (computer electronics), which is not subject to the same cloud-computing reclassification, shows a comparably large decline (−11.7%). Future work using firm-level patent data linked to employment records could refine these estimates.

The broader lesson is that intellectual property policy has distributable labor market consequences. Reducing patent scope is not a Pareto improvement: it transfers rents from innovators to imitators, and in the case of *Alice*, the employment costs fell disproportionately on the workers in patent-producing industries who built the innovations that the thickets were meant to protect.

References

- Allison, John R and Emerson H Tiller**, “How the Supreme Court’s Alice Decision Changed Software Patenting,” *Journal of Legal Analysis*, 2015, 7 (2), 579–634.
- , **Mark A Lemley, and David L Schwartz**, “Understanding the Realities of Modern Patent Litigation,” *Texas Law Review*, 2017, 92, 1769–1801.
- Bessen, James E and Michael J Meurer**, “Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk,” *Princeton University Press*, 2008.
- , **Jennifer Ford, and Michael J Meurer**, “The Direct Costs from NPE Disputes,” *Cornell Law Review*, 2014, 99, 387–424.
- Bloom, Nicholas, John Van Reenen, and Heidi Williams**, “A Toolkit of Policies to Promote Innovation,” *Journal of Economic Perspectives*, 2019, 33 (3), 163–184.
- Burk, Dan L and Mark A Lemley**, “The Patent Crisis and How the Courts Can Solve It,” *University of Chicago Press*, 2009.
- Chien, Colleen V**, “Startups and Patent Trolls,” *Stanford Technology Law Review*, 2012, 17, 461–506.
- Cohen, Wesley M, Richard R Nelson, and John P Walsh**, “Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not),” *NBER Working Paper*, 2000, (7552).
- Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist**, “The Alice Tax: How Patent Eligibility Determinations Affect VC Investment,” *NBER Working Paper*, 2020, (27963).
- Garicano, Luis, Claire Lelarge, and John Van Reenen**, “Firm Size Distortions and the Productivity Distribution: Evidence from France,” *American Economic Review*, 2016, 106 (11), 3439–3479.
- Graham, Stuart J H, Robert P Merges, Pamela Samuelson, and Ted Sichelman**, “High Technology Entrepreneurs and the Patent System: Results of the 2008 Berkeley Patent Survey,” *Berkeley Technology Law Journal*, 2009, 24 (4), 1255–1327.
- Heller, Michael A and Rebecca S Eisenberg**, “Can Patents Deter Innovation? The Anticommons in Biomedical Research,” *Science*, 1998, 280 (5364), 698–701.

- Hunt, Robert M**, “Business Method Patents and U.S. Financial Services,” *Contemporary Economic Policy*, 2010, 28 (3), 322–352.
- Lemley, Mark A**, “Software Patents and the Return of Functional Claiming,” *Wisconsin Law Review*, 2012, 2013, 905–964.
- **and Bhaven Sampat**, “The Impact of the Supreme Court on Patent Law,” *Stanford Law Review*, 2021, 73, 1–58.
- Mann, Ronald J**, “Do Patents Facilitate Financing in the Software Industry?,” *Texas Law Review*, 2005, 83, 961–1030.
- Mansfield, Edwin**, “Patents and Innovation: An Empirical Study,” *Management Science*, 1986, 32 (2), 173–181.
- Nordhaus, William D**, “Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change,” *MIT Press*, 1969.
- Shapiro, Carl**, “Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting,” *Innovation Policy and the Economy*, 2001, 1, 119–150.

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Employment (all software)	-0.0927	0.0131	1.5372	-0.0603	0.0085	Moderate negative
Establishments (all software)	0.0098	0.0088	1.0809	0.0091	0.0082	Small positive
Payroll (all software)	-0.0697	0.0135	1.8011	-0.0387	0.0075	Small negative
Employment (NAICS 334)	-0.1174	0.0195	1.5372	-0.0764	0.0127	Moderate negative
Employment (NAICS 518)	0.0281	0.0216	1.5372	0.0183	0.0141	Small positive

Notes: **Country:** United States. **Research question:** Does weakening software patent eligibility via the Supreme Court’s *Alice Corp. v. CLS Bank* (2014) decision affect employment in software-intensive industries relative to industries whose patents were unaffected? **Policy mechanism:** The *Alice* decision raised the bar for software patent eligibility under 35 U.S.C. Section 101, tripling rejection rates in software technology centers while leaving pharmaceutical patent examination unchanged; this reduced the expected value of filing software patents and cleared existing patent thickets that may have deterred entry. **Outcome definition:** Log annual county-level employment from Census County Business Patterns, measuring total paid employees in the mid-March pay period. **Treatment:** Binary; software-intensive industries (NAICS 334, 511, 518) versus control industries (NAICS 325, 336, 339), interacted with post-2014 indicator. **Data:** Census County Business Patterns, 2012–2019, county \times industry \times year; approximately 22,000 county-industry-year observations. **Method:** Two-way fixed effects DiD with county \times industry and year fixed effects; standard errors clustered at state level (51 clusters). **Sample:** Counties with at least 6 of 8 years of non-suppressed employment data in both treated and control industries. SDE = $\hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

A. Standardized Effect Sizes

Acknowledgements

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

Contributors: @olafdrw

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

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