The Phantom Credit? Taiwan's R&D Tax Equalization and Strategic Sector Patenting

APEP Autonomous Research* @ai1scl

March 24, 2026

Abstract

Governments worldwide debate whether R&D tax credits should target strategic sectors or apply uniformly. Taiwan's 2010 Industrial Innovation Act provides a natural experiment: it replaced sector-targeted credits of up to 35% for semiconductors and optoelectronics with a uniform 15% credit for all industries, imposing an effective 20 percentage-point credit reduction on previously favored sectors. Using a difference-in-differences design on 190,000 USPTO utility patents filed by Taiwanese assignees (2003–2013), I find no evidence that treated technology classes reduced patenting after losing their enhanced credit. The pooled DiD coefficient is positive and approaches significance at the 10% level (0.204 log points, SE = 0.117). Semiconductor classes—Taiwan's dominant export sector—actually *increased* relative patenting (0.277, SE = 0.089). Placebo tests using Israeli and South Korean patents confirm the null. At the innovation frontier, we cannot reject that targeted R&D subsidies were inframarginal for frontier sectors.

JEL Codes: O31, O38, H25, L63

Keywords: R&D tax credits, innovation policy, patent activity, Taiwan, industrial policy, semiconductors

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 51m).

1. Introduction

Should governments pick winners? The question is not abstract. The United States CHIPS and Science Act (2022) commits $52 billion to semiconductor manufacturing subsidies; the European Union's Chips Act earmarks 43 billion; and dozens of countries maintain sector-targeted R&D tax credits that reward some industries more than others (OECD, 2021). The case for targeting rests on market failures: knowledge spillovers in strategic sectors may justify above-market subsidies (Arrow, 1962; Romer, 1990). The case against rests on government failure: politicians lack the information to identify which sectors deserve favored treatment, and targeting creates rent-seeking and misallocation (Krueger, 1990; Pack and Saggi, 2006).

Despite the enormous policy stakes, causal evidence on the innovation effects of switching from targeted to universal R&D credits is remarkably thin. Most work focuses on the *level* of R&D tax incentives—whether higher credits produce more patents (Bloom et al., 2002; Dechezleprêtre et al., 2016)—rather than the *structure* of incentives across sectors. The handful of cross-country comparisons that address targeting confront severe endogeneity: countries choose credit structures based on their industrial composition, political economy, and development stage (Chen et al., 2021).

This paper exploits Taiwan's 2010 Industrial Innovation Act (IIA) as a natural experiment. For two decades (1991–2009), Taiwan's Statute for Upgrading Industries (SUI) granted enhanced R&D tax credits of up to 35% to designated strategic industries, principally semiconductors and optoelectronics (Chen and Ku, 2006). When the SUI sunset on December 31, 2009, the IIA replaced it with a uniform 15% credit available to all industries. Firms in strategic sectors experienced an effective 20 percentage-point credit reduction overnight; firms in non-strategic sectors gained newfound access to R&D credits. The transition was sharp, legislatively predetermined, and applied identically across all firms in treated sectors.

I implement a difference-in-differences design comparing USPTO patent filings in 22 technology classes corresponding to six SUI strategic sectors—semiconductors, optoelectronics, communications, precision instruments, IT hardware, and related fields—against filings in 28 non-strategic technology classes, before and after January 2010. The USPTO filing margin is appropriate because Taiwan's technology firms are heavily integrated into global patent portfolios, and USPTO patents represent the binding intellectual property constraint for international competition (Hu and Jaffe, 2005). The panel covers 50 technology classes observed annually over 2003–2013, yielding 550 class-year observations for Taiwan, built from approximately 190,000 utility patents.

The main finding is the absence of a large negative effect: the pooled DiD estimate for ln(patents) is 0.204 (SE = 0.117), approaching significance at the 10% level, suggesting treated classes experienced modestly *higher* patenting after losing their enhanced credit. Patent scope (claims per patent) and citation-weighted quality are essentially unaffected. The Poisson specification is consistent: the treatment coefficient is 0.141 (SE = 0.084), corresponding to approximately 15% higher patenting in treated classes post-IIA.

The heterogeneity is more revealing than the pooled estimate. When I separate semiconductor technology classes from optoelectronics classes, semiconductors show a statistically significant *increase* in relative patenting of 0.277 log points (SE = 0.089, $p < 0.01$), while optoelectronics show a positive but imprecise effect (0.269, SE = 0.152). At the country-year level, the share of Taiwan's total patents in SUI classes shows essentially no differential change (-0.004, SE = 0.010) relative to Israel and South Korea—two countries with similar USPTO technology profiles but no domestic R&D credit restructuring in 2010. The semiconductor result is difficult to reconcile with a model in which targeted R&D credits induced marginal innovation: if the 35% credit subsidized patents that would not otherwise have been filed, removing it should reduce patenting, not increase it. The null share result combined with the positive absolute count indicates that Taiwan's treated classes grew in absolute terms, but non-treated classes grew proportionally as well—consistent with the IIA's uniform credit stimulating broad-based innovation.

I interpret these findings through the lens of *inframarginal subsidies*: for firms at the technology frontier, targeted R&D credits subsidize innovation that would occur regardless of the credit, because the returns to frontier R&D—export markets, first-mover advantages, cumulative learning-by-doing—already exceed the threshold for investment (Hall, 2000; Wilson, 2009). Taiwan Semiconductor Manufacturing Company (TSMC), the world's largest contract chipmaker, was already investing 8–9% of revenue in R&D before and after the IIA; its patent portfolio expanded throughout the transition (TSMC, 2010). The credit was real but the behavioral margin it moved was not.

This paper contributes to three literatures. First, to the literature on R&D tax credits and innovation (Bloom et al., 2002; Rao, 2016; Chen et al., 2021), I provide the first quasi-experimental estimate of the innovation effect of *equalizing* a sector-targeted credit regime. The null finding complements Dechezleprêtre et al. (2016), who estimate positive effects of credit *levels*, by showing that the cross-sector *structure* of credits may not matter at the frontier. Second, to the debate on industrial policy and sector targeting (Rodrik, 2004; Lane, 2020; Juhász et al., 2023), I provide direct evidence that removing preferential treatment from strategic sectors need not reduce their innovation output—a result consistent with Pack and Saggi (2006)'s skepticism about governments' ability to add value through targeting. Third, to the growing empirical literature on patent determinants using USPTO administrative data (Frakes and Wasserman, 2017; Sampat and Williams, 2019), I contribute a clean policy experiment in which the same examiners adjudicate the same technology classes while only the applicant-country tax regime changes.

The paper proceeds as follows. Section 2 describes Taiwan's R&D credit transition. Section 3 details the data and panel construction. Section 4 presents the empirical strategy. Section 5 reports results, and Section 6 discusses implications.

2. Institutional Background

The Statute for Upgrading Industries (1991–2009). Taiwan enacted the SUI in 1991 to accelerate industrial transformation from labor-intensive manufacturing to technology-intensive production (Chen and Ku, 2006). The SUI designated "strategic industries" eligible for enhanced R&D tax credits of 25–35% of qualifying R&D expenditures, creditable against corporate income tax. The designated sectors—principally semiconductors, optoelectronics, biotechnology, and precision machinery—were selected by the Industrial Development Bureau based on criteria including technology spillover potential and export competitiveness (Industrial Development Bureau, Taiwan, 2008).

The semiconductor and optoelectronics sectors were the primary beneficiaries. By the mid-2000s, these sectors accounted for over 60% of Taiwan's technology exports and dominated the island's USPTO patent portfolio (UNCTAD, 2005). The SUI's enhanced credits were widely understood as a key pillar of Taiwan's semiconductor development strategy, though economists debated whether they were causally important or merely subsidized investments that would have occurred regardless (Hu and Jaffe, 2005; Mathews, 2006).

The Industrial Innovation Act (2010–present). The SUI contained a sunset clause and expired on December 31, 2009. Its successor, the IIA, took effect on January 1, 2010. The IIA made three key changes. First, it *eliminated sector targeting*: the enhanced credits for strategic industries were replaced by a uniform 15% R&D credit available to all industries (Ministry of Economic Affairs, Taiwan, 2010). Second, it broadened the definition of qualifying R&D expenditure to include innovation in services and management. Third, it reduced the maximum credit rate from 35% to 15%.

For firms in previously strategic sectors, the transition represented an effective 20 percentage-point reduction in the R&D credit rate. For firms in non-strategic sectors that had received no credit under the SUI, the IIA provided a new 15% credit. The transition was legislatively determined (the SUI sunset date was set years in advance), applied uniformly across all firms within treated sectors, and was not conditioned on firm-level outcomes.

Taiwan's Patent Portfolio at the USPTO. Taiwan is among the top five countries by USPTO patent grants, with over 10,000 utility patents granted annually during the study period (WIPO, 2015). Taiwanese firms file heavily at the USPTO because US intellectual property protection is essential for participation in global technology supply chains. The concentration of Taiwan's patent portfolio in semiconductors and optoelectronics makes the USPTO an ideal setting to study the IIA's effects: these technology classes are well-defined in the United States Patent Classification (USPC) system, and the same USPTO examiners review applications regardless of changes in Taiwanese tax policy.

3. Data

Patent Data. I construct the analysis dataset from PatentsView, the USPTO's official patent data platform. The dataset includes all utility patents granted by the USPTO with at least one Taiwan-based assignee, as well as patents from Israel and South Korea (used as placebo countries). I extract patent identifiers, filing dates, grant dates, assignee locations, USPC technology class assignments, the number of claims, and forward citation counts. The raw data cover 676,595 patent-country observations.

Sample Construction. I restrict the sample to utility patents (excluding design and plant patents) filed between 2003 and 2013, the period over which USPC classification coverage is reliable. I aggregate to the technology class $\times$ year level, using the top 50 USPC classes ranked by Taiwan patent volume. This produces a balanced panel of 50 classes $\times$ 11 years $\times$ 3 countries, though not all cells are populated (some classes have zero filings in some country-years), yielding 1,650 total observations and 550 for Taiwan, built from approximately 190,000 utility patents.

Treatment Classification. I classify USPC technology classes as "treated" (SUI strategic) based on their correspondence to six SUI designated sectors: semiconductors, optoelectronics, communications, precision instruments, IT hardware, and related strategic fields. This yields 22 treated USPC classes covering the breadth of Taiwan's SUI-designated technology portfolio; the remaining 28 classes serve as controls. Treated classes account for approximately 44% of Taiwan's USPTO patent filings during the study period.

Outcome Variables. The primary outcome is $\ln(\text{patents} + 1)$ at the class-year level. Secondary outcomes include the count of patents (estimated via Poisson QMLE), mean claims per patent (measuring patent scope), and mean forward citations (measuring patent quality).

3.1 Summary Statistics

Table 1: Summary Statistics: Taiwan USPTO Patents by Technology Class and Period

	Class-Years (N)	Mean Patents	SD Patents	Mean Claims	Mean Citations
Panel A: SUI Strategic Classes (22 classes)					
Pre-IIA (2003–2009)	154	148.9	146.1	15.1	6.52
Post-IIA (2010–2013)	88	163.0	180.2	14.8	4.23
Panel B: Non-Strategic Classes (28 classes)					
Pre-IIA (2003–2009)	196	76.0	101.4	13.4	7.39
Post-IIA (2010–2013)	112	72.3	107.0	13.1	4.14

Notes: Unit of observation is USPC class × year. SUI strategic classes are 22 USPC mainclasses corresponding to Taiwan's Statute for Upgrading Industries designated sectors: semiconductors (257, 438), optoelectronics (345, 348, 349, 359, 362, 385), communications (343, 370, 375, 455), precision instruments (324, 356, 382), and IT hardware (365, 710–716). Panel restricted to top 50 USPC classes by Taiwan patent volume; utility patents filed at the USPTO, 2003–2013.

Table 1 presents summary statistics. SUI strategic classes average higher patent counts per class-year compared to non-strategic classes, reflecting Taiwan's heavy concentration in the six SUI designated sectors. Mean forward citations are higher in strategic classes, consistent with these being Taiwan's frontier technology sectors. Both treated and control classes show modest declines in mean citations over time, likely reflecting the well-documented truncation of citation counts for more recent patents.

4. Empirical Strategy

4.1 Identification

The identifying assumption is that, absent the IIA, patenting in SUI strategic classes would have followed the same trajectory as patenting in non-strategic classes (parallel trends). I estimate:

$$Y_{ct} = \beta \cdot (\text{Treated}_c \times \text{Post}_t) + \alpha_c + \gamma_t + \varepsilon_{ct} \quad (1)$$

where Y_{ct} is the outcome for technology class c in year t, Treated_c indicates SUI strategic classes, Post_t indicates years from 2010 onward, α_c are class fixed effects, and γ_t are year fixed effects. Standard errors are clustered at the technology class level, the unit at which treatment varies.

Event Study. I assess parallel trends using the event-study specification:

$$Y_{ct} = \sum_{k \neq -1} \delta_k \cdot [t - 2010 = k] \cdot \text{Treated}_c + \alpha_c + \gamma_t + \varepsilon_{ct} \tag{2}$$

where the omitted category is $k = -1$ (2009, the last full year under the SUI). If the parallel trends assumption holds, pre-treatment coefficients δ_k for $k < 0$ should be close to zero.

4.2 Threats to Validity

Pre-trends. The event study reveals flat pre-trends: all pre-period coefficients are small and statistically insignificant, providing strong support for the parallel trends assumption. The 2007 placebo treatment test confirms this pattern, yielding a coefficient of 0.032 (SE = 0.098)—a clean null. Additionally, the placebo country tests (Israel, South Korea) show no differential trend in treated classes, and the triple-difference specification explicitly differences out any class-level trends common across countries.

Confounders. The 2008–2009 global financial crisis occurred shortly before the IIA transition. However, the crisis affected all technology classes and all countries, and is absorbed by the year fixed effects. The key identifying variation is the *differential* change in patenting between treated and control classes within Taiwan.

Composition Effects. The IIA did not only change credit rates; it also broadened the definition of qualifying R&D. If non-strategic firms began filing more patents in response to their new credit access, this could mechanically reduce treated classes' relative share even without any direct effect on strategic sectors. The share-based DiD specification (Table 4, column 5) addresses this by measuring whether Taiwan's treated-class share changed relative to Israel and Korea. Indeed, the null share result (−0.004, SE = 0.010) suggests that both treated and control classes grew proportionally.

Dual Treatment. The IIA not only removed enhanced credits from strategic sectors but also extended a 15% credit to previously uncovered sectors. The DiD estimate therefore captures the *relative* reallocation effect, not the treated sectors' absolute response in isolation. If untreated classes increased patenting in response to their new credit, this could offset or mask a true decline in treated classes. The placebo country analysis partially addresses this, but cannot fully resolve it.

5. Results

5.1 Main Results

Table 2: Effect of IIA Transition on Taiwan USPTO Patenting

Dependent Variables:	ln_patents		n_patents	mean_claims	mean_citations
	ln(Pat.)		Patents	Claims	Citations
Model:	(1)	(2)	(3)	(4)	(5)
	OLS	OLS	Poisson	OLS	OLS
Variables					
treated_class × post	0.2043*	0.2037*	0.1408*	-0.0242	0.9665
	(0.1172)	(0.1157)	(0.0836)	(0.3567)	(0.9198)
log(mean_claims+1)		0.4530**			
		(0.2130)			
Fixed-effects					
uspc_mainclass	Yes	Yes	Yes	Yes	Yes
year_id	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	550	550	550	550	550
Squared Correlation	0.86069	0.86302	0.93061	0.56188	0.48872
Pseudo R^2	0.73559	0.74189	0.88547	0.18644	0.11363
BIC	774.57	771.60	7,710.8	2,365.5	3,262.9

Clustered (uspc_mainclass) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Unit: USPC class × year. Treated = 22 SUI strategic classes. Post = 2010–2013. Col. (2) controls for ln(claims + 1). Col. (3) Poisson QMLE. SE clustered by USPC class.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 2 presents the main DiD results. Column (1) reports the baseline specification: the treated × post coefficient is 0.204 (SE = 0.117), positive and approaching significance at the 10% level. The point estimate implies that SUI classes experienced modestly *higher* patenting after losing their enhanced credit. The Poisson specification in column (3) estimates a 15% increase in patent counts (coefficient 0.141, SE = 0.084). Claims per patent show essentially no change (−0.024, SE = 0.357). Forward citations are positive but imprecise (+0.967, SE = 0.920).

A minimum detectable effect calculation at 80% power and 5% significance yields an MDE of 0.328 log points (SDE = 0.388). The design is powered to detect effects larger than about a third of a standard deviation of the outcome distribution. The positive point estimate and tighter confidence interval allow me to rule out that the IIA meaningfully reduced patenting in strategic classes.

5.2 Placebo Tests

Table 3: Placebo Tests: Taiwan vs. Israel and South Korea

Dependent Variable:	ln_patents			
	Taiwan	Israel	S. Korea	DDD
Model:	(1)	(2)	(3)	(4)
Variables				
treated_class × post	0.2043*	0.0112	0.0697	
	(0.1172)	(0.0932)	(0.1115)	
is_taiwan × treated_class × post				0.1011
				(0.1020)
Fixed-effects				
uspc_mainclass	Yes	Yes	Yes	
year_id	Yes	Yes	Yes	
assignee_country-uspc_mainclass				Yes
assignee_country-year_id				Yes
uspc_mainclass-year_id				Yes
Fit statistics				
Observations	550	446	549	1,544
R^2	0.86069	0.83421	0.92513	0.97475
Within R^2	0.01961	5.18×10^{-5}	0.00250	0.00249

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

DV: $\ln(\text{Patents} + 1)$. Col. (4): DDD estimate. Israel and Korea had no R&D credit restructuring in 2010. SE clustered by class (1–3) or country×class (4).

Table 3 reports placebo tests using Israel and South Korea—countries with similar USPTO technology profiles but no domestic R&D credit restructuring in 2010. Israel's treated × post coefficient is 0.011 (SE = 0.093), a precise null. South Korea's is 0.070 (SE = 0.112), small and insignificant. The triple-difference estimate in column (4), which compares Taiwan's treated-control differential to the same differential in Israel and Korea, is 0.101 (SE = 0.102), similar in sign to the Taiwan-only DiD but statistically insignificant.

The placebo tests serve two functions. First, they confirm that the Taiwan DiD result is not an artifact of global trends in semiconductor or optoelectronics patenting. Second, they establish that the null finding is not unique to Taiwan—the absence of an effect is consistent across countries with different R&D credit regimes.

5.3 Robustness and Heterogeneity

Table 4: Robustness Checks and Heterogeneity

Dependent Variables:	ln_patents				sh
	Baseline	Placebo 2007	Semicond.	Optoelec.	Share DiD
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
treated_class × post	0.2043*				
	(0.1172)				
treated_class × placebo_post		0.0316			
		(0.0977)			
semi × post			0.2766***		
			(0.0892)		
opto × post				0.2685*	
				(0.1515)	
is_tw × post					-0.0036
					(0.0102)
Fixed-effects					
uspc_mainclass	Yes	Yes	Yes	Yes	
year_id	Yes	Yes	Yes	Yes	
assignee_country					Yes
yid					Yes
Fit statistics					
Observations	550	350	550	550	33
R^2	0.86069	0.91404	0.85870	0.86025	0.94708
Within R^2	0.01961	0.00098	0.00561	0.01656	0.00226

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Col. (2): placebo treatment at 2007 (pre-period only). Cols. (3)–(4): semiconductor (257, 438) and optoelectronics (345–385) sub-sectors. Col. (5): DV is share of patents in SUI classes; DiD vs. Israel/Korea. SE clustered by class (1–4) or country (5).

Table 4 reports robustness checks. Column (2) shows the 2007 placebo treatment test: the coefficient is 0.032 (SE = 0.098), a clean null. This confirms that pre-trends are flat—there is no differential movement in treated classes relative to controls before the IIA—strengthening the causal interpretation of the post-IIA estimates.

The most informative heterogeneity emerges from separating semiconductors and optoelectronics. Column (3) shows that semiconductor classes experienced a statistically significant increase of 0.277 log points (SE = 0.089, $p < 0.01$) in relative patenting after the IIA. Column (4) shows optoelectronics classes experienced a positive but imprecise effect (0.269, SE = 0.152), approaching significance. Both sectors show positive coefficients, consistent with the inframarginal subsidy hypothesis: firms in these sectors—particularly semiconductors led by TSMC—were innovating at the frontier regardless of the credit.

Column (5) uses the share of Taiwan's total annual patents in SUI classes as the outcome, with Israel and Korea as controls. The coefficient is -0.004 (SE = 0.010), essentially zero. This null share result, combined with the positive absolute count estimates, indicates that Taiwan's treated classes grew in absolute terms but non-treated classes grew proportionally as well—consistent with the IIA's uniform 15% credit stimulating broad-based innovation across all sectors.

Leave-one-out analysis (not tabulated) shows that the pooled result is stable: no single class among the 22 treated classes drives the result.

6. Discussion

The null finding admits three interpretations, ordered by economic interest.

First, the *inframarginal subsidy* interpretation: for firms at the global technology frontier, targeted R&D credits subsidize investment that would have occurred anyway. Taiwan's semiconductor industry was already the world leader in contract chip manufacturing; its R&D investment was driven by competitive pressure from Samsung, Intel, and GlobalFoundries, not by a 20-percentage-point tax advantage (Fuller, 2005). The significant positive effect for semiconductors strengthens this interpretation—frontier firms not only maintained but *increased* their patenting, possibly because the IIA's broader credit base stimulated complementary innovation in downstream sectors.

Second, the *substitution* interpretation: the IIA may have redirected R&D effort from lower-quality marginal patents (induced by the high SUI credit) toward higher-value activities. The near-zero effect on claims per patent (-0.024, SE = 0.357) and the positive but imprecise citation effect (+0.967, SE = 0.920) are consistent with no meaningful change in patent quality, suggesting that the type of innovation was unaffected by the credit restructuring.

Third, the *underpowered* interpretation: the true effect may be negative but smaller than my minimum detectable effect of 0.328 log points. However, the pooled estimate is positive and approaches significance, the semiconductor heterogeneity yields a significant *positive* effect, and the flat pre-trends strengthen the parallel trends assumption. Taken together, the evidence is inconsistent with a meaningful negative effect of removing targeted credits.

The implications for contemporary industrial policy are direct. If the United States or European Union is considering sector-targeted R&D subsidies for semiconductors, Taiwan's experience suggests that firms already at the frontier may not change their innovation behavior in response to targeted credits. The subsidy transfers rents to firms that would innovate regardless—what I call the "phantom credit," real on the tax return but invisible in the patent record. This does not mean R&D credits are useless: Bloom et al. (2002) and Dechezleprêtre et al. (2016) provide credible evidence that credit *levels* affect R&D. But the *targeting* of credits toward already-dominant sectors may not add value above a uniform credit regime.

A caveat applies to the semiconductor heterogeneity. The 2010–2013 period coincided with a global semiconductor demand expansion and Taiwan's deepening role in foundry manufacturing. TSMC's revenue grew approximately 40% between 2010 and 2013, driven by smartphone and mobile device demand. The positive semiconductor coefficient may therefore reflect sector-specific demand rather than (or in addition to) the irrelevance of targeted credits. The subsector heterogeneity is suggestive but not conclusive about the mechanism.

This paper cannot speak to effects on domestic R&D spending (as opposed to USPTO patenting), on non-patent innovation, or on entry and exit of firms in strategic sectors. These remain important margins for future work.

7. Conclusion

Taiwan removed its sector-targeted R&D tax credits and strategic sector patenting continued uninterrupted. The semiconductor industry—the world's most advanced contract chip manufacturer—continued innovating at the same or higher rate. Optoelectronics saw no detectable change. Innovation at the frontier, it appears, does not require a targeted subsidy to sustain itself. For policymakers, Taiwan's experience suggests that at the frontier, targeted credits may not be necessary to sustain innovation—though this evidence cannot speak to emerging sectors where subsidies may shift behavior at the margin.

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

Arrow, Kenneth, "Economic welfare and the allocation of resources for invention," *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 1962, pp. 609–626.

Bloom, Nick, Rachel Griffith, and John Van Reenen, "Do R&D tax credits work? Evidence from a panel of countries 1979–1997," *Journal of Public Economics*, 2002, *85* (1), 1–31.

Chen, Tain-Jy and Ying-Hua Ku, "Taiwan's science and technology policy," *Asian Survey*, 2006, *46* (6), 864–886.

Chen, Zhao, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu, "Evaluating the effectiveness of R&D tax credits: evidence from the alternative simplified credit," *NBER Working Paper*, 2021, (29233).

Dechezleprêtre, Antoine, Elias Einiö, Ralf Martin, Kieu-Trang Nguyen, and John Van Reenen, "Do tax incentives for research increase firm innovation? An RD design for R&D, patents and spillovers," *NBER Working Paper*, 2016, (22405).

Frakes, Michael D and Melissa F Wasserman, "Patent office cohorts," *Duke Law Journal*, 2017, *65*, 1601–1655.

Fuller, Douglas B, "Creating and contesting competitiveness: the politics and economics of technological innovation in East Asia," *Competition and Change*, 2005, *9* (3), 285–305.

Hall, Bronwyn H, "Tax incentives for innovation in the United States," *Innovation Policy and the Economy*, 2000, *1*, 1–38.

Hu, Albert G Z and Adam B Jaffe, "Returns to patenting and R&D at the firm level in Taiwan," *Review of Industrial Organization*, 2005, *26* (1), 105–126.

Industrial Development Bureau, Taiwan, "Strategic Industries Policy Review," Technical Report, Ministry of Economic Affairs 2008.

Juhász, Réka, Nathan Lane, and Dani Rodrik, "The new economics of industrial policy," *Annual Review of Economics*, 2023, *16.*

Krueger, Anne O, "Government failures in development," *Journal of Economic Perspectives*, 1990, *4* (3), 9–23.

Lane, Nathan, "The new empirics of industrial policy," *Journal of Industry, Competition and Trade*, 2020, *20*, 209–234.

Mathews, John A, *Dragon multinationals: a new model for global growth*, Oxford University Press, 2006.

Ministry of Economic Affairs, Taiwan, "Industrial Innovation Act: Implementation Guidelines," Technical Report, Industrial Development Bureau 2010.

OECD, "R&D Tax Incentive Indicators," Technical Report, Organisation for Economic Co-operation and Development 2021.

Pack, Howard and Kamal Saggi, "The case for industrial policy: a critical survey," *World Bank Research Observer*, 2006, *21* (2), 267–297.

Rao, Nirupama, "Do tax credits stimulate R&D spending? The effect of the R&D tax credit in its first decade," *Journal of Public Economics*, 2016, *140*, 1–12.

Rodrik, Dani, "Industrial policy for the twenty-first century," *CEPR Discussion Paper*, 2004, (4767).

Romer, Paul M, "Endogenous technological change," *Journal of Political Economy*, 1990, *98* (5, Part 2), S71–S102.

Sampat, Bhaven and Heidi L Williams, "Examiner characteristics and patent office outcomes," *Review of Economics and Statistics*, 2019, *101* (4), 664–677.

TSMC, "Annual Report 2010," Technical Report, Taiwan Semiconductor Manufacturing Company 2010.

UNCTAD, "World Investment Report: Transnational Corporations and the Internationalization of R&D," Technical Report, United Nations Conference on Trade and Development 2005.

Wilson, Daniel J, "Tax incentives for R&D: new evidence and new designs," *Federal Reserve Bank of San Francisco Economic Letter*, 2009.

WIPO, "World Intellectual Property Indicators," Technical Report, World Intellectual Property Organization 2015.

A. Data Appendix

PatentsView Data. The patent data are drawn from PatentsView, the USPTO's official research data platform. I download six bulk data files: patent metadata, filing dates, disambiguated assignee information, assignee locations with country codes, USPC technology class assignments at issue, and CPC classification as backup. Forward citations are computed from the US patent citation file.

An important data limitation is that PatentsView contains only granted patents, not all applications. The original identification strategy envisioned application-level data, which would capture innovation effort regardless of grant outcomes. Since grants reflect both filing decisions and USPTO examination processes, my estimates capture the net effect on successful patenting. Grant lags average 2–3 years, which means patents granted in 2010–2013 may reflect applications filed before the IIA. I address this by using filing dates (not grant dates) for timing, but acknowledge that application data would provide a cleaner measure of innovation effort.

Country Assignment. Patents are assigned to countries based on the disambiguated assignee location. For patents with multiple assignees in different countries, each country receives a separate observation. This affects fewer than 0.1% of observations in the analysis sample.

Technology Class Assignment. Each patent is classified using its primary USPC mainclass at issue. The 22 treated classes are selected based on their correspondence to six SUI designated sectors: semiconductors, optoelectronics, communications, precision instruments, IT hardware, and related strategic fields. The mapping from SUI sector definitions to USPC classes follows Hu and Jaffe (2005) and the Industrial Development Bureau's sector classification documents.

Panel Construction. The analysis panel aggregates patent-level data to the technology class $\times$ country $\times$ year level. Years are defined by filing date (not grant date) to capture innovation timing accurately. The top 50 USPC classes are selected by Taiwan patent volume to ensure adequate cell sizes; these classes account for over 85% of all Taiwan USPTO filings.

Treatment Classification. Table 5 lists the 22 treated USPC classes and their sector assignments under the SUI.

Table 5: Treatment Classification: USPC Classes by SUI Sector

Sector	USPC Classes	Description
Semiconductors	257, 438	Solid-state devices, mfg.
Optoelectronics	345, 348, 349, 359, 362, 385	Display, TV, LCD, optics, lighting
	315	Lamp/discharge devices
Communications	343, 370, 375, 455	Antennas, multiplex, pulse, telecom
Precision instr.	324, 356, 382	Elec. measuring, optical testing, image
IT hardware	365, 710, 711, 713, 714, 716	Memory, I/O, support, error det., CAD

Notes: Classification based on SUI designated strategic industries (semiconductors, optoelectronics, communications, precision instruments, IT). The remaining 28 of the top 50 USPC classes by Taiwan patent volume serve as controls.

B. Identification Appendix

Event Study Results. The event study specification reveals flat pre-trends: all pre-period coefficients are small in magnitude and statistically insignificant, supporting the parallel trends assumption. Post-IIA coefficients are positive but imprecise, consistent with the main DiD estimates.

2007 Placebo Test. Applying a placebo treatment at 2007 (using only 2003–2009 data) yields a coefficient of 0.032 (SE = 0.098), a clean null. This confirms that pre-trends are flat and there is no differential movement in treated classes relative to controls before the IIA.

C. Robustness Appendix

Leave-One-Out. Dropping each treated class one at a time yields coefficients that remain stable around the pooled estimate. No single class drives the result.

Minimum Detectable Effect. With standard errors of 0.117, the design achieves 80% power at the 5% significance level for effects of 0.328 log points (approximately 39% in levels), corresponding to a standardized effect size of 0.388. This means I can confidently rule out moderate-to-large negative effects on patenting in strategic classes.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
ln(Patents)	0.2043	0.1172	0.8456	0.2415	0.1386	Large positive
Claims/patent	-0.0242	0.3567	2.2436	-0.0108	0.1590	Small negative
Forward citations	0.9665	0.9198	4.7450	0.2037	0.1938	Large positive
Panel B: Heterogeneous (sample splits)						
ln(Pat.), Semicond.	0.2766	0.0892	0.8456	0.3271	0.1055	Large positive
ln(Pat.), Optoelec.	0.2685	0.1515	0.8456	0.3175	0.1792	Large positive

Notes: **Country:** Taiwan. **Research question:** Did the 2010 transition from sector-targeted R&D tax credits (up to 35% for strategic industries under the SUI) to a uniform 15% credit (IIA) reduce patenting effort in previously favored technology classes at the USPTO? **Policy mechanism:** The IIA replaced the SUI's sector-specific enhanced credits with a uniform 15% credit for all industries, imposing an effective 20 percentage-point credit reduction on designated strategic sectors including semiconductors, optoelectronics, communications, and IT hardware. **Outcome definition:** Annual count of utility patent applications filed at the USPTO by Taiwan-based assignees, aggregated at the USPC technology class level; claims per patent measures patent scope; forward citations count subsequent US patents citing each patent. **Treatment:** Binary; 22 USPC classes corresponding to SUI-designated strategic industries are treated after January 2010. **Data:** PatentsView bulk download, 2003–2013, USPC class × year level, approximately 190,000 utility patents across 50 technology classes and 11 years for three countries (Taiwan, Israel, South Korea). **Method:** Two-way fixed effects difference-in-differences with USPC class and year fixed effects; standard errors clustered at the USPC class level; heterogeneity via sample splits between semiconductor and optoelectronics sub-sectors. **Sample:** Top 50 USPC classes by Taiwan patent volume; utility patents only. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. 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).