Engineering Germany's Industrial Success

Professional Societies and the Social Dynamics of Innovation

Stefan Dietrich

September 10, 2026

Few economies have ever reached the global technological frontier; Germany did, moving from industrial laggard in 1850 to world leader in chemistry, electrical machinery, and precision instruments by 1900. I study the social organization of knowledge behind that catch-up: the German Engineering Society, a voluntary professional body whose local chapters brought working engineers together through lectures, factory visits, and a technical journal. Exploiting the staggered founding of 22 chapters in a difference-in-differences design on high-value patents, I find that chapter founding raises patenting by four percent a year — cumulating over fifteen years to roughly 60 percent of the average unit's patent stock — with flat pre-trends and no effect on patent classes outside the Society's scope. Chapters amplified capacity already in place: patenting rose most where members concentrated, in technologies of prior regional specialization, and among owner-engineers rather than credentialed graduates. Localized knowledge spillovers, organized through a professional society, emerge as a distinct input to innovation, complementary to formal education and patent rights.

Boston University. Email: sjdietri@bu.edu. I thank Martin Fiszbein, James Feigenbaum, Natalia Ramondo, and all seminar participants at Boston University for helpful comments. All errors are my own.

1. Introduction

Innovation is local and social. Inventors cluster, knowledge travels more easily within a region than across regions, and the institutions that bring practitioners into contact shape what gets invented where. The standard account of the most dramatic catch-up of the industrial age makes little use of this. Between 1850 and 1900, German industry went from laggard to leader, moving from imitation to invention and defining the state of the art in fields from synthetic dyes to dynamos and fine optics, and that rise is usually credited to formal inputs: schools and universities that trained engineers, and a patent law that protected their inventions. This paper is about the layer that account leaves out — the social organization that connected them. Trained engineers and enforceable patents do not by themselves produce commercially valuable inventions; the knowledge of the laboratory and the factory floor has to be brought into contact. In late-19th-century Germany, much of that coordination ran through a private professional body, the German Engineering Society: a voluntary association in the language of the social-capital literature, and a vehicle for localized knowledge spillovers in the language of the innovation literature.

The Society's local chapters organized monthly technical lectures, factory visits, and machinery demonstrations; its weekly journal reached every member across the Empire; and chapter leadership lay with factory owners and senior practitioners rather than professors. Whether these chapters *caused* local innovation, or merely appeared where innovation was already accelerating, is the question this paper takes up. The setting is well suited to answering it: national unification (1871) and a unified patent law (1877) fixed patent rules, education policy, and the aggregate supply of engineers, so the remaining variation is local. I exploit the staggered founding of 22 chapters between 1877 and 1918, looking at the production of high-value patenting — patents renewed beyond their tenth year in the Streb, Baten, and Yin (2006) corpus. Because chapter effects build gradually, two-way fixed effects attenuates; the doubly-robust Callaway and Sant'Anna (2021) estimator resolves that.

The main result is that chapter founding raises annual high-value patenting by roughly 4.2 percent on average over the post-treatment window, against flat pre-trends across the fifteen pre-founding years. The effect is not a one-off jump but a build-up: from roughly zero in the founding year to nine percent fifteen years later, and compounding the event-study path over that window cumulates to roughly 60 percent of the average unit's patent stock, a substantial contribution.

Three further sets of results sharpen the interpretation. First, the effect is the Society's, not background industrialization's: splitting the corpus by the Society's technical scope, patenting accelerates on the 49 in-scope classes — chemistry, machinery, electrical engineering, transport, metallurgy, mechanized textiles — and not on the 40 out-of-scope classes such as windmills, watermills, and traditional crafts, a boundary fixed ex ante from the class descriptions. Second, the mechanism is amplification of capacity already in place: units with denser 1900 membership patent more after founding (elasticity 0.68 within chapter), post-founding sectoral composition tracks pre-founding specialty, and the patenters are owner-practitioners rather than Dipl.-Ing. graduates (Gispen, 1989). Third, the effect persists: counties near pre-1918 chapter headquarters carry a manufacturing-GDP premium of 24–27 log points through 1935–1966, fading on aggregate GDP per worker by the present — suggestive rather than causal, but aligned with the patent-level mechanism.

The estimate is robust. Across fourteen specifications — alternative heterogeneity-robust estimators (Sun and Abraham, 2021; Borusyak, Jaravel, and Spiess, 2024; Arkhangelsky et al., 2021; Chaisemartin and D'Haultfœuille, 2024), sample restrictions, PPML on raw counts, leave-one-chapter-out estimation, and permutation inference — the doubly-robust estimates cluster between 4.2 and 4.8 percent. Two donut-hole designs address host-city confounding and spatial spillovers, and the same spatial donut applied to engineering-school foundings returns flat noise: the chapter effect is not a downstream consequence of the educational shock that produced the chapter's members.

The paper contributes to three strands of literature. The first is on engineers and growth. Maloney and Valencia Caicedo (2022) show that the density of engineers — not of lawyers, and not of schooling generally — predicts subsequent income across countries and across US counties, using the 1862 Morrill land-grant colleges for identification; Hanlon (2025) documents, from British patent records, how invention shifted over the 19th century from artisan tinkerers to credentialed professional engineers. Both treat the supply of engineers as the operative margin. But a given supply can be organized in very different ways — in consulting partnerships, in state bureaus, in professional societies — and neither design can separate the engineers from the organizational form in which they worked. This paper's setting holds patent law, education policy, and the national supply of engineers fixed, and measures the organizational layer directly: the arrival of a local chapter, holding constant the engineers already in place.

The second strand studies formal inputs to innovation capacity — the input that endogenous growth theory takes as given (Romer, 1990; Aghion and Howitt, 1992). Squicciarini and Voigtländer (2015) show that upper-tail human capital, proxied by *Encyclopédie* subscriptions, predicts French city growth once industrialization raised the return to advanced knowledge; Moser (2005) shows that patent institutions shape the direction more than the volume of innovation. For 19th-century Germany specifically, Cinnirella and Streb (2017) establish the role of human capital, Dittmar and Meisenzahl (2022) the effect of research universities, and Donges, Meier, and Silva (2023) the effect of Napoleonic-era institutional reforms, the latter with a within-country design close in spirit to the one used here. All of these inputs are formal and largely state-provided: schools, universities, courts, patent offices. The chapter network was neither, but a voluntary, member-financed body. My estimates thus identify an additional, quantitatively significant channel for growth.

The third strand is on the local and social organization of knowledge. Knowledge spillovers are strongly localized (Jaffe, Trajtenberg, and Henderson, 1993; Audretsch and Feldman, 1996); inventors cluster, and clusters raise their productivity (Carlino and Kerr, 2015); and associational social capital — the density of voluntary organizations — predicts regional economic performance and innovation (Putnam, Leonardi, and Nanetti, 1993; Akçomak and Weel, 2009; Satyanath, Voigtländer, and Voth, 2017). Economic historians find the same forces at work before the modern data begin: Allen (1983) on collective invention among the ironmasters of England's Cleveland district, Kelly, Mokyr, and Ó Gráda (2023) on the tacit mechanical skill carried by British artisans, and Mokyr (2002) on the academies, societies, and journals through which academia and fabricants communicated. The closest antecedent is Cinnirella, Hornung, and Koschnick (2025), who show that the economic societies of 18th- and early-19th-century Germany predict patenting in their host regions decades later. I take the question to the societies of the engineers, the organization at the center of the second industrial revolution, when science-based industry raised the return to moving knowledge between academy and factory floor. The staggered rollout delivers what a cross section lacks: their evidence establishes that knowledge-network organization matters; mine identifies the causal effect and the channel — amplification of existing local capacity — through which it operated.

Outline. Section 2 presents the setting, the data, and the motivating facts; Section 3 the empirical strategy; Section 4 the main results and robustness; Section 5 the mechanism; Section 6 the long run; Section 7 concludes.

2. Historical Setting and Data

2.1. The catch-up puzzle and the engineering society

By 1850 Germany was a politically fragmented agrarian economy decades behind Britain in industrial output. By 1900 it had overtaken Britain in high-value patenting per capita and led global markets in chemistry, electrical machinery, and precision instruments. The institutional framework conditioning this transition came together in three steps: polytechnic schools were established from 1825 onward (eleven institutions by 1880, upgraded to technical universities by 1899); state engineering schools followed from the 1880s, with 29 institutions by 1900 providing practical training to most working engineers (Gispen, 1989); and political unification (1871) together with a unified patent law (1877, modeled on the British system) created the national innovation system. What this framework did not provide — and what the German Engineering Society filled — was a coordination layer between the academy and the factory floor.

The German Engineering Society (Verein Deutscher Ingenieure, VDI) was founded in May 1856 in central Germany by twenty-three young engineers. It operated through a weekly technical journal distributed to every member across the Empire, and through a network of local chapters that organized monthly technical lectures, factory visits, and machinery demonstrations. Chapter leadership was dominated not by university professors but by factory owners and senior practitioners (Gispen, 1989); the society was, in practice, a coordination device for industrial as much as for academic knowledge. Membership grew from 172 in 1856 to more than 27,000 in 1925; chapter coverage expanded from 5 to 50. Figure 1 maps the chapter-headquarters network on 1925 administrative boundaries. The VDI also influenced national policy through the 1877 patent law and through its national standards committee (Gispen, 1989), but the variation I exploit is the local rollout, not the national footprint.

The engineers the Society organized came out of a two-track training system, and the division between the tracks structured the profession. The technical universities taught engineering as applied science and, from 1899, granted the Dipl.-Ing. degree with formal parity to the classical universities; their graduates gravitated toward academic, government, and senior corporate positions. The state engineering schools — 29 institutions by 1900 — supplied the far larger stream of working engineers trained in drawing offices and workshops. Gispen (1989) describes the resulting split as one between a *school culture* of credentialed graduates and a *shop culture* of owner-engineers and senior practitioners who had learned the trade on the factory floor. Unlike law or medicine, the profession had no licensure and no protected title; an engineer's standing rested on what he could demonstrate to other engineers, not on a state credential.

The chapter was the Society's instrument for making that demonstration routine. A local chapter maintained a meeting calendar of technical lectures and discussion evenings, organized excursions to member factories and machinery demonstrations, kept a technical library, and fed local material to the national journal. For a machine-builder in a provincial town, the chapter was the point of access to the Empire-wide stock of practical knowledge — and, in the other direction, the channel through which his own workshop solutions reached the national profession. It is this coordination function, operating between the academy and the factory floor rather than inside either, that the staggered chapter rollout allows me to measure.

Figure 1: Local chapters of the German Engineering Society

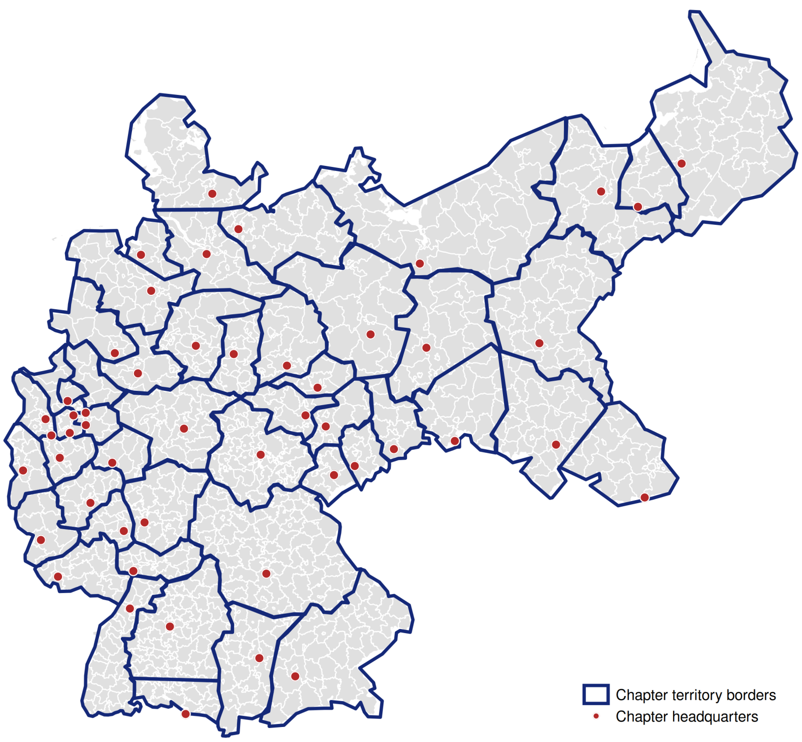

Notes. Chapter headquarters with 1925 territorial boundaries overlaid on 1900 statistical-unit borders.

2.2. Data

The patent outcome is drawn from the Streb, Baten, and Yin (2006) corpus of patents filed at the Imperial Patent Office between 1877 and 1918 and renewed beyond ten years ($N = 39{,}360$ of roughly 380,000 filings). The renewal hurdle is meaningful because maintenance fees were substantial: cumulative fees through year ten amounted to 2,300 Marks, roughly two-and-a-half years of industrial wages (Bry, 1960). Patents surviving the year-ten renewal therefore signal commercial value rather than mere filing activity. The corpus is organized into 89 technology classes. The classification records the application domain of the invention rather than the inventor's professional discipline: a chemical process applied to textile dyeing, for instance, is recorded under textiles, not chemistry. I aggregate the 89 classes into two principal buckets, defined ex ante by whether a class falls within the Society's documented professional scope. The first comprises the 49 classes covering second-industrial-revolution technologies (chemistry, machinery, electrical engineering, transport, metallurgy, mechanized textiles) — the fields in which the Society and its journal were active — and accounts for roughly three-quarters of the corpus. The second comprises 40 classes spanning windmills, watermills, mechanical clocks, traditional crafts, and small consumer goods, technologies outside that scope; this bucket serves as the placebo outcome throughout. The classification is made from the class descriptions alone and fixed before any outcome is estimated. The chemists split from the VDI in 1887 to form a separate society and the electrical engineers split in 1893, biasing my estimates on those sectors downward through the post-split panel years.

I use the 1900 and 1925 VDI membership directories (15,513 and 27,176 hand-transcribed individuals; 26,869 unique individuals after within-name-and-city deduplication across the two directories), each listed with name, professional title, employer, and city of residence. For the engineer-level analysis in Section 5 I link the directory to the Streb–Baten corpus by name using the `reclink2` probabilistic record-linkage routine (Wasi and Flaaen, 2015), producing 644 person-level links. The analysis uses 1900 statistical units as the primary geographic resolution. The German Empire contained 1,020 such units; 493 enter the canonical difference-in-differences sample once always-treated units (covered by chapters founded by 1877) are dropped and units outside any 1925 chapter territory are excluded; the units of the three chapters founded after 1918 remain in the sample as never-treated controls. Chapter territories are taken from the 1925 yearbook and assigned to statistical units via majority-area shapefile overlay; patents are geocoded via the inventor's residence address. The eastern territories (East and West Prussia, Posen, parts of Silesia) were predominantly rural and structurally less suited to a technical-professional treatment; the main estimate strengthens marginally when they are excluded (Appendix C).

Historical population comes from Rösel (2023); the 1871 cross-section forms the basis for pre-treatment density and an interaction with year provides a time-varying trend control. The railway network is mapped at the unit-year level from Cederman, Pengl, and Girardin (2025). Long-run outcome data come from Peters (2022) for the 1933–1966 period at the modern-county level (manufacturing GDP, manufacturing employment, GDP per worker) and from the INKAR regional indicators (BBSR, 2019) for 2019. Coverage is West Germany throughout the long-run exercises, since post-1945 division and central planning in the East leave no comparable output series.

2.3. Spatial co-distribution and summary statistics

A first look at the data establishes the cross-sectional fact that motivates the causal design. Figure 2 maps engineers and high-value patents at the statistical-unit level. The spatial co-distribution is tight: regions with more engineers also produced more patents. This is the regularity that any causal claim must navigate around, since chapters were founded where engineers already clustered.

Figure 2: Spatial co-distribution of engineers and high-value patenting

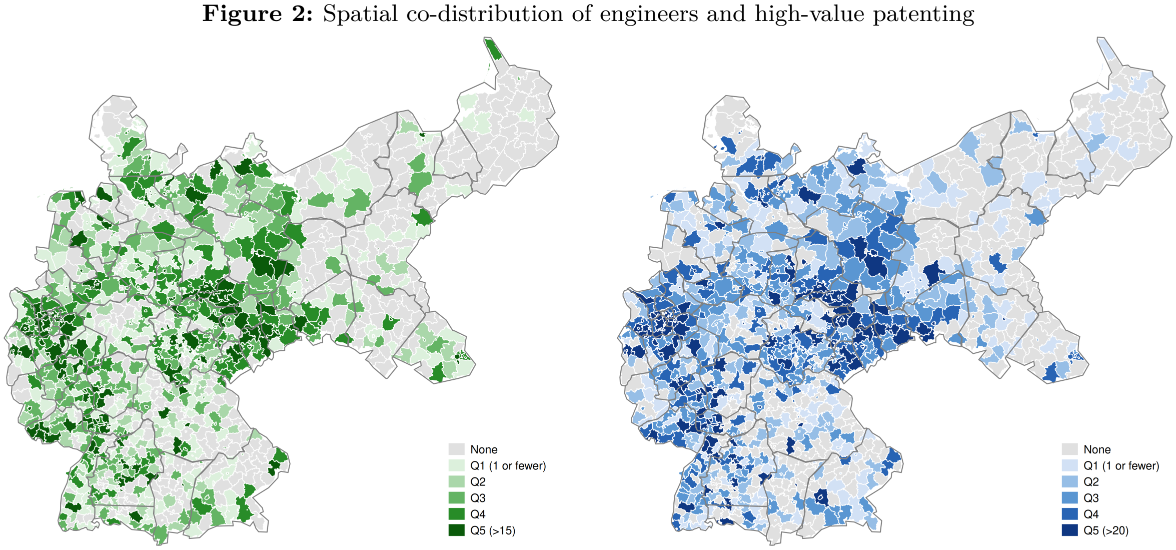

(a) VDI members, 1900 **(b)** High-value patents, 1877–1918

Notes. Quintile choropleths at the statistical-unit level on the 1900 boundary file.

Table 1 reports summary statistics: Panel A covers all 1,020 statistical units of Imperial Germany, Panel B the chapter network. The canonical difference-in-differences sample of Section 3 comprises 493 of these units in an annual panel from 1877 to 1918 (N = 20,706 unit-year observations); within this sample the patent outcome averages 0.30 high-value patents per unit-year, with a right skew that motivates the $\ln(1+x)$ functional form used throughout; PPML on raw counts is reported in Appendix C.

Table 1: Summary statistics

	Mean	Median	SD	Min	Max
Panel A: All 1900 Statistical Units (full Kaiserreich, N = 1020)					
Total high-value patents (1877–1918)	24.9	2	196.1	0	5,668
Engineering	8.5	0	91.1	0	2,802
Chemical	5.9	0	49.2	0	842
Mechanical engineering	3.6	0	22.1	0	581
Patents by Individuals	12.4	1	86.7	0	2,507
Patents by Firms	12.4	0	113.3	0	3,161
Engineering Society members (1900)	13.2	1	61.5	0	1,478
Population, 1871 (000s)	28.7	24	44.6	0	1,029
Railroad km (1885)	32.4	28	27.5	0	226
Panel B: Local Chapters — All Bezirksvereine (N = 50)					
Founding year	1882	1878	19.6	1856	1924
Members (1900)	297	259	266	61	1,612

Notes. Panel A: all 1,020 statistical units of Imperial Germany (1900 boundaries). Panel B: the 50 local chapters.

2.4. Two stylized facts

Two cross-sectional facts motivate the causal design. Both establish that there is something to identify; neither delivers identification.

The first is the within-country counterpart to the cross-country regularity that Maloney and Valencia Caicedo (2022) document. Table 2 reports the cross-sectional OLS of $\ln(1+\text{patents 1877–1918})$ on $\ln(1+\text{engineers 1900})$ at two administrative resolutions: modern municipality (N = 10,736) and modern county (N = 315). The coefficient lies between 0.49 and

0.71 across columns, with state fixed effects, 1871 population, and pre-treatment geographic controls. A ten-percent increase in 1900 engineer density predicts a five- to seven-percent increase in high-value patenting over the next four decades. The association is stable across resolutions, robust to dropping the top one percent of units by patent count, and quantitatively comparable to the engineer-share-and-growth relationship that Maloney and Valencia Caicedo (2022) document on cross-country data.

Table 2: Engineer density and high-value patenting, 1877–1918: cross-sectional OLS

	Municipality (N = 10,736)		District (N = 315)	
	(1) All	(2) Engineering	(3) All	(4) Engineering
ln(1+Engineers 1900)	0.709***	0.488***	0.636***	0.588***
	(0.023)	(0.023)	(0.047)	(0.044)
ln(population 1871)	0.049***	-0.000	0.496***	0.415***
	(0.008)	(0.005)	(0.136)	(0.118)
Controls:				
Geography	✓	✓	✓	✓
Soil & climate	✓	✓	✓	✓
State FE	✓	✓	✓	✓
R^2	0.606	0.582	0.657	0.626

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The second fact is that the cross-sectional pattern persists into long-run economic outcomes. Regressing West-German manufacturing GDP (Peters, 2022) on 1900 engineer density at the modern-county level yields coefficients near 0.30 in 1935 and 0.21 in 1966 — a persistence of three generations through both world wars. The analogous coefficient on aggregate GDP per worker is 0.14 in 1935 and fades to roughly 0.03 by 2019. Manufacturing responds several times more strongly than aggregate output, and the manufacturing effect survives both wars. Appendix Table 10 reports the underlying regressions. Section 6 returns to this pattern with a design-based counterpart.

These two facts establish a robust regularity but do not identify a causal effect. Chapters were founded in places that were already economically important, with stronger industrial fabric and denser population. The cross-sectional coefficient could reflect a causal effect of organization on patenting, a selection effect from chapter placement, or both. The remainder of the paper uses within-region variation in the timing of chapter founding to separate them.

3. Empirical Strategy

3.1. Staggered chapter rollout

Twenty-two chapters were founded between 1877 and 1918. The earliest switchers in the identification window were Schleswig-Holstein and East Prussia (both 1880); the latest was Mosel (1912). Figure 3 plots the resulting identification map. Each polygon is a chapter territory, labeled by founding year. Gray territories belong to the chapters that contribute no within-panel timing variation: the 25 chapters already founded when the patent panel opens in 1877, treated throughout the window and dropped from the treated set, and the three chapters founded after 1918 — Elbing and Lübeck (1921) and Osnabrück (1924) — whose statistical units serve as never-treated controls. Orange territories are the 22 switching chapters founded between 1877 and 1918 whose staggered timing drives identification. The orange-shaded territories span the full geographic extent of the Empire. Identification in what follows comes from comparing patenting in statistical units within an orange territory about to be covered by a newly-founded chapter against patenting in units not yet covered. Appendix C charts the switching dates (Figure 7).

3.2. Specification and identification assumptions

The baseline specification is

$$\ln(1 + \text{Patents})_{it} = \beta\, D_{it} + \boldsymbol{\gamma}'\mathbf{X}_{it} + \alpha_i + \delta_t + \varepsilon_{it}, \tag{1}$$

where i indexes statistical units, t indexes years from 1877 to 1918, and $D_{it} = 1$ once the chapter covering unit i has been founded. The control vector $\mathbf{X}_{it}$ contains time-varying railway kilometers and a pre-treatment population trend (1871 population × year); α_i and δ_t are statistical-unit and calendar-year fixed effects. The panel comprises $N = 20{,}706$ unit-years across 493 statistical units. Identification requires no anticipation, parallel trends conditional on $\mathbf{X}$, and no spillover across territorial boundaries. The first two assumptions are testable

Figure 3: The identification map

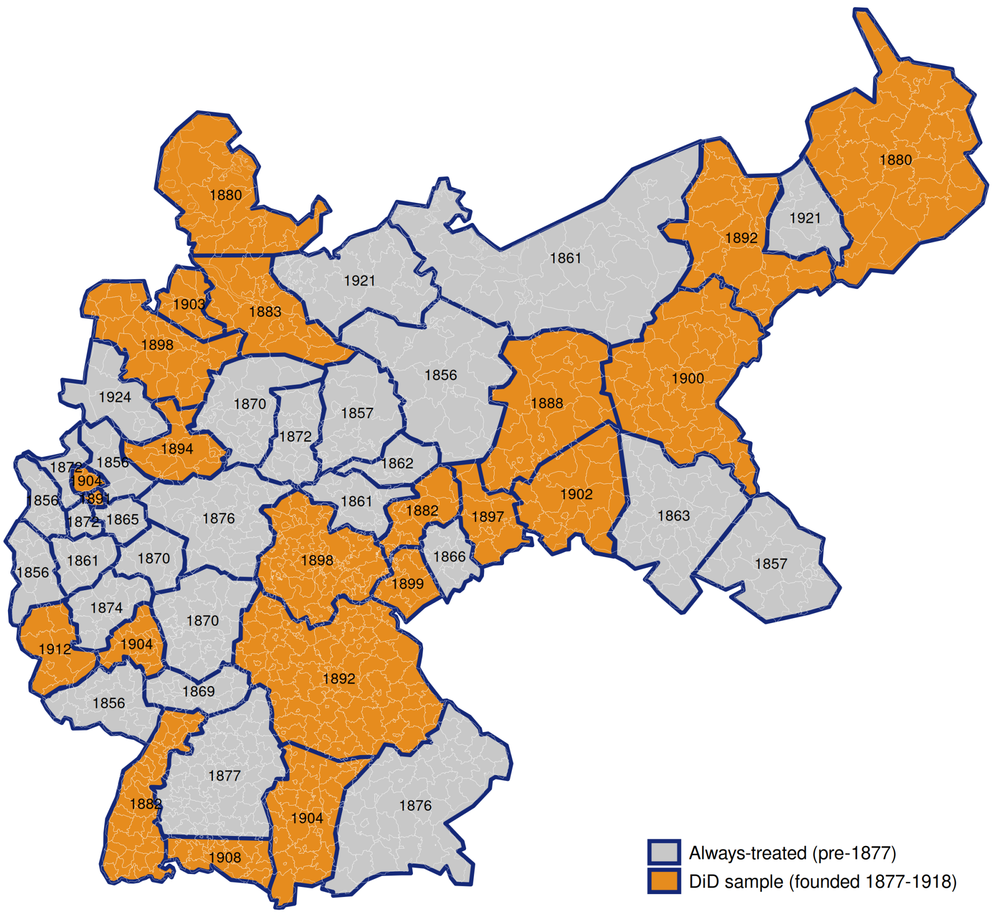

Notes. Chapter territories on the 1900 statistical-unit boundary file, labeled by chapter founding year. Gray: non-switching territories — the 25 chapters founded by the 1877 start of the patent panel (always-treated, dropped from the treated set) and the three chapters founded after 1918 (never-treated). Orange: the 22 switching chapters founded over 1877–1918 whose staggered timing identifies the chapter effect.

through the event-study aggregation; the third is examined through a donut-hole design in Section 4.

3.3. Two-way fixed effects under staggered adoption

Standard two-way fixed effects is biased under staggered adoption with heterogeneous and dynamic treatment effects. The Goodman-Bacon (2021) decomposition shows that TWFE is a weighted average of all possible 2×2 comparisons between treated and control unit-years, including "forbidden" comparisons in which already-treated cohorts serve as controls; when treatment effects build up post-founding the forbidden comparisons enter with negative weight and attenuate the TWFE coefficient. On the canonical sample the TWFE coefficient on D_{it} is 0.014 — roughly a factor of three below the CS-DiD estimate of 0.042 reported below. The Goodman-Bacon decomposition (Appendix B) shows that 99.8 percent of TWFE weight on this sample comes from legitimate timing-group comparisons, since I drop always-treated chapters by construction, so the forbidden-comparison block is empirically negligible. The TWFE–CS-DiD gap on this sample is therefore not a forbidden-comparison artifact; it is the absence of doubly-robust inverse-propensity reweighting in TWFE, which pools heterogeneous treatment dynamics across cohorts without rebalancing on covariates.

3.4. Callaway–Sant'Anna and a corroborating estimator

I use the Callaway and Sant'Anna (2021) estimator as the workhorse. The estimator builds cohort-time-specific average treatment effects

$$\mathrm{ATT}(g,t) = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \,|\, G_i = g,\, t \geq g]\,, \tag{2}$$

where g indexes a chapter's founding year and t indexes calendar time, then aggregates across (g,t) pairs after estimation. I implement the doubly-robust inverse-propensity-weighted variant using not-yet-treated units as controls. The estimator is consistent if either the propensity score for being in cohort g conditional on $\mathbf{X}$ is correctly specified or the outcome regression on never-treated controls is correctly specified. Three aggregations are reported: the simple ATT, the cohort-size-weighted post-treatment average, and the event-study aggregation. Standard errors come from the multiplier bootstrap with 999 replications.

The CS-DiD estimator rests on conditional parallel trends, which is testable in the pre-period but untestable in the post-period. I bring in the Chaisemartin and D'Haultfœuille (2024) switchers-vs-stayers estimator (dCDH) as a corroborating estimator on a different identifying assumption. dCDH compares treatment-status switchers to stayers at each panel year and does not apply inverse-propensity reweighting. On the canonical sample the dCDH simple ATT is 0.048 (SE 0.005) with a flat pre-period (coefficient -0.0003) and a joint placebo test that does not reject.

4. Results

4.1. Main results

Figure 4 and Table 3 report the central result. Chapter founding raises annual high-value patenting in the canonical 22-switcher sample by roughly 4.2 percent on average over the post-treatment window (simple ATT = 0.042, multiplier-bootstrap SE 0.014, $p = 0.003$). The pre-trend block is flat: none of the 15 pre-period coefficients is significant at the five-percent level, the average pre-period effect is essentially zero, and the joint test of parallel trends does not reject. The post-treatment block rises monotonically from roughly zero at $t = 0$ (the founding year), to three percent at $t = +5$, to six to eight percent at $t = +10$, to nine to ten percent at $t = +15$. The gradual build-up matches a simple recruitment-and-meeting interpretation: chapters did not switch on at full strength; they grew their membership and meeting calendar and gradually became embedded in the local industrial fabric.[1]

The decomposition by the Society's technical scope delivers the within-design placebo. On the 49 in-scope classes — chemistry, machinery, electrical engineering, transport, metallurgy, mechanized textiles — the simple ATT is 0.037 (SE 0.011), close to the main estimate. On the 40 out-of-scope classes — windmills, watermills, mechanical clocks, traditional crafts — the simple ATT is 0.016 (SE 0.010), and I cannot reject the null at conventional levels. Where the Society's fields of activity lay, patenting accelerated; where they did not, it did not. I define the boundary from the class descriptions alone and fix it before observing any of the patent outcomes; the contrast is therefore a prediction the data could have rejected and did not.[2]

A subsample restriction to peacetime years (1877–1913) yields a simple ATT identical to the full window; the result is not driven by wartime patent-office disruption.

How large is this effect? The main ATT of roughly four percent per year is large but not implausibly so. Over a 15-year post-treatment window, compounding the event-study co-

[1] The Streb–Baten corpus dates patents by grant year, not filing year. Under the 1877 Imperial Patent Act, substantive pre-grant examination and the mandatory eight-week public opposition period typically introduced a lag of one to two years between filing and grant: a patent applied for in year g would not appear in the data until $g + 1$ or $g + 2$. Part of the observed $t = +1$ to $t = +2$ ramp therefore reflects this mechanical dating lag rather than treatment-effect dynamics; the recruitment-and-meeting interpretation remains the leading interpretation of the slope at longer horizons.

[2] The pre-industrial bucket coefficient is positive but statistically indistinguishable from zero ($t \approx 1.5$); a residual general-equilibrium effect of the post-1880 industrial boom on all patenting cannot be ruled out, but the magnitude is well below the in-scope bucket and I fixed the boundary before estimation.

Figure 4: Event-study estimates: chapter founding and high-value patenting

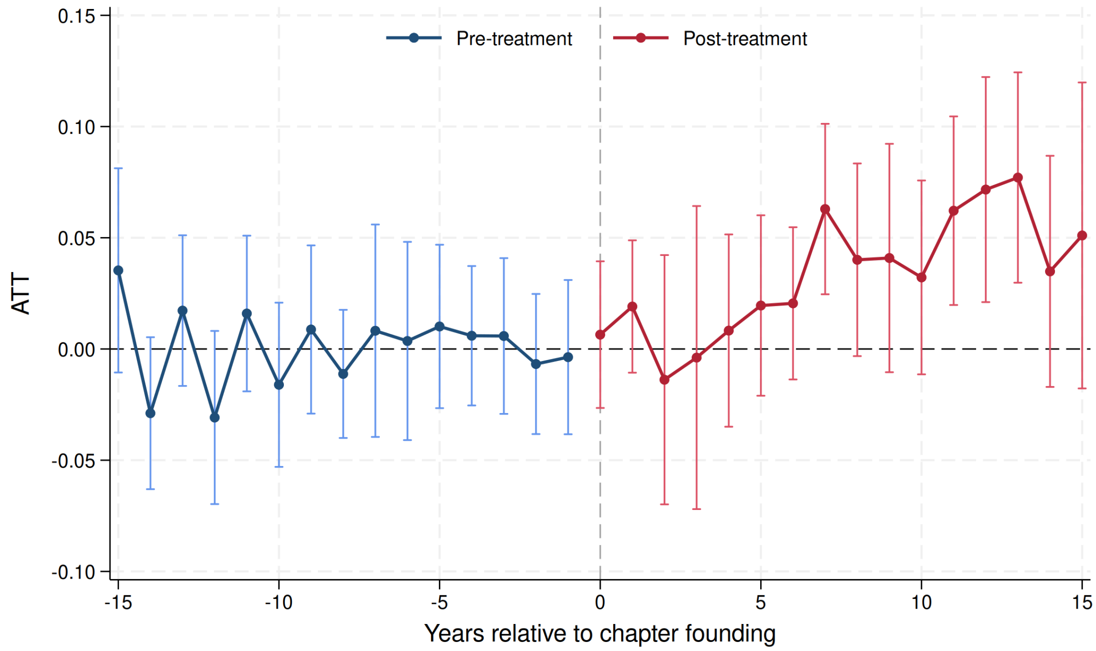

Notes. Callaway–Sant'Anna event study on $\ln(1 + \text{patents})_{it}$, ± 15 event-time window, canonical 22-switcher sample. Vertical bars: 95% confidence intervals from the multiplier bootstrap with 999 replications.

Table 3: Chapter founding and high-value patenting: main estimates

	Total (all 89 cl.)	**2nd Ind. Rev.** (49 cl.)	**Pre-Industrial** (40 cl., placebo)
ATT (CS-DiD, ± 15)	0.0419^{***}	0.0366^{***}	0.0157
	(0.0141)	(0.0114)	(0.0104)
Implied effect	**+4.2%/yr**	**+3.7%/yr**	+1.6%
Sample: 22 switching chapters, 493 stat units, 1877–1918 (20,706 obs)			
Method: Callaway–Sant'Anna DiD (`dripw`); controls: rail + pop×year			

Notes. Callaway–Sant'Anna ATT for total high-value patents and the decomposition into second-industrial-revolution technologies (49 classes) and pre-industrial / artisanal technologies (40 classes; placebo). Multiplier-bootstrap standard errors in parentheses. Sample: 22 switching chapters, 1877–1918, $N = 20{,}706$ unit-years across 493 statistical units. Method: Callaway–Sant'Anna DiD; controls: rail + pop × year.

efficients yields a cumulative effect of roughly 60 percent of the sample-mean patent stock per statistical unit. The cohort with the longest observed post-period — chapters founded in the early 1880s — shows event-study coefficients near 10 percent by the end of the window, consistent with the linear extrapolation of the average build-up. For context, Donges, Meier, and Silva (2023) find effects of comparable order on related innovation outcomes from Napoleonic-era legal reforms, and Cinnirella, Hornung, and Koschnick (2025) document large cross-sectional differences in patenting tied to economic-society membership; the magnitudes I estimate sit within the range that the literature on knowledge-organization institutions in industrial-era Europe has documented.

4.2. Robustness

Table 4 reports fourteen specifications grouped into three categories: alternative heterogeneity-robust estimators, sample restrictions, and specification-and-inference adjustments. The CS-DiD–family point estimates lie between 0.020 and 0.070 across the table; PPML on raw counts implies roughly 15 percent on counts. The permutation-test p-value of 0.12 (two-sided) and the synthetic-DiD insignificance are the two rows that fall outside conventional significance thresholds; I discuss both below.

Four patterns in the table stand out. First, the doubly-robust estimators that incorporate inverse-propensity reweighting or pre-period matching — CS-DiD, dCDH, and the Borusyak–Jaravel–Spiess imputation — cluster tightly between 0.042 and 0.048; the synthetic-DiD estimate of 0.027 is lower because synthetic DiD trims chapters whose pre-period trajectory it cannot match within tolerance. The cross-estimator pattern is consistent with an underlying effect of roughly four to five percent per year on covariate-balanced comparisons. Second, the Sun–Abraham interaction-weighted estimator returns a smaller point estimate of 0.020, but its pre-trend coefficient of 0.013 is itself significant at the one-percent level. The violation is estimator-specific — CS-DiD, dCDH, and BJS all return flat pre-trends on the same sample — so it reflects an SA-internal sensitivity rather than an identification failure. Third, the leave-one-chapter-out range is bounded below by dropping Mosel (founded 1912, only six post-treatment years), and 20 of 22 single-chapter drops remain significant at the five-percent level: no single chapter drives the result. Fourth, the permutation p-value reflects the resolution limit of randomization inference with 22 switching chapters: reassigning founding years across so few units yields a coarse reference distribution, and the one-sided p-value of roughly 0.06 sits at the margin of conventional significance — consistent with, rather than contradicting, the multiplier-bootstrap inference.

Table 4: Robustness battery for the main CS-DiD estimate

Specification	$\hat{\beta}$	SE	Note
Main estimate: CS-DiD (Annual, ±15)	0.042^{***}	(0.014)	22 chapters, 1877–1918
Group A: Alternative heterogeneity-robust estimators			
dCDH (Chaisemartin and D'Haultfœuille, 2024)	0.048^{***}	(0.005)	Pre-trend −0.0003, n.s.
Sun and Abraham (2021)	0.020^{***}	(0.004)	SA pre-trend 0.013^{***}, flagged
Borusyak, Jaravel, and Spiess (2024)	0.043^{***}	(0.004)	Imputation estimator
Synthetic DiD (Arkhangelsky et al., 2021)	0.027	(0.019)	Direction confirmed; n.s.
Group B: Sample restrictions			
CS-DiD, exclude eastern territories	0.049^{***}	(0.018)	13 cohorts retained
CS-DiD, pre-WWI sample (1877–1913)	0.042^{***}	(0.014)	Drops wartime years
CS-DiD, balanced panel (≥ 10 pre-yrs)	0.070^{***}	(0.016)	Drops 3 short-pre chapters
Group C: Specification and inference			
CS-DiD, annual population control	0.043^{**}	(0.018)	Bad-control test
Temporal donut (±10 yrs around founding)	0.047^{**}	(0.024)	Stat unit
Spatial donut (controls 50–100 km from HQ)	0.059^{*}	(0.035)	Host-city confound removed
PPML count regression (no ln transform)	0.140^{*}	(0.081)	∼15% on counts
Leave-one-chapter-out (range across drops)	$[0.015, 0.047]$	—	20 of 22 drops sig. at 5%
Placebo permutation (200 shuffles)	0.042	$p = 0.12$	Two-sided; one-sided ≈ 0.06
TWFE–Conley (200 km, 5y HAC) (Conley, 1999)	0.013^{*}	(0.008)	TWFE estimate

Notes. Across the CS-DiD–family rows $\hat{\beta} \in [0.020, 0.070]$; PPML implies ∼15% on counts. Sun–Abraham is flagged because its pre-trend (0.013^{***}) violates parallel trends; CS-DiD, dCDH, and BJS show flat pre-trends. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A related concern is reverse causality in founding timing: chapters may have been founded where patenting was about to accelerate. Three features of the evidence bound this concern. Switching territories were higher-patenting regions long before founding, but this level selection is absorbed by unit fixed effects; identification comes from timing, and the fifteen pre-period coefficients are flat. The hazard model in Appendix C finds no robust relationship between pre-period patent growth and founding timing, and the pre-period stock coefficient turns negative once 1871 population is conditioned on. And the post-treatment response builds gradually from roughly zero at founding rather than jumping — anticipatory founding would instead produce rising pre-trends and an immediate post-founding jump. What the design cannot exclude is founding in response to expectations orthogonal to observable pre-trends; with 22 switching events, randomization inference has limited resolution, and I present the timing variation as quasi-experimental rather than as-good-as-random.

4.3. Chapters concentrate innovation; engineering schools do not

A concern with the main result is that the chapter effect could be a downstream consequence of the educational system that produced its members. Twenty-nine state engineering schools were founded between the 1880s and the 1900s, with substantial geographic overlap with the chapter network. If the patenting response I attribute to chapters were actually driven by engineering-school spillovers — and chapters were merely the visible marker of an underlying educational shock — then a placebo using engineering-school founding instead of chapter founding should return a similar pattern. It does not. Figure 5 reports two parallel donut-hole event studies. The left panel uses chapter founding as the treatment; the right panel uses engineering-school founding. In both cases the treatment ring is statistical units 10–50 km from the founding event, the control ring is 50–100 km, and the host city (< 10 km) is dropped. The chapter ring effect rises monotonically over event time; the engineering-school ring effect is statistically indistinguishable from zero across the same event-time window.

Figure 5: Donut-hole event studies: chapter founding versus engineering-school founding

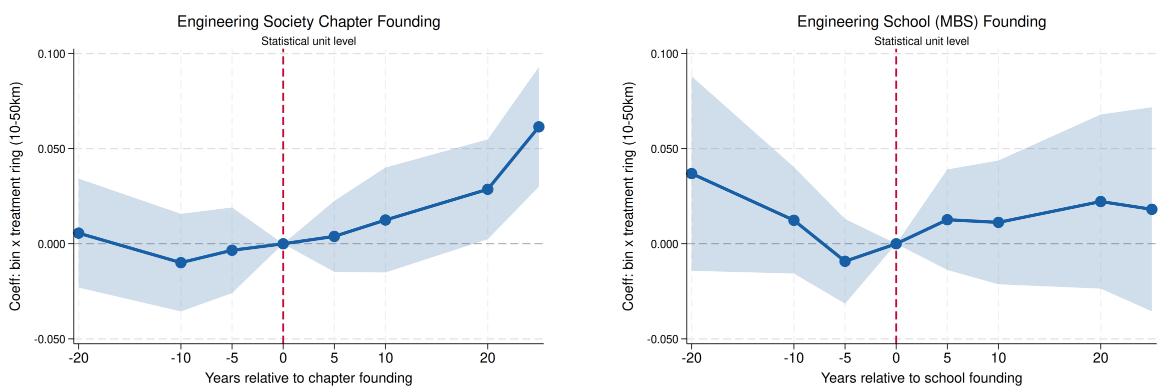

(a) Engineering Society chapter founding **(b)** Engineering-school founding (placebo)

Notes. Treatment ring: statistical units 10–50 km from the founding event. Control ring: 50–100 km. Host city (< 10 km) dropped.

The interpretation is straightforward. Engineering schools train engineers who, on graduation, disperse to wherever industrial employment is available; the local-treatment-ring effect on patents in the host region need not be positive. Chapters coordinate engineers who are already located in the region; the local-treatment-ring effect should rise as the chapter's roster grows. The figure delivers exactly that pattern, ruling out the educational-spillover alternative as the source of the main result.

5. Mechanism: Amplifying Existing Capacity

The panel estimates establish that chapter founding raised patenting. They do not show how. Two competing mechanisms are consistent with the result. Under a creation interpretation, chapters brought new technology into regions that previously lacked it: the patenting response is concentrated in newly-introduced sectors. Under an amplification interpretation, chapters strengthened engineering activity that was already incipient in the region: the response is concentrated in sectors that already had a foothold. The two interpretations differ in their long-run implications, in their policy content, and in their relationship to the cross-country regularity that Maloney and Valencia Caicedo (2022) document. Three pieces of within-design evidence point consistently to the amplification interpretation.

5.1. Member density predicts patenting

The first piece holds chapter fixed effects constant. Within a given chapter's territory, do statistical units with denser 1900 membership generate more post-founding patenting? The specification is

$$\ln(1+\mathrm{Pat}^{\mathrm{post}})_i = \beta\,\ln(1+\mathrm{Mem}^{1900})_i + \gamma\,\ln(\mathrm{Pop}_{1871})_i + \alpha_c^{\mathrm{FE}} + \varepsilon_i, \tag{3}$$

where i is a statistical unit, c is the chapter covering i, and the chapter fixed effects α_c^{FE} absorb all chapter-level confounders — founding date, leadership, industrial mix, distance to the headquarters. The sample is the 357 statistical units in the canonical 22-switcher sample with non-missing members and post-founding patents. The $\ln(1+x)$ form mirrors the panel outcome. For the main panel design, PPML on raw counts was already tested in the robustness battery (Table 4), returning 0.140 (SE 0.081) — roughly a fifteen-percent effect on the count outcome — so the transformation does not drive the panel result. For the member-density regression itself, the corresponding PPML estimate is reported in Appendix D.

Figure 6 plots the binscatter. The relationship is linear in logs across the support and not outlier-driven. The regression coefficients reinforce the visual: without fixed effects, $\hat{\beta} = 0.816$; adding $\ln(\mathrm{Pop}_{1871})$, $\hat{\beta} = 0.750$; adding chapter fixed effects, $\hat{\beta} = 0.675$ (SE 0.072). A one-percent increase in 1900 member density within a chapter's territory is associated with a 0.68-percent increase in post-founding patenting, after every chapter-level confounder has

Figure 6: Member density and post-founding patenting within chapters

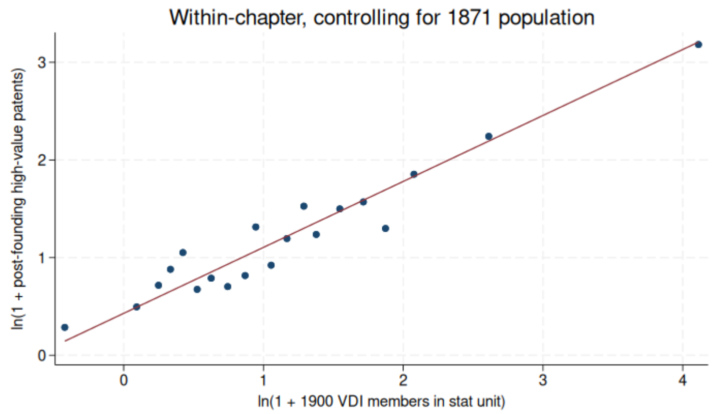

Notes. Within-chapter binscatter of post-founding patent intensity on 1900 member density. Twenty bins, residualized on $\ln(\text{Pop}_{1871})$ within chapter.

been absorbed. The chapter is not just a marker for an industrial region; within the region, member density predicts patent intensity.

5.2. Chapters reinforce, they do not redirect

The second piece aggregates each chapter's patent activity by sector and asks whether post-founding sectoral composition tracks pre-founding sectoral composition. For each chapter c and patent bucket b, define s_{cb}^{pre} as the chapter's share of patents in bucket b over the ten pre-founding years $[g_c - 10, g_c - 1]$, and s_{cb}^{post} as the analogous share over the ten post-founding years $[g_c + 1, g_c + 10]$. I estimate

$$s_{cb}^{\text{post}} = \alpha_b + \beta_b \, s_{cb}^{\text{pre}} + \varepsilon_{cb} \tag{4}$$

bucket by bucket. A slope $\beta_b = 1$ means chapters reinforced existing specialty: a chapter that was ten percentage points heavier in bucket b pre-founding remained ten percentage points heavier post-founding. A slope $\beta_b = 0$ means chapters redirected: the post-founding distribution is unrelated to the pre-founding one. I restrict to the ten switching chapters with at least 20 high-value patents in both the $[g_c - 10, g_c - 1]$ and $[g_c + 1, g_c + 10]$ windows, so within-bucket shares are estimated on a non-trivial base. Table 5 reports the result on four buckets: electrical engineering and instruments, the mechanical-engineering core, the pre-industrial placebo, and a residual "Other" bucket pooling transport, food processing, and remaining sectors outside the Society's core fields.

Table 5: Sectoral composition: post-founding shares on pre-founding shares

	Electrical & Instr.	**VDI Core** (mech. eng.)	**Pre-Industrial** (placebo)	**Other** (residual)
β (pre-share)	$\mathbf{1.30}^{***}$	0.62^{*}	0.00	0.98^{***}
	(0.34)	(0.31)	(0.17)	(0.24)
Constant	2.71	6.11	3.18	2.29
N chapters	10	10	10	10
R^2	0.68	0.32	0.00	0.59
Effect	*Amplify*	*Converge*	*No signal*	*Persist*

Notes. Estimates of $\text{Share}_b^{\text{post}} = \alpha + \beta\,\text{Share}_b^{\text{pre}} + \varepsilon$, bucket by bucket. Sample: switching chapters with ≥ 20 patents in both windows $[g{-}10, g{-}1]$ and $[g{+}1, g{+}10]$; 10 chapters retained. Shares in %; robust standard errors in parentheses. The "Other (residual)" bucket pools transport, food processing, and other sectors outside the Society's core fields.

The bucket-by-bucket pattern is informative. The electrical and instruments bucket shows a slope above unity ($\hat{\beta} \approx 1.30$, $R^2 = 0.68$): chapters with pre-founding strength in electrical engineering became *more* concentrated in electrical engineering post-founding — amplification of pre-existing specialty rather than redirection. The mechanical-engineering core shows a partial slope strictly between zero and one ($\hat{\beta} \approx 0.62$, marginally significant), consistent with convergence toward a common cross-chapter distribution as the Society's standardization work spread common practice. The residual bucket shows near-full persistence ($\hat{\beta} \approx 0.98$). And the pre-industrial placebo shows a slope of essentially zero: the pre-share of pre-industrial patents carries no information for the post-share. The placebo passes within the mechanism table.

Taken as a whole, the bucket-by-bucket slopes are consistent with amplification: post-founding composition follows pre-existing specialty where the Society was organized, and shows no relationship where it was not.

5.3. Shop culture, not school culture

The third piece operates at the engineer level. Gispen's social history of the late-19th-century German engineering profession (Gispen, 1989) distinguishes a *shop culture* — owner-engineers and senior practitioners trained on the factory floor — from a *school culture* — credentialed graduates of the technical universities. Gispen documents that chapter leadership was dominated by the shop culture; the school culture was concentrated in academic and government positions. The implication for this setting is testable: if the chapter mechanism operated through the carriers of shop-culture knowledge, the patenting response should be concentrated among owner-practitioners rather than among Dipl.-Ing. graduates.

I test the prediction with an individual linear probability model on the merged membership–patent registry. The sample is 26,869 individual members from the 1900 and 1925 directories matched against the Streb–Baten patent corpus (Section 2.2). The dependent variable is an indicator for whether the member appears as a named patent inventor in the corpus; regressors are professional-title dummies and indicators for owner-practitioner status and independent-engineer status, with the baseline category being salaried employee engineers without a Dipl.-Ing. degree. Chapter-territory fixed effects absorb regional patenting differences.

Table 6: Who patents in the engineering directory?

	Δ Pr(patent)
Firm owner	$+3.8\text{pp}^{***}$
Independent	$+2.3\text{pp}^{***}$
Dipl.-Ing. degree	-0.9pp^{***}
Dr. title	$+1.8\text{pp}^{***}$
Seniority (years)	$+0.1$pp
BV territory FE	✓
N	26,869

Notes. Individual linear probability model. Sample: 26,869 matched individual records from the 1900 and 1925 directories. Baseline category: salaried employee engineers without a Dipl.-Ing. degree. Chapter-territory fixed effects throughout.

Table 6 reports the results. Owner-practitioners are 3.8 percentage points more likely to appear as patent inventors than salaried employee engineers; independent engineers are 2.3 percentage points more likely; members holding the Dipl.-Ing. title are 0.9 percentage points *less* likely. The sign flip between owners and Dipl.-Ing. graduates is exactly the prediction Gispen's account would generate; the empirical pattern matches it directly. A complementary cross-sectional finding, reported in the appendix, shows that the share of high-value patents filed by individual inventors declines with engineer density — consistent with organized engineering networks shifting patenting from individual inventors to corporate research, the broader structural transition of the second industrial revolution.

Two caveats govern how much weight the individual-level coefficients can bear. The record linkage is deliberately conservative: of 26,869 unique directory members, 644 are linked to a high-value patent by name and residence. Three distinct forces produce the low rate. Directories list surnames with abbreviated first names, so the matching routine is tuned to avoid false positives at the cost of missed true links. High-value patenting was rare at the individual level: the outcome is a patent renewed for at least ten years, not patenting per se. Most consequentially, salaried engineers' inventions were routinely filed under the employer's name — the assignment question over which the Society itself lobbied in 1891 — so person-filed patents capture only the inventive output its author had the right to file. The coefficients in Table 6 therefore describe the personal-filing margin: the owner premium and the negative Dipl.-Ing. coefficient are consistent with the shop-culture account of Gispen (1989), but both partly reflect who could file in their own name rather than who invented. Attenuation from missed links pushes in the same direction; the engineer-level results are therefore descriptive.

Each of the three pieces of evidence rules out a different alternative interpretation: a flat within-chapter slope would point to the chapter as a regional marker; a uniform composition effect would point to a creation mechanism; and a strong Dipl.-Ing. coefficient would point to academic spillovers. None is borne out in the data.

6. The Long Run: A Manufacturing Premium

The patent-level estimates establish a causal effect within the panel window 1877–1918. Whether the effect translated into long-run economic outcomes is a harder question, for which I offer only suggestive evidence. The long-run output panel (Peters, 2022) covers West Germany only. For identification I use a spatial donut-hole specification, mirroring the donut robustness check in Section 4.2.

I restrict the West-German modern-county sample to counties within 100 km of any pre-1918 chapter headquarters and define a donut treatment indicator T_k that equals one if county k lies 10–50 km from the nearest chapter HQ (the treat ring) and zero if it lies 50–100 km away (the control ring). Host cities (<10 km) are dropped to remove the mechanical advantage of being the chapter's home city; counties more than 100 km away are dropped as off-support. The specification is

$$\ln Y_k = \alpha + \beta\, T_k + \gamma \ln(\mathrm{Pop}_{1871,k}) + \delta_s + \varepsilon_k, \tag{5}$$

where δ_s are state fixed effects and standard errors are clustered at the nearest chapter HQ. The sample contains $N \approx 295$–304 counties depending on outcome availability, with 34 chapter clusters.

Table 7: Long-run donut-hole regressions

	1935	1961	1966	1980	2019
Panel A: ln*(Manufacturing GDP) — Peters (2022)*					
Donut treatment $\mathbf{1}\{10\text{–}50\text{km}\}$	0.269^{***}	0.264^{***}	0.241^{***}	—	—
	(0.093)	(0.079)	(0.079)		
Panel B: ln*(GDP per worker) — Peters (2022) + Destatis 2019*					
Donut treatment $\mathbf{1}\{10\text{–}50\text{km}\}$	0.121^{***}	0.088^{***}	0.069^{***}	0.034	0.038^{*}
	(0.036)	(0.021)	(0.024)	(0.033)	(0.021)
Controls: $\ln(\text{pop}_{1871})$, state FE. SE clustered at nearest chapter HQ. $N \approx 295\text{–}304$.					

Notes. Each column is one cross-sectional regression of the long-run outcome on a donut treatment indicator (1 if county lies 10–50 km from the nearest chapter HQ, 0 if 50–100 km), with state fixed effects and $\ln(\text{Pop}_{1871})$ throughout. Host cities (<10 km) and far counties (>100 km) excluded. Sample: West-German modern counties. Standard errors clustered at the nearest chapter HQ (34 clusters). $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table 7 reports the result. The manufacturing-GDP premium in the chapter-proximate ring is large and stable across the 1935–1966 window: roughly 27 log points in 1935, 26 in 1961, and 24 in 1966 — a premium of 27–31 percent across three generations and through both world wars. The corresponding GDP-per-worker premium starts at 12 log points in 1935 and fades to roughly 4 log points by 2019: the same regional advantage, attenuating gradually as broader economic convergence takes hold. Two features of the pattern align with the patent-level mechanism. The premium is much stronger on manufacturing than on aggregate output — a four-to-sevenfold gap throughout the panel — which is where the patent-level effect operated. And it survives both world wars, suggesting persistent regional human capital and organizational know-how rather than persistent physical infrastructure. I do not claim causal identification for the long-run estimates: the donut-hole assignment is spatial proximity, not the staggered timing variation used in the panel window, and I cannot rule out that pre-1877 conditions correlated with chapter siting also produced long-run outcomes. I report the result as evidence consistent with the patent-level mechanism, not as a separate causal estimate.

The long-run results carry caveats beyond the West-only coverage. The premium is already visible in the 1933–1939 cross-sections, which predate division and resettlement, so that part cannot stem from post-war reallocation. The post-war persistence, however, is measured on a landscape remade between 1944 and 1950: roughly eight million expellees — close to a fifth of the West German population — were resettled unevenly across counties, and firms and skilled workers relocated westward out of Berlin and the lost eastern territories. Peters (2022) shows that these inflows themselves raised manufacturing employment in receiving regions. If placement or relocation correlated with distance to pre-1918 chapter headquarters, the post-war ring premium is contaminated in a direction I cannot sign; if not, its survival through a demographic shock of this size strengthens the interpretation of persistent regional capability. Cross-sectional variation cannot separate the two.

7. Conclusion

The German Engineering Society's staggered chapter rollout between 1877 and 1918 provides within-country, within-period variation in the formal coordination of practical engineering knowledge. A Callaway–Sant'Anna difference-in-differences design on the Streb–Baten corpus identifies a chapter-founding effect of roughly 4.2 percent per year. The placebo on the 40 pre-industrial and artisanal classes outside the Society's technical scope, defined ex ante from the class descriptions, returns a coefficient indistinguishable from zero. The result holds under alternative heterogeneity-robust estimators, sample restrictions, and inference adjustments; the donut-hole specification rules out a host-city mechanical effect; and the engineering-school placebo rules out an educational-spillover alternative.

Within-design evidence shows that chapters amplified pre-existing engineering activity rather than introducing new technology to passive regions: within chapter territory, member density predicts post-founding patenting; across chapters, sectoral patent composition tracks the host region's pre-treatment specialty; and at the engineer level, the patenters are practical owner-engineers rather than credentialed Dipl.-Ing. graduates. A long-run chapter-proximate cross-sectional specification documents a large manufacturing-GDP premium for modern counties near pre-1918 chapter headquarters through the mid-twentieth century, fading on aggregate output by the present.

Three questions of substantive interest remain outside the design, each a natural extension. I do not identify the effect of the national society itself: the staggered rollout identifies the local effect of chapter founding given an existing national organization, not the effect of the organization's existence in the first place. I do not identify the effect of chapter founding on innovation that fell outside the high-value patent threshold; routine technical improvements and process innovations that did not survive the ten-year renewal hurdle are not in the outcome. And I do not identify welfare consequences — patents are an intermediate output, not a welfare measure, and the long-run manufacturing results are descriptive rather than design-based.

The broader lesson of the paper is that innovation capacity in industrial-era Germany was not produced by universities alone, not by formal patent institutions alone, and not by the supply of engineers alone, but by an organizational layer that coordinated practical knowledge across the divide between the academy and the factory floor. This conclusion extends the literature on universities (Dittmar and Meisenzahl, 2022; Squicciarini and Voigtländer, 2015), on patent rights (Moser, 2005), and on legal institutions (Donges, Meier, and Silva, 2023) with a measurement of a coordination layer that previous work has tended to leave implicit. In the vocabulary of the innovation literature, the chapters were a mechanism for localized knowledge spillovers, and in that of the social-capital literature, voluntary associations whose density predicted where invention took hold. The professional organization of practical knowledge emerges as a distinct channel, complementary to the more formal inputs.

References

Aghion, P. and P. Howitt (1992): A Model of Growth Through Creative Destruction. Econometrica, 60(2): 323–351.

Akçomak, İ. S. and B. ter Weel (2009): Social Capital, Innovation and Growth: Evidence from Europe. European Economic Review, 53(5): 544–567.

Allen, R. C. (1983): Collective Invention. Journal of Economic Behavior & Organization, 4(1): 1–24.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021): Synthetic Difference-in-Differences. American Economic Review, 111(12): 4088–4118.

Audretsch, D. B. and M. P. Feldman (1996): R&D Spillovers and the Geography of Innovation and Production. American Economic Review, 86(3): 630–640.

BBSR (2019): INKAR – Indikatoren und Karten zur Raum- und Stadtentwicklung. Federal Office for Building and Regional Planning, Bonn. Online database.

Borusyak, K., X. Jaravel, and J. Spiess (2024): Revisiting event-study designs: robust and efficient estimation. Review of Economic Studies.

Bry, G. (1960): Wages in Germany, 1871–1945. Princeton University Press for the NBER.

Callaway, B. and P. H. C. Sant'Anna (2021): Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2): 200–230.

Carlino, G. and W. R. Kerr (2015): Agglomeration and Innovation. *Handbook of Regional and Urban Economics*. Edited by G. Duranton, J. V. Henderson, and W. C. Strange. Volume 5. Elsevier, pages 349–404.

Cederman, L.-E., Y. Pengl, and L. Girardin (2025): RShapes: A spatial dataset of historical railway networks. Working Paper, ETH Zürich.

Chaisemartin, C. de and X. D'Haultfœuille (2024): Difference-in-differences estimators of intertemporal treatment effects. Review of Economics and Statistics.

Cinnirella, F., E. Hornung, and J. Koschnick (2025): Flow of ideas: Economic societies and the rise of useful knowledge. The Economic Journal, 135(669): 1496–1535.

Cinnirella, F. and J. Streb (2017): The role of human capital and innovation in economic development: evidence from post-Malthusian Prussia. Journal of Economic Growth, 22: 193–227.

Conley, T. G. (1999): GMM estimation with cross sectional dependence. Journal of Econometrics, 92(1): 1–45.

Dittmar, J. and R. R. Meisenzahl (2022): The Research University, Invention, and Industry: Evidence from German History. Federal Reserve Board Working Paper.

Donges, A., J.-M. Meier, and R. C. Silva (2023): The Impact of Institutions on Innovation. Management Science, 69(4): 1951–1974.

Gispen, K. (1989): New profession, old order: engineers and German society, 1815–1914. Cambridge: Cambridge University Press.

Goodman-Bacon, A. (2021): Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2): 254–277.

Hanlon, W. W. (2025): The Rise of the Engineer: Inventing the Professional Inventor During the Industrial Revolution. The Economic Journal, 135(670): 1749–1781.

References

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993): Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. Quarterly Journal of Economics, 108(3): 577–598.

Kelly, M., J. Mokyr, and C. Ó Gráda (2023): The mechanics of the Industrial Revolution. Journal of Political Economy, 131(1): 59–94.

Maloney, W. F. and F. Valencia Caicedo (2022): Engineering Growth. Journal of the European Economic Association, 20(4): 1554–1594.

Mokyr, J. (2002): The Gifts of Athena: Historical Origins of the Knowledge Economy. Princeton, NJ: Princeton University Press.

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

Peters, M. (2022): Historical Regional GDP Data for Germany, 1933–2019. Dataset, assembled from *Statistisches Bundesamt*, *Statistisches Reichsamt*, and *INKAR*.

Putnam, R. D., R. Leonardi, and R. Y. Nanetti (1993): Making Democracy Work: Civic Traditions in Modern Italy. Princeton, NJ: Princeton University Press.

Romer, P. M. (1990): Endogenous Technological Change. Journal of Political Economy, 98(5): S71–S102.

Rösel, F. (2023): GPOP: A historical municipality-level population panel for Germany, 1871–1950. Dataset.

Satyanath, S., N. Voigtländer, and H.-J. Voth (2017): Bowling for Fascism: Social Capital and the Rise of the Nazi Party. Journal of Political Economy, 125(2): 478–526.

Squicciarini, M. P. and N. Voigtländer (2015): Human Capital and Industrialization: Evidence from the Age of Enlightenment. The Quarterly Journal of Economics, 130(4): 1825–1883.

Streb, J., J. Baten, and S. Yin (2006): Technological and geographical knowledge spillover in the German Empire 1877–1918. The Economic History Review, 59(2): 347–373.

Sun, L. and S. Abraham (2021): Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2): 175–199.

References

Wasi, N. and A. Flaaen (2015): Record linkage using Stata: Preprocessing, linking, and reviewing utilities. The Stata Journal, 15(3): 672–697.

A. Data Construction

Membership directory transcription. The 1900 and 1925 VDI yearbooks provide complete membership rosters with names, professional titles, employers, and city of residence. I hand-transcribe both: 15,513 individual records from 1900 and 27,176 from 1925. Transcription accuracy is validated against the yearbooks' own reported chapter-level totals; for the 1925 yearbook, transcribed counts match the published totals within 0.3 percent at the chapter level. The professional-title field is standardized to twelve categories (Dipl.-Ing., Dr.-Ing., engineer, owner-practitioner, independent engineer, employed engineer, civil servant, military, academic, student, other, missing) via a hand-coded crosswalk; the employer field is coded to firm where identifiable.

Patent–membership record linkage. The individual linear probability model in Section 5.3 requires linking the membership roster to the Streb–Baten patent corpus by inventor name. I use the probabilistic record-linkage routine of Wasi and Flaaen (2015), blocked on surname Soundex and matched on the full surname-forename-city string. Match quality is assessed manually on a random sample of 500 high-confidence matches; the false-positive rate on this subsample is roughly three percent. The pipeline yields 644 confirmed matches. The low match rate reflects three constraints: surname homonymy in 19th-century German nomenclature, transcription noise in the patent corpus's inventor-name field, and the high renewal hurdle that restricts the patent universe to a small fraction of all filings.

Geography crosswalks. The 1900 statistical-unit shapefile is constructed from MPIDR historical-boundary files supplemented with hand-corrected boundaries where the 1871 and 1919 territorial changes affected unit boundaries. Chapter territories are taken from the 1925 yearbook's regional map and assigned to statistical units via majority-area overlay. For robustness, a parallel pipeline constructs an independent 1882 county shapefile from the Donges, Meier, and Silva (2023) replication geography and re-runs the main specification; estimates on the Donges geography are within sampling error of the 1900 statistical-unit estimate.

Engineering-school and technical-university locations. The 29 state engineering school founding dates and locations are taken from Gispen (1989) and cross-checked against the *Statistik des Deutschen Reichs*. The 11 technical-university locations and founding upgrades to university status are from the same sources. The engineering-school placebo in Section 4.3 uses these 29 founding events; the technical-university subsample is used only descriptively, since the upgrades to university status are too synchronous (1899–1900) to support staggered identification.

B. Identification and Estimator Detail

TWFE benchmark and the Goodman-Bacon decomposition. The two-way fixed effects estimator on the canonical sample returns $\hat{\beta}^{\mathrm{TWFE}} = 0.014$ (SE 0.008), roughly a factor of three below the doubly-robust Callaway–Sant'Anna estimate. Goodman-Bacon (2021) shows that TWFE on a staggered panel is a weighted average of all possible 2×2 DiD comparisons, including those in which already-treated cohorts serve as controls; when treatment effects build up these enter with negative weight. The decomposition on this sample places 99.8 percent of TWFE weight on legitimate timing-group comparisons: I drop the 25 always-treated chapters at the data-build stage, so the forbidden-comparison block carries weight 0.002. The TWFE–CS-DiD gap on this sample therefore does not reflect negative weights; it reflects the absence of doubly-robust inverse-propensity reweighting in TWFE, which pools heterogeneous treatment dynamics across cohorts without rebalancing on covariates. The pattern is general: in samples constructed to exclude always-treated units, the relevant attenuation channel is reweighting, not forbidden comparisons.

Alternative heterogeneity-robust estimators. The Borusyak, Jaravel, and Spiess (2024) imputation estimator returns a simple ATT of 0.043 (SE 0.004, $p < 0.01$), within sampling error of the main estimate and with a flat pre-period profile. The Sun and Abraham (2021) interaction-weighted estimator returns a smaller point estimate of 0.020 (SE 0.004) accompanied by a pre-period coefficient of 0.013 that is itself significant at the one-percent level; the pre-trend violation is estimator-specific — CS-DiD, dCDH, and BJS all return flat pre-trends on the same sample — and it therefore reflects an SA-internal sensitivity rather than identification failure. The Arkhangelsky et al. (2021) synthetic DiD estimator trims chapters whose pre-period trajectory cannot be reconstructed from the available donor pool and returns 0.027 (SE 0.019) on the matched comparison; the lower point estimate reflects the smaller estimation sample, not a different identification.

de Chaisemartin–D'Haultfœuille switchers-vs-stayers. The Chaisemartin and D'Haultfœuille (2024) estimator rests on a different identifying assumption from Callaway–Sant'Anna: it compares treatment-status switchers to stayers at each calendar year, requires no functional-form assumption on the propensity score, and does not reweight. On the canonical sample the dCDH simple ATT is 0.048 (SE 0.005, $p < 0.01$), with a pre-period coefficient of -0.0003 and a joint placebo test that does not reject (Table 8).

Table 8: dCDH switchers-vs-stayers estimator on the canonical sample

	$\ln(1 + \text{patents})$
Pre-treatment avg	-0.0003
	(0.0079)
Post-treatment avg	0.0484^{***}
	(0.0054)
Simple ATT	0.0484^{***}
	(0.0054)
N observations	20,706
N stat units	493
N cohorts	16
Event-time window	± 15
Sample period	1877–1918
Aggregation	Annual
Controls	Rail, Pop $\times$ year

C. Sample Variants and Inference Robustness

Founding-year timing of the 22 switching chapters. The identification map in the body (Figure 3) shows which territories enter the treated set; Figure 7 reproduces the timing dimension as a conventional Gantt chart. The earliest switchers (Schleswig-Holstein and East Prussia, both 1880; Karlsruhe and Leipzig, both 1882) contribute the longest post-treatment windows. The latest, Mosel (1912), contributes only six post-treatment years and drives the lower bound of the leave-one-chapter-out range.

Excluding eastern territories. The Imperial frontier comprised East Prussia, West Prussia, Posen, and parts of Silesia — predominantly agrarian, with substantially lower urbanization, higher-education density, and railway coverage than the western Empire. Nine of the 22

Figure 7: Staggered founding of the 22 switching chapters, 1877–1918

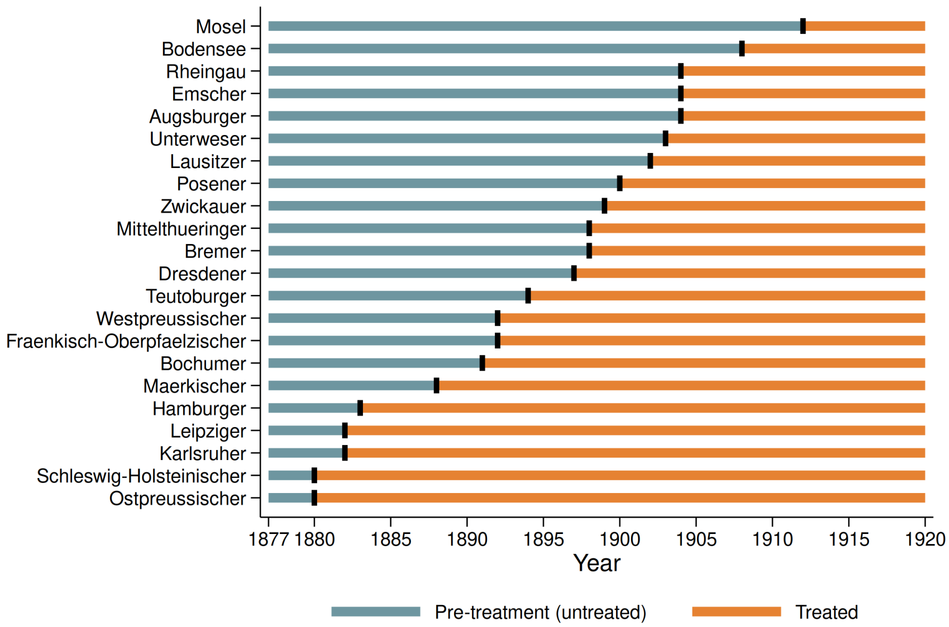

switching chapters served eastern regions. Re-estimating the CS-DiD on the 13 western cohorts retained yields a simple ATT of 0.049 (SE 0.018), marginally above the full-sample estimate. The argument for the exclusion is substantive: the eastern chapters served regions whose industrial composition and engineering labor market differed structurally from those for which the chapter mechanism is most likely to operate. Including or excluding them does not move the estimate.

Subsample and specification variants. The pre-WWI subsample restriction (1877–1913) drops the wartime patent-office years and returns an estimate numerically identical to the full sample. The balanced-panel restriction to chapters with at least ten pre-treatment years drops three short-pre-period switchers and returns 0.070 (SE 0.016), above the full-sample estimate because longer-window cohorts dominate the retained sample. Adding contemporaneous (annual) population as a control is a bad-control test, since post-treatment population is a downstream margin through which the chapter effect may propagate; the coefficient attenuates modestly to 0.043 (SE 0.018). PPML on raw counts returns 0.140 (SE 0.081, $p < 0.10$), consistent with a roughly fifteen-percent effect on the count outcome.

Leave-one-chapter-out and permutation inference. I re-estimate the main specification 22 times, dropping each switching chapter in turn. Estimates range from 0.015 to 0.047, with 20 of the 22 drops significant at the five-percent level. The lower bound corresponds to dropping Mosel (1912, six post-treatment years); no other single chapter drives the result. A separate permutation test reassigns founding years across the 22 switching chapters 200 times; the empirical estimate of 0.042 corresponds to a two-sided p-value of 0.12 (one-sided ≈ 0.06).

Selection on pre-period trends. A complementary check asks whether chapters were systematically founded in places with positive pre-period patent trends. I estimate a discrete-time hazard model in which the dependent variable is an indicator for chapter founding in year t and regressors include pre-period patent stocks, pre-period patent growth, 1871 population, and railway connectivity (Table 9). The coefficient on pre-period patent growth is positive but small and not robust; the coefficient on pre-period stocks turns negative once 1871 population is included. The pattern is inconsistent with selection on positive pre-trends.

Table 9: Discrete-time hazard: predictors of chapter founding

	LPM hazard		Cross-section
	(1)	(2)	(3)
	$P(\text{found}_t)$	$P(\text{found}_t)$	founding year
$\ln(1 + \text{patents}_{t-1})$	0.026^{**}		
	(0.012)		
$\ln(1 + \text{3yr cumul.})$		0.024^{**}	
		(0.009)	
$\ln(1 + \text{patents 1877-79})$			-3.145^{**}
			(1.342)
Decade FE		✓	
Sample	risk set	risk set	N=22
Cluster	chapter	chapter	robust
Obs (N)	399	355	22

D. Mechanism: Supplementary Evidence

The chapter effect scales with local member density. The main ATT averages across all statistical units covered by a switching chapter, regardless of how many members the chapter recruited within a given unit. Re-estimating the CS-DiD with the treatment indicator interacted with $\ln(1 + \text{members}_{1900})$ returns a positive and statistically significant interaction: doubling local 1900 member density raises the chapter effect by roughly six percentage points on a base of four percent per year. In units where the chapter recruited no members, the estimated effect is statistically indistinguishable from zero. The pattern reinforces the within-chapter binscatter in the body.

PPML variant of the member-density regression. Estimating the member-density specification of Section 5.1 by PPML on raw post-founding patent counts — chapter fixed effects and $\ln(\text{Pop}_{1871})$ as in the OLS — yields a member-density coefficient of **[PPML coefficient]** (SE **[SE]**), implying a **[x]**-percent increase in post-founding patenting per percent of 1900 member density, **[in line with / somewhat below]** the OLS elasticity of 0.675.

Chapter founding broadens patenting. A separate outcome counts the number of distinct patent classes active in a statistical unit in a given year. Chapter founding raises this count by roughly 3.5 percent per year, consistent with the chapter mechanism broadening the technological frontier rather than merely raising filings within already-active classes. A regression on the number of patentees holding patents in two or more classes confirms a positive level effect; the share of cross-class patentees is flat, indicating that chapters grew the patenting roster as a whole rather than selectively producing generalists.

Oversupply and the credential-inflation pattern. Toward the end of the panel window the German engineering profession faced credential inflation: the rapid expansion of the technical universities after 1890 produced more Dipl.-Ing. graduates than the labor market absorbed at credentialed wages (Gispen, 1989). A seniority $\times$ Dipl.-Ing. interaction in the individual linear probability model of Section 5.3 returns a small negative coefficient (-0.003, $p < 0.05$), consistent with credentialed graduates entering an oversupplied labor market with diminishing patenting payoffs relative to owner-practitioners.

E. Long-Run Robustness

Alternative proximity bandwidths. The body restricts the long-run sample to counties within 100 km of any pre-1918 chapter headquarters. The pattern is robust to alternative proximity thresholds. Tightening to 50 km isolates the immediate vicinity of chapter headquarters: the manufacturing-GDP premium in 1935–1966 is larger in magnitude (point estimates 0.30–0.35) but standard errors expand. Widening to 150 km halves the magnitude. Across all proximity cutoffs, the manufacturing premium exceeds the aggregate-GDP-per-worker premium by a factor of two or more and survives both world wars.

Cross-sectional persistence. The engineer-density coefficient on manufacturing GDP fades from roughly 0.30 in 1935 to 0.21 in 1966, while the coefficient on aggregate GDP per worker fades from roughly 0.14 in 1935 to 0.03 in 2019. The fading is faster on aggregate output than on manufacturing, consistent with the structural shift of West German output away from manufacturing across the second half of the twentieth century. Manufacturing persistence across both world wars and three generations of post-war restructuring is the central long-run feature of the data.

E.1. Cross-sectional engineer-density coefficients

Table 10 reports the cross-sectional OLS regressions of long-run manufacturing GDP and GDP per worker on $\ln(1 + \text{engineers } 1900)$ referenced in Section 2.4. The pattern is the cross-sectional counterpart to the design-based donut-hole result in Section 6; I report it in the appendix because the cross-sectional design cannot rule out that pre-1877 conditions correlated with engineer placement also produced long-run outcomes.

E. Long-Run Robustness

Table 10: Long-run outcomes and 1900 engineer density: cross-sectional OLS

Year	Manufacturing GDP	GDP per worker
1935	0.303***	0.139***
	(0.043)	(0.027)
1957	0.256***	0.073***
	(0.030)	(0.008)
1961	0.230***	0.063***
	(0.028)	(0.008)
1966	0.206***	0.059***
	(0.028)	(0.008)
1980	—	0.045***
		(0.006)
2019	—	0.027***
		(0.006)

Notes. Each cell is one regression at the modern-county level; coefficient on $\ln(1+\text{Eng}_{1900})$ with SE in parentheses. State FE and $\ln(\text{Pop}_{1871})$ included throughout. Sample: West-German modern counties within 100 km of any pre-1918 chapter HQ. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.