Factories of Ideas? Big Business and the Golden Age of American Innovation

Pier Paolo Creanza*

July 10, 2026

[Click here for the most recent version]

Abstract

This paper studies the Great Merger Wave (GMW) of 1895–1904—the largest consolidation event in U.S. history—to identify how Big Business affected American innovation. Between 1880 and 1940, the U.S. experienced a golden age of breakthrough discoveries in chemistry, electronics, and telecommunications that established its technological leadership. Using newly constructed data linking firms, patents, and inventors, I show that consolidation substantially increased innovation. Among firms already innovating before the GMW, consolidation led to an increase of 6 patents and 0.6 breakthroughs per year—roughly four-fold and six-fold increases, respectively. Firms with no prior patents were more likely to begin innovating. By contrast, consolidations that were announced but ultimately abandoned—"paper" mergers—saw no comparable increase, supporting a causal interpretation. The establishment of corporate R&D laboratories served as a key mechanism driving these gains. Building a matched inventor–firm panel, I show that lab-owning firms enjoyed a productivity premium not due to inventor sorting, robust within size and technology classes. To assess whether firm-level effects translated into broader technological progress, I examine total patenting within technological domains. Overall, the GMW increased breakthroughs by 13% between 1905 and 1940, with the largest gains in science-based fields (30% increase).

*Princeton University: pcreanza@princeton.edu.

I owe a large debt of gratitude to Naomi Lamoreaux for sharing with me the scans of Ralph Nelson's GMW worksheets, which made this paper possible. I also thank Enrico Berkes, Anne Marie Knott, Konstantin Kosenko, Natalya Vinokurov, and Yishay Yafeh for generously sharing their data. I am deeply grateful to my dissertation committee—Leah Boustan, Ilyana Kuziemko, and Owen Zidar—for their invaluable support and steady guidance. Thanks also to Zachary Bleemer, Pietro Buri, Kaan Cankat, Pauline Carry, Ellora Derenoncourt, Allison Green, Simon Jäger, Amy Kim, David Lee, Franco Malpassi, Daniel Milo, Kobi Mizrahi, Myera Rashid, Bernardo Ribeiro, João Tampellini, and Danielle Graves Williamson, for their helpful comments. I further thank participants at the Princeton Labor Workshop, the Yale Economic History Lunch, the 2025 NBER Summer Institute (Development of the American Economy), and the 2025 EHA Annual Meeting. Financial support from the Princeton Industrial Relations Section and the Economic History Association is also gratefully acknowledged.

1 Introduction

Did Big Business power the golden age of American innovation? U.S. productivity growth peaked in the first half of the twentieth century, propelled by technological progress (Field 2003; 2013; Gordon 2016). Influential views credit two institutions that rose together at the turn of the century: the modern corporation and its industrial research laboratory. Schumpeter (1942) casts the large enterprise as the most powerful engine of technological progress, Chandler (1977) places the large managerial enterprise at the center of America's industrial development, and DeLong (2022, p. 3) counts both among "the keys" that unlocked modern economic growth. Yet these narratives rest largely on qualitative evidence, and their central causal claim is difficult to test: innovative firms tend to *become* big, so correlations between scale and innovation are hard to interpret, and major exogenous shifts in market structure are rare (Cohen 2010; Bryan and Williams 2021). This paper exploits one such transformational episode.

The Great Merger Wave (GMW) of 1895–1904 was the largest merger and acquisition event in U.S. economic history. At its peak in 1899, merged assets totaled 12.5 percent of GNP—equivalent to US$(2024) 3.7 trillion today—as more than 2,600 firms combined to form corporate giants that dominated their respective industries.[1] As Figure 1 shows, this period also witnessed a remarkable rise in breakthrough innovation, especially in fields like chemistry and electronics (Mowery and Rosenberg 1998; Field 2003; Gordon 2016). Whether the two phenomena are causally connected, however, has remained an open question, owing to data limitations and identification challenges.

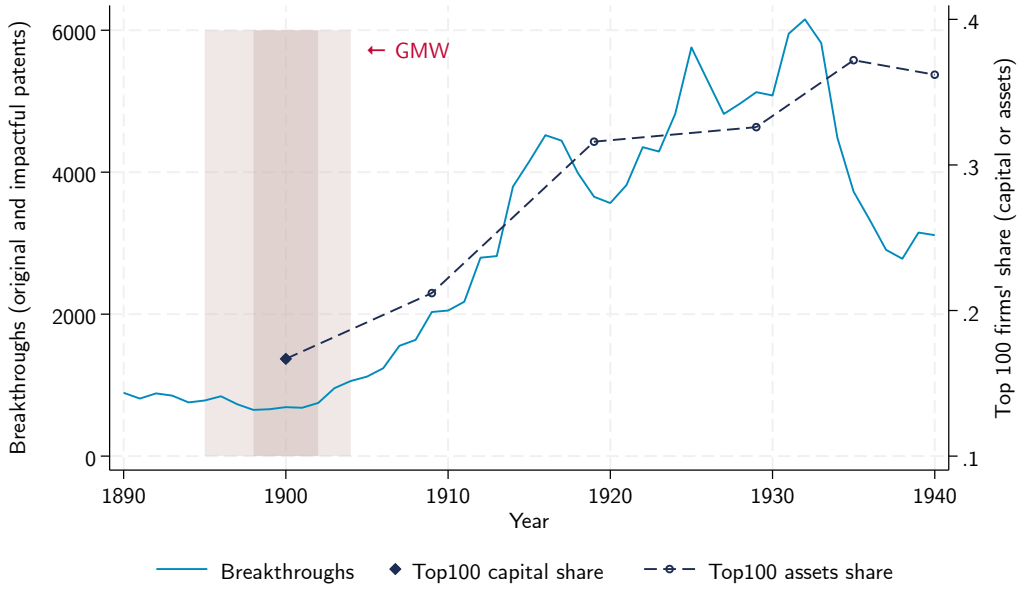
This paper overcomes these issues by using the GMW as a quasi-experiment for studying the innovation effects of consolidation. A central identification concern is that consolidations might have been selectively organized around firms with greater innovation potential. Instead, GMW activity was driven by economic pressures and legal incentives unrelated to firms' innovative prospects. The deflationary Depression of 1893–1897 triggered severe price competition and overcapacity, as wholesale prices fell over 15 percent. Mergers concentrated in industries with low profit margins and high fixed costs (Lamoreaux 1985) and involved a large number of firms.[2] Crucially, the legal environment made consolidation necessary for firms seeking to restrict output. Court rulings beginning in 1895 clarified that while the Sherman Act prohibited cartels and price agreements, mergers remained legal (Bittlingmayer 1985). Firms had no alternative but to consolidate if they wanted to control supply and stabilize prices (Stigler 1950).[3] Corporate R&D was only nascent during this era, and the literature finds little evidence that technological synergies motivated mergers. Empirically, this is reflected in flat pre-consolidation innovation trends at levels of both GMW firms and affected tech-

1. The estimate is from Golbe and White (1988), adjusted for inflation. The second-largest M&A peak was recorded in 2000 and valued at about US$(2024) 1.8 trillion (Huang, Yang, and Zhao 2025). Examples of GMW consolidations include U.S. Steel (the largest U.S. corporation at the time and the first billion-dollar company), DuPont, and International Harvester.

2. The average consolidation involves 9.8 firms (Table 1). Recent evidence also highlights the role of tariffs: industries more exposed to arbitrary tariff hikes saw greater merger activity (Ahumada 2025).

3. Four in five consolidations combined direct horizontal competitors only, with constituent firms overwhelmingly active in the same technological area (Table 1).

Figure 1: Innovation and the Rise of Big Business, 1880–1940



Note: The figure shows the relationship between the size of top firms and breakthrough innovation in the U.S. from 1890 to 1940. The solid line represents the annual count of breakthrough innovations, which are defined following Kelly et al. (2021) as patents that are both novel relative to predecessors and influential for subsequent inventions (top decile of a text-based importance measure). The dashed line and diamond marker show the share of total manufacturing assets and capital, respectively, controlled by the 100 largest manufacturing firms. Concentration data comes from Kwon, Ma, and Zimmermann (2024) and the 1900 U.S. Census of Manufacturing.

nological sectors (Figures 7 and 17) and limited selection on firm observables (Appendix Table C3).

Exploiting this quasi-experimental variation, I employ three complementary empirical strategies to investigate how large-scale consolidation affected U.S. innovation. First, I estimate firm-level effects, presenting the evidence against three comparisons: consolidating firms' own pre-merger trajectories, consolidations that were announced but ultimately abandoned—"paper" mergers, in the language of contemporaries—and non-merging firms in the same broad technological area. Consolidation led to large innovation surges: among firms already patenting before the wave, patenting rose approximately four-fold relative to its pre-merger level, an increase not shared by matched non-merging firms or by abandoned deals. Breakthrough discoveries, defined as patents that are both highly novel and highly influential (Kelly et al. 2021), rose six-fold. Second, to examine the organizational mechanism behind these gains, I construct new inventor–firm linked data. In an Abowd, Kramarz, and Margolis (1999) framework, I show that corporate R&D laboratories—adopted at much higher rates by consolidating firms—conferred genuine productivity advantages, rather than simply attracting superior inventors. Third, to rule out that firm-level gains merely reflect reallocation or crowd-out, I trace the GMW's effects across entire technological domains, asking

whether consolidation's gains spilled over to other innovators or displaced them. Consolidation accelerated breakthroughs in fields closer to the scientific frontier like chemistry and electronics, but slowed them farther away from it. The net effect saw breakthroughs increase by 13 percent between 1905 and 1940.

Throughout, the treatment is consolidation itself, and the implicit counterfactual is less merger and acquisition activity under stricter antitrust laws. Consolidation made firms simultaneously larger and more dominant—a joint increase that I label "bigness" when describing outcomes. These dimensions combine advantages of scale, scope, and market power that are inherently intertwined and difficult to regulate separately. Because the GMW increased them together, the estimated effects capture their compound effect—the full bundle of advantages associated with Big Business in contemporary and historical debates—rather than the contribution of any single dimension. Crucially, the analysis identifies the establishment of corporate R&D laboratories as a key mechanism driving innovation gains by big firms. Because the treatment bundles scale, scope, and market power, I do not separate their individual contributions; what the evidence does point to is the organization of corporate research—the industrial laboratory—as a central channel through which consolidation raised innovation, rather than product-market power acting on incentives alone.[4]

To conduct this analysis, I construct a new dataset linking corporate structure, innovation outcomes, and individual inventors from 1875 to 1955. This contribution has three main components. First, I digitize handwritten worksheets compiled by Nelson (1959), which document mergers and acquisitions from 1895 to 1930, and National Research Council surveys on corporate R&D laboratories (1920–1946). Second, drawing on both existing sources and original data collection, I track over 23,000 ownership changes and identify roughly 13,000 distinct enterprises with information on M&A activity, R&D labs, and subsidiaries. To match these firms to innovation outcomes, I also reconcile over 137,000 unique firm assignees in the patent record. Third, I disambiguate one million inventors, matching them to firms and corporate R&D labs. The resulting inventor-firm panel spans eight decades and, to my knowledge, is the first to extend before 1940.[5] In addition, I assemble two ancillary datasets from the contemporary financial press: a control group of 30 paper mergers—consolidations announced between 1891 and 1907 but ultimately abandoned—whose constituent firms I link to the patent record, and the earliest public announcement date of every consolidation and abandoned deal.

Using this dataset, I identify the firm-level effects of consolidation on innovation through a difference-in-differences framework comparing firms that merged during the GMW to select non-merging enterprises. The analysis distinguishes between two margins of innovative activity. On the intensive margin, I examine consolidated firms with at least one patent before 1895. I construct pre-merger innovation outcomes for GMW firms as the sum of individual constituent firms' outcomes. Then, I build the control group

4. Other dimensions sometimes associated with bigness include diversification, access to finance, and managerial practices. Where feasible, I examine some of these additional channels. For instance, restricting to purely horizontal consolidations (Appendix Figure C6) yields similar estimates, suggesting vertical integration and diversification are not primary drivers.

5. The closest effort is by Akcigit et al. (2022), whose panel begins in 1940.

from medium-sized non-merging firms attested in the patent record.[6] On the extensive margin, I analyze consolidated firms with no prior patenting activity, comparing them to manufacturing enterprises of similar size listed in the 1900 *Moody's Manual of Industrial Securities* that neither patented before 1895 nor participated in the merger wave.

My first set of results shows that consolidated firms significantly increase their innovative activity. Firms with pre-merger patents (intensive margin) experience sustained innovation gains: patenting rises by approximately 6 patents annually, while breakthrough innovations increase by 0.56 per year. Though these absolute increases may seem modest, they reflect the generally low level of patenting by U.S. firms prior to 1900, amounting to a 310 and 536 percent increase relative to pre-merger levels respectively. Merging firms with no prior patenting activity (extensive margin) have a 23 percentage point higher likelihood of starting to patent than control firms. Paper mergers sharpen the causal interpretation of these gains: patenting by abandoned deals stays flat around its pre-announcement level, and completed consolidations add, on average, about 2.2 patents per deal-year relative to them. The lower magnitude reflects a pooled sample with a lower baseline rather than a weaker effect.

Corporate R&D laboratories emerge as the primary mechanism behind substantial firm-level innovation gains from consolidation. Consolidating firms substantially increased their adoption of dedicated research infrastructure: firms on the intensive margin of patenting experience a 16 percentage point increase in the probability of establishing at least one laboratory, while the likelihood rises by 4.5 percentage points for extensive margin firms. Because patent assignment may reflect a variety of relationships between firm and inventor, I use the geographic proximity of inventors to laboratory facilities to assess whether labs were substantively involved in innovation, building on Nicholas (2009). Decomposing the sources of breakthrough innovation gains for intensive margin firms reveals that approximately 50 percent comes from patents filed by inventors within 50km (30mi) of an R&D lab—likely representing in-house research—while an additional 30 percent comes from patents assigned to lab-owning firms but filed by inventors beyond this radius, consistent with labs not only generating innovations directly but also playing a crucial role in evaluating and acquiring external inventions.

Yet a fundamental question remains: Do laboratories actually enhance firm-level innovative productivity, or do they merely attract superior talent and correlate with other firm characteristics? To investigate this, I extend the analysis beyond GMW firms, building a matched inventor–firm panel spanning 1875–1950. I first measure whether an observational productivity premium exists, and find that lab-owning firms are more innovative on average. Then, I estimate an Abowd, Kramarz, and Margolis (1999)—AKM hereafter—decomposition that exploits inventor mobility to separate inventor ability from firm-specific innovation productivity.[7] This approach allows me to explore whether

6. To address mechanical size differences, I reweight control firms to obtain a comparable distribution of pre-1895 patent counts. Alternative inference approaches directly aggregate control firms rather than reweighting them. I implement synthetic control methods (constructing optimal weighted combinations of control firms for each treatment unit), synthetic difference-in-differences (Arkhangelsky et al. 2021), and placebo mergers (randomly selecting control firm combinations matching treatment units' pre-period innovation). These approaches yield comparable estimates (Appendix D).

7. Bhaskarabhatla et al. (2021) perform a similar decomposition using contemporary data but do not

the observed premium reflects better inventors, more innovative firms choosing to establish labs, or genuine productivity gains from laboratory organization.

My second set of results suggests that R&D labs provide substantial innovative productivity advantages. Firm-specific effects account for approximately 33 percent of total explained variation in innovative productivity, with this share rising after 1905 as corporate R&D labs became more common. Lab-owning firms exhibit significantly higher firm productivity, but they do not employ systematically more productive inventors, rejecting the hypothesis that labs operated as a talent magnet. The laboratory premium exists even when controlling for firm size and technological specialization, and firms experience significant productivity gains following laboratory establishment. Likewise, individual inventors become significantly more productive upon joining lab-owning enterprises, raising their quality-weighted output by about 0.12 log points. Together, these results indicate that laboratories represented a fundamental organizational innovation, transforming R&D from individual invention to systematic, collaborative research processes that underpinned big firms' outsized innovative output.

Firm-level innovation gains, however, may not translate into broader technological progress if they merely reallocate inventive activity or crowd-out innovation by other actors. I therefore shift focus from individual firms to total innovation within technological domains, examining whether exposure to consolidations spilled over to—or displaced—innovation by other actors. I distinguish between established technologies (those with at least one patent by 1895) and emerging domains (those with no pre-1895 activity), where a technology is defined as a group of related patent classes. For established technologies, I define consolidation exposure based on whether GMW firms held any patents in that domain before 1895. I then compare technology by exposure status in a standard difference-in-differences approach. For technological domains emerging after 1895, I exploit the hierarchical structure of patent classification: emerging technologies inherit exposure from their parent subclass if related pre-existing domains had GMW firm activity before 1895. Using survival analysis, I model the timing of each domain's first patent and breakthrough, exploiting variation in GMW exposure within patent classes.

My third set of results indicates that consolidation raised breakthroughs by 13.2 percent overall (1905–1940). Yet this average effect masks stark differences once technologies are classified by closeness to the scientific frontier. Following established historical assessments of R&D intensity and reliance on scientific knowledge (Mowery and Rosenberg 1998; Chandler 1990; Arora et al. 2024), I distinguish science-based fields (chemistry, electronics, telecommunications) from non-science-based domains.[8] In established, science-based fields, exposure generates large increases in breakthrough patenting that persist even when excluding GMW firms, indicating positive spillovers. By contrast, in emerging, non-science-based fields, exposure significantly reduces the likelihood of

focus on the role of R&D laboratories.

8. I group patent CPC sections into three categories based on these assessments. The science-based technology group encompasses chemistry, metallurgy, scientific instruments, computing, electronics, and telecommunications (sections C, G and H). The engineering and industrial group spans mechanical engineering, manufacturing processes, vehicles, weapons, heating systems and cross-cutting technologies (sections F, B and Y). The infrastructure and consumer-oriented group includes agriculture, food processing, medical devices, construction, textiles and apparel (sections A, E and D).

achieving a first breakthrough. Overall, back-of-the-envelope calculations suggest that the GMW increased breakthroughs in science-based technologies by 30 percent, while reducing them in non-science-based domains by nearly 7 percent.

Related literature and contribution. This paper contributes to several strands of literature. First, it advances the economic history of American innovation. While prior scholarship has carefully examined the determinants of the Great Merger Wave—highlighting the roles of price competition, tariff policy, and legal incentives in driving consolidation (Lamoreaux 1985; Bittlingmayer 1985)—its consequences remain quantitatively unstudied. Despite influential historical narratives linking Big Business to America's technological ascendancy (Chandler 1977; Gordon 2016; DeLong 2022) and detailed case studies of notable firms (Hounshell and Smith 1988; Jenkins 1975; Wise 1985), no work has causally identified how GMW consolidation shaped innovation or indeed any major economic outcome. This paper provides the first quantitative causal evidence that GMW consolidation drove technological progress, offering concrete support for prevailing narratives while challenging skeptical views that question whether large and dominant corporations genuinely advanced American innovation (Noble 1979; Lamoreaux 2000). In so doing, it complements recent quantitative studies of other pivotal moments in U.S. innovation history, such as the wartime research mobilization of the 1940s (Gross and Sampat 2023) and the Apollo program (Kantor and Whalley 2023).

This contribution sits within the wider economic history of a period of exceptional technological dynamism (Field 2003; 2013; Gordon 2016) and organizational transformation (Chandler 1959; 1977; 1990). This era witnessed dramatic changes in market concentration (Kwon, Ma, and Zimmermann 2024), firm size (Collins and Preston 1961; Navin 1970), and the geography of innovation (Andrews and Whalley 2022). Prior research has analyzed how historical shocks shaped innovation outcomes—World War I (Moser and Voena 2012), the Great Depression (Nanda and Nicholas 2014; Babina, Bernstein, and Mezzanotti 2023; Lampe and Moser 2016), the Bell breakup (Watzinger et al. 2020; Watzinger and Schnitzer 2022), and migration flows (Moser, Voena, and Waldinger 2014; Moser and San 2020)—as well as changes in inventor characteristics and mobility (Nicholas 2010; Akcigit, Grigsby, and Nicholas 2017). By comparison, this paper provides the first economy-wide quantitative assessment linking the GMW's corporate consolidation to long-term innovation outcomes across U.S. manufacturing.

A second contribution is to the literature on corporate R&D laboratories as organizational drivers of innovation. How firms structure their innovation activities—whether through internal development, external sourcing, or hybrid arrangements—has profound implications for their innovative capacity (Arora, Fosfuri, and Gambardella 2001). Prior work has explored the internal structure and management of firm R&D (Henderson and Clark 1990; Cohen and Levinthal 1990; Henderson and Cockburn 1996; Argyres and Silverman 2004; Arora, Belenzon, and Rios 2011), spatial dynamics of innovation (Bikard and Marx 2020), and scientist incentives within firms (Sauermann and Cohen 2010). Scholars have also documented the decline of corporate science in recent decades (Arora, Belenzon, and Patacconi 2018). In economic history, the emergence of industrial research is richly documented—largely through observational evidence (Mowery 1983;

1984; Nicholas 2003; 2009; 2011)—with debate continuing over lab-formation drivers (MacGarvie and Furman 2005; Arora et al. 2024) and productivity effects (Jewkes, Sawers, and Stillerman 1958; Nicholas 2009; Hartog et al. 2024).

This paper demonstrates a causal link from bigness to R&D laboratory establishment. It also shows that the widely observed lab premium is not driven by the sorting of superior inventors into lab-owning firms, and is consistent with a genuine firm-level productivity effect of laboratory organization. These findings provide robust quantitative support for historical narratives portraying corporate labs as engines of American innovation (Hughes 2004; Gertner 2013; DeLong 2022), in contrast to scholars who have questioned whether laboratories genuinely drove technological progress (Reich 1985; Lamoreaux, Sokoloff, and Sutthiphisal 2011; Gruber and Johnson 2019). This contribution rests on the first application of an AKM model to historical inventor productivity: while AKM models have been widely used to study labor earnings and, more recently, patenting outcomes in post-1975 data (Bhaskarabhatla et al. 2021), no prior work has applied them to historical innovation settings.

A third contribution is to the literature on firm size, market concentration, and innovation. Modern growth theory recognizes that Schumpeterian and Arrovian forces operate simultaneously, generating a generally non-monotonic relationship between product-market competition and innovation. Aghion et al. (2005) show both theoretically and empirically that competition and innovation can follow an inverted-U pattern, with innovation rising and then declining as competition intensifies. Empirical studies provide extensive correlational evidence linking large and dominant firms to innovation (Atkinson and Lind 2019; Braguinsky et al. 2023), and antitrust debates often hinge on whether breaking up incumbents will spur technological progress (Federico, Morton, and Shapiro 2020; Shapiro 2019). Yet the evidence remains mixed (Cohen 2010; Bryan and Williams 2021). Even in similar settings, scholars reach divergent conclusions about startup acquisitions (Phillips and Zhdanov 2013; Cunningham, Ederer, and Ma 2021; Fons-Rosen, Roldan-Blanco, and Schmitz 2021), horizontal competition (Blonigen and Pierce 2016; Haucap, Rasch, and Stiebale 2019; Kang 2023; Comanor and Scherer 2013), or oligopoly models (Goettler and Gordon 2011; Igami 2017; Igami and Uetake 2020). R&D itself can be defensive—patenting to deter entry rather than to advance the frontier (Gilbert and Newbery 1982; Argente et al. 2020)—though spillovers often outweigh business-stealing motives (Griliches 1992; Bloom, Schankerman, and Van Reenen 2013; Jones and Summers 2020).

This paper provides new causal evidence on the innovation effects of consolidation—and, with it, of the joint increase in firm size and market dominance—during a transformational historical episode. While I cannot directly test specific theoretical models of market concentration and innovation—as precisely measuring industry-level concentration in this setting remains challenging—the results offer important insights for the literature. Section 7 interprets my findings through the lens of economic theory to assess which mechanisms they most plausibly reflect.

Finally, the findings also carry implications for current debates about concentration and innovation. Recent work documents defensive innovation strategies by modern incumbents that dampen creative destruction (Akcigit and Ates 2023; Fernández-Villaverde, Yu, and Zanetti 2025). The positive effects documented here—particularly

in science-based frontier domains—likely reflect a historical context in which the U.S. lacked substantial federal funding for research (Gruber and Johnson 2019) and American universities often lagged their European peers in key fields (Graham and Diamond 1997). In this environment, Big Business was arguably the only institution with both the resources and incentives to sustain long-term R&D. While extrapolation must be cautious, the results strongly suggest that the net effects of big firms on innovation are highly heterogeneous and dependent on the knowledge base of affected technologies.

Outline. The paper proceeds as follows. Section 2 describes the historical background. Section 3 describes the data. Section 4 analyzes the GMW's firm-level effects. Section 5 investigates the R&D lab mechanism. Section 6 examines the GMW's spillover effects on U.S. innovation. Section 7 discusses the findings, and Section 8 concludes.

2 Historical Background

2.1 The Rise of Big Business and Industrial Consolidation

The late 19th century American economy had developed crucial preconditions for large-scale industrial consolidation. National market integration via transcontinental railroads and telegraph networks created unprecedented opportunities for economies of scale (Donaldson and Hornbeck 2016). Institutional innovations—including general incorporation laws (Langlois 2023) and deepening capital markets (Baskin and Miranti 1997)—removed previous barriers to nationwide industrial enterprises.

The Sherman Act of 1890, designed to limit anti-competitive behavior, arguably accelerated industrial consolidation through judicial interpretations that incentivized mergers (Bittlingmayer 1985). Before the Act, American industrialists had repeatedly attempted to suppress what they viewed as destructive price competition through informal collusive agreements, but these consistently failed due to strong incentives for individual firms to deviate from the arrangement (Ellison 1994). The 1895 *E. C. Knight* case (concerning the Sugar Trust) first tested the Sherman Act and held that manufacturing consolidations lay beyond its reach, on the ground that manufacturing was not interstate commerce. The illegality of cartels and price agreements among independent firms was established instead by the later *Trans-Missouri* (1897) and *Addyston* (1898) decisions. Taken together, these rulings directed antitrust enforcement at cartels rather than mergers, closing off loose combination as an alternative to consolidation and leaving a permissive environment for merger activity. In February 1898, an editorial in the trade publication *Iron Age* commenting on *Addyston* made explicit that the legal environment was incentivizing mergers:

> "The new decision is one which may gravely affect some of the arrangements now in force among manufacturers in different lines, in which some control over prices is sought by concerns otherwise acting independently in the conduct of their business. At first sight it looks as though *this decision must drive*

them to actual consolidation, which is really more apt to be prejudicial to public interests than the losses and temporary agreements which it condemns."[9]

At the same time, the Depression of 1893–1897 led to substantial deflation and triggered severe price competition, making consolidation an increasingly attractive option to restrict supply under the new legal environment. Wholesale prices fell by more than 15 percent during the Depression before rebounding as consolidation activity accelerated (Appendix Figure C1). In a seminal study, Lamoreaux (1985) demonstrates that consolidation was more likely in industries with lower profit margins and higher fixed costs, consistent with mergers being a defensive response to price wars in more exposed sectors of the economy.[10] Wall Street financiers played an active role in promoting consolidations. Figures like J.P. Morgan and John W. Gates extracted substantial promotional profits for organizing mergers, reaching up to 20 percent of the new firm's capitalization (Du Boff and Herman 1989; Markham 1955).

The Great Merger Wave of 1895–1904 caused a structural transformation of American industry. Though intensity varied, a degree of M&A activity occurred in most manufacturing sectors (Nelson 1959). Unlike later merger waves characterized by diversification or vertical integration, the Great Merger Wave was overwhelmingly horizontal—combining direct competitors into dominant market players. Of the 93 consolidations with traceable market shares, 72 controlled at least 40 percent of their respective markets and 42 controlled at least 70 percent (Lamoreaux 1985).[11] The wave was not only the largest M&A event in U.S. economic history, but also exceptional by international standards. The UK and Germany also saw outbursts of M&A activity between 1880 and World War I, but on a considerably smaller scale (Tilly 1982; Hannah 1974; Kling 2006).

Figure 2 shows consolidation activity in the U.S. between 1880 and 1920, revealing the sharp timing and unparalleled scale of the Great Merger Wave.

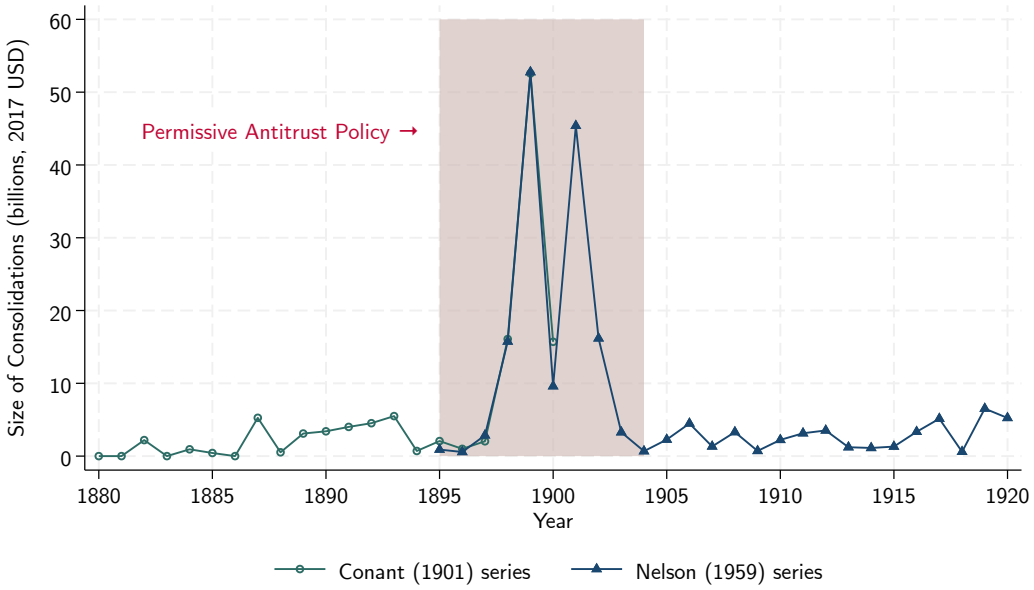
Despite the varied performance of individual consolidations, the Great Merger Wave significantly and enduringly reduced competition and increased industrial concentration (Porter 2006). Many mergers that aimed for near-total market control struggled to maintain it: as Lamoreaux (1985) documents, aggressive post-merger price increases often triggered entry by new competitors, leading to a partial erosion of market share. By 1932, Livermore (1935) had classified about one-third of these consolidations as "failures" (Lamoreaux 2000). This designation, however, did not necessarily imply collapse or exit, but rather persistent underperformance on benchmarks such as earnings, stock prices, or the need for repeated recapitalizations and restructurings. Yet these setbacks did not bring a return to pre-merger levels of competition. Even where monopolistic dominance was short-lived, the industrial landscape remained fundamentally altered: in sector after sector, consolidation gave rise to durable oligopolistic structures, with the

9. As reported in Bittlingmayer (1985), emphasis added.

10. Appendix C.1 discuss the industry-level evidence by Lamoreaux (1985). Ahumada (2025) complements it by showing that arbitrary tariffs contributed to more intensive consolidation activity, consistent with a model in which higher tariffs incentivize mergers when import prices constrain domestic ones.

11. The lower number of consolidations covered in Lamoreaux (1985) reflects both double counting in the Nelson figures (e.g. large consolidations like U.S. Steel took several progressively larger mergers to achieve their final form) and data being unavailable for some mergers (likely smaller or less successful).

Figure 2: Consolidation Activity and the Great Merger Wave, 1880–1920



Note: The figure plots the value of industrial consolidations in the United States between 1880 and 1920 (in billions of 2017 USD). The blue triangle and green circle series show consolidation activity from different historical sources: Nelson (1959) and Conant (1901), respectively. Vertical lines mark key events that shaped the merger wave. The sharp spike in merger activity between 1895 and 1904 defines the Great Merger Wave.

merged firm often retaining a leading—if no longer hegemonic—position (Lamoreaux 1985; 2000).[12]

The permissive merger environment ended abruptly as public concern over monopoly power fueled the Populist and Progressive movements. Beginning in 1902, the Roosevelt administration launched a series of landmark antitrust suits against major consolidations, marking a sharp reversal in federal enforcement policy. The Supreme Court's 1904 *Northern Securities* decision—ordering the dissolution of a railroad holding company created by J.P. Morgan—signaled the close of the era of unregulated consolidation and effectively brought the GMW to an end. This new stance was reinforced by the 1911 breakups of Standard Oil and American Tobacco. Seeking to restrain Big Business and address fears of economic exclusion and political manipulation, in 1914 Congress enacted the Clayton Act to prohibit mergers that would "substantially lessen competition, or tend to create a monopoly" (Lamoreaux 2019). It also established the Federal Trade Commission to enforce these provisions.

12. Indeed, Ma et al. (2025) find that today's largest U.S. firms disproportionately originate from the 1880–1920 cohort, likely a reflection of the GMW's long-run imprint on American Big Business.

2.2 A Golden Age of American Innovation

Throughout the 19th century, U.S. innovation was largely the purview of individual independent inventors. Patents served as a key mechanism for monetizing inventions thanks to a robust, and increasingly national, market for patent rights (Lamoreaux and Sokoloff 1999; 2001; Lamoreaux, Sokoloff, and Sutthiphisal 2011). Notable patent-protected inventions like the 1840 telegraph (Samuel A. Morse), the 1851 sewing machine (Isaac M. Singer) and the 1876 telephone (Alexander G. Bell) formed the basis of new and very successful companies dedicated to their commercialization.

From patents begetting firms, the U.S. gradually shifted to an innovation system where firms beget patents (Nicholas 2010). In 1880, independent inventors accounted for about 77 percent of patents and firms less than 10 percent. By 1940, firms were assigned 57 percent of patents and inventors 37 percent. The pattern is even starker for breakthrough innovations: between 1880 and 1940 the firm share went from 19 to 78 percent.[13] The Great Depression contributed to the decline of independent invention, as younger inventors increasingly joined corporate research departments (Babina, Bernstein, and Mezzanotti 2023).

Firms took on a more active role in organizing the inventive process, often through the establishment of dedicated research facilities and the employment of educated personnel (Mowery and Rosenberg 1989; Mowery 1990). While pioneering examples like Thomas Edison's New Jersey laboratory date back to the 1880s, the R&D lab as a specialized organizational unit devoted to systematic experimentation and discovery became more common around and after 1900.[14] By the 1920s, scientific methods and professional management had become prevalent (Arora et al. 2024). Data from the National Research Council show that in 1940 about 2,300 firms owning R&D labs were employing about 71,800 personnel, with a large majority having scientific or technical education. Bell Labs, for example, employed over 3,600 staff by 1940, including hundreds of Ph.D. scientists working on fundamental and applied research problems (Gertner 2013).

In a context of nearly absent federal support for R&D, industry's relationship with academia combined collaboration and competition. Firms relied on universities for trained engineers and scientists, but also built internal research infrastructure to compensate for weak public science (Arora et al. 2024). Some firms recruited prominent academic scientists into industry by offering superior equipment, research freedom, and generous salaries (Hounshell and Smith 1988; Wise 1985).[15] Others formalized partner-

13. My calculations based on data from CUSP (Berkes 2018) and Kelly et al. (2021).

14. The functions of early R&D labs are clearly illustrated by the 1902 establishment of DuPont's Eastern Laboratory. Hounshell and Smith (1988, p. 19) quotes the new laboratory's, director, Charles L. Reese, as he outlined six objectives for the new institution:

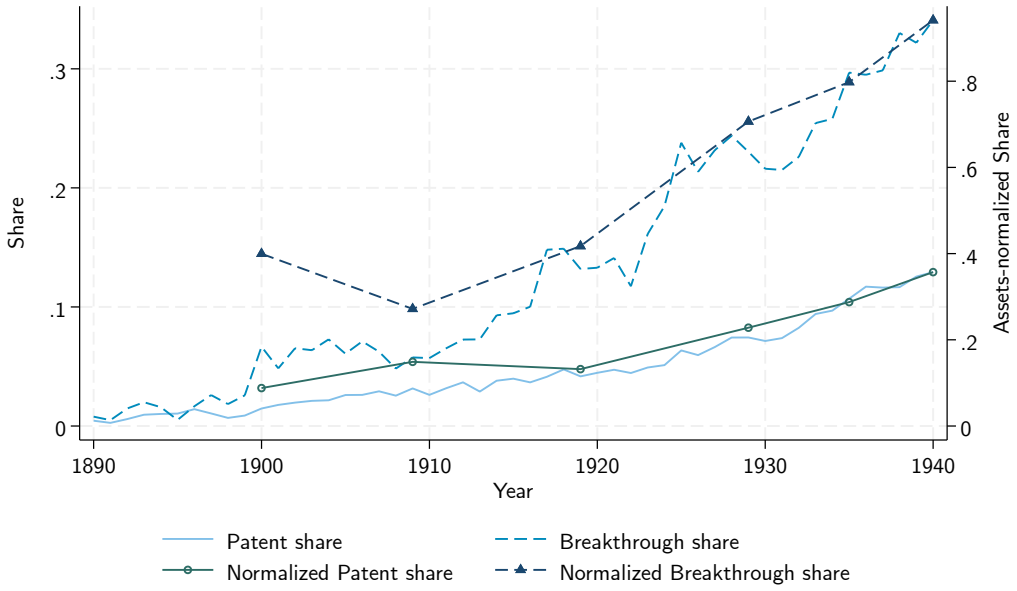
> (1) To improve as far as possible the chemical operations now employed. (2) To investigate the explosives now being manufactured, to revise their Formulas and to put them on a scientific basis. (3) To devise or discover new explosives for general and specific purposes. (4) To keep in touch with all new and improved processes which have any bearing on operations connected with the explosives industry. (5) To investigate new explosives brought forward by outsiders or suggested by members of the Company. (6) To train young chemists and keep the plants supplied with technical, assistants.

15. Prominent examples include Irving Langmuir at GE and Wallace Carothers at DuPont.

ships with universities through fellowships, contract research, and advisory boards (Furman and MacGarvie 2007). These ties became significant sources of funding for some institutions: by 1919, industry gifts to the University of Michigan exceeded government research funds by a factor of three (University of Michigan 1919).

The era of the R&D lab witnessed remarkable technological progress, with breakthrough innovation reaching historical peaks in the 1920s and the early 1930s (see Figure 1). Among the significant innovations emerging from corporate labs were synthetic materials (nylon and neoprene from DuPont), communications technology (radio and early television at Radio Corporation of America), and electrical innovations (household appliances and power systems from General Electric) (Rosenberg 1994). Field (2011) argues that the 1930s were particularly productive despite the Depression, with technological advances in manufacturing, transportation, and communications laying groundwork for subsequent economic growth.

Figure 3: Share of Patents and Breakthrough Innovations from Large Firms, 1890-1940



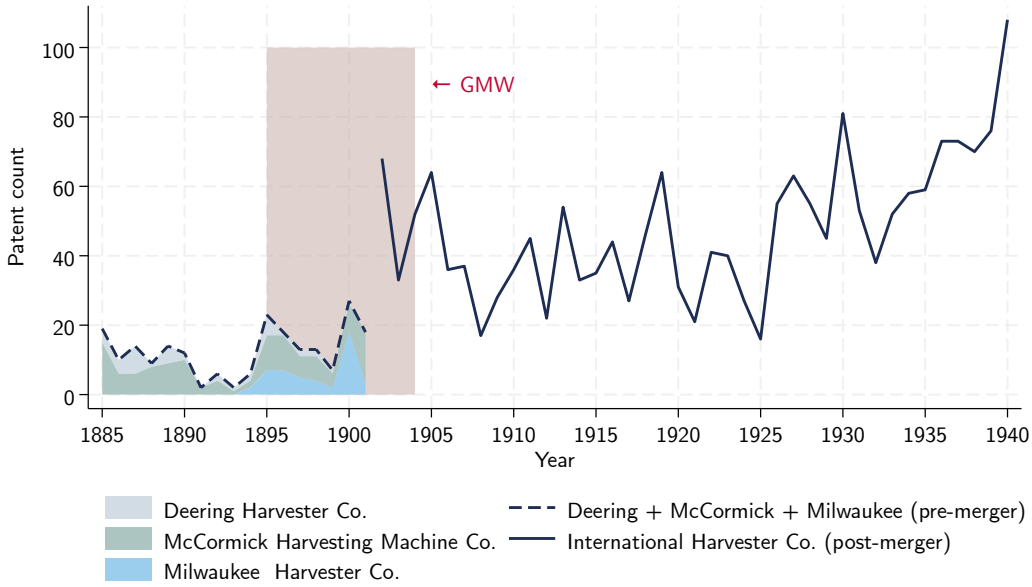
Note: The figure shows the increasing dominance of large firms in U.S. innovative activity between 1890 and 1940. The solid line shows the share of all U.S. patents assigned to large firms; the dashed line reports their share of breakthrough innovations. The trend persists when normalizing patent shares by asset shares (from Figure 1). Normalized series are represented by circles (patents) and triangles (breakthroughs). Large firms are identified using the list in Collins and Preston (1961).

Between 1890 and 1940, large firms grew to account for a large and increasingly disproportionate share of innovative activity (Figure 3). Companies like DuPont, Eastman-Kodak and GE came to dominate their respective areas of technological specialization. Yet scholars debate whether this concentration of innovative activity served broader technological progress (Nicholas 2003; Lamoreaux 2000; Mowery 1990). In particular, work by Reich (1977, 1980, 1985) has highlighted R&D and patenting practices by dominant firms like GE and Bell that seemed—often overtly—directed at preempting competitors

and erecting barriers to entry, rather than commercializing useful new or improved products. Consistently with this view, Watzinger et al. (2020) and Watzinger and Schnitzer (2022) show positive innovation effects from the demise of the Bell System. By contrast, Chandler (1977) and DeLong (2022) argue that large firms developed crucial organizational infrastructure that greatly benefited development and discovery. Positive accounts of the achievements of corporate R&D in specific enterprises are not lacking, including DuPont (Hounshell and Smith 1988), Bell Labs (Gertner 2013), Eastman-Kodak (Jenkins 1975) and General Electric (Wise 1985).

2.3 A Look at the International Harvester Company

Figure 4: Patenting Activity of International Harvester and its Predecessors, 1885-1940



Note: The figure tracks patenting activity before and after the 1902 formation of International Harvester through the merger of McCormick Harvesting Machine Co., Deering Harvester Co., and three smaller manufacturers—Milwaukee Harvesting Machine Co., Plano Manufacturing Co., the Warder, Bushnell, and Glessner Co. The pre-merger total (dashed line) shows the combined patenting output of the merging firms, namely McCormick, Deering and Milwaukee, as the other two firms did not patent.

The International Harvester Company provides a useful case study. Formed in 1902 through the merger of fierce competitors McCormick and Deering Harvester Companies plus three smaller manufacturers, it embodied the classic pattern of the GMW: horizontal consolidation facilitated by Wall Street financiers, clearly motivated by a desire to end price competition (Kramer 1964).[16] The merger achieved its immediate goal—

16. According to Kramer (1964, p. 301)'s rich analysis of historical records and correspondence, the "overwhelming weight of evidence points to elimination of competition in order to control output and prices, and thus ultimately to increase profits as the most important motive of the merger." The merger's

International Harvester commanded approximately 90 percent of total domestic production of grain binders and about 80 percent of mowers, the two major types of harvesting machines. An investigation by the Bureau of Corporations found substantial price increases between 1903 and 1911 and much greater profitability in product lines that were more monopolized (Kramer 1964).

Figure 4 tracks the innovative activity of International Harvester and its predecessor firms over time. For a pre-merger benchmark, I use the combined patenting of the merging firms. Prior to consolidation, the constituent firms collectively averaged 12.8 patents annually. Post-merger, this figure more than tripled to 48.1 patents per year. Breakthrough innovations also increased markedly, rising from an average of 0.18 to 1.49 annually. By 1940, International Harvester operated R&D facilities in eight locations employing hundreds of specialized scientific personnel. Among its major contributions, International Harvester introduced the first general-purpose tractor, the Farmall (1924), which transformed mechanized agriculture by replacing single-use machines with a versatile, mass-produced platform for plowing, planting, and cultivation.

The case of International Harvester illustrates how industrial consolidation during the Great Merger Wave could coincide with substantial increases in innovation. The remainder of the paper investigates whether this pattern generalizes—whether mergers systematically caused firms to become more innovative—and explores the mechanisms and broader consequences of this radical transformation of American industry.

3 Data

To study the relationship between Big Business and innovation, I collect extensive data covering firms, inventors and patents between 1875 and 1955. This effort results in three key contributions. First, a detailed dataset describing M&A activity during the Great Merger Wave and beyond. Second, a firm linkage system that harmonizes identities across disparate sources and tracks evolving ownership, enabling accurate attribution of innovation to the correct business entities. Third, a long-run panel of all U.S. inventors matched to patent and firm information.

3.1 Sources

My analysis draws on several data sources—discussed in greater detail in Appendix A.

Merger activity. I digitize the detailed handwritten worksheets compiled by Ralph Nelson for his seminal study of U.S. merger history (Nelson 1959).[17] Although Nelson

execution was critically facilitated by J.P. Morgan & Co. There had been at least two earlier attempts at consolidation, one of the key roadblocks being disagreements between the McCormick and Deering families as both wanted to control the merged organization. Eventually, George W. Perkins of the Morgan firm structured and executed an agreement, at no small cost for the manufacturers—overall J.P. Morgan & Co. got 4.2 million dollars in commission and promotional fees (about 3.5 percent of capitalization), and outsize decision power (Kramer 1964).

17. I am very grateful to Naomi Lamoreaux for sharing with me the scans of the originals.

originally published only aggregate statistics, his raw data provides granular information on each consolidation and acquisition, including the names of both acquiring and acquired firms—critical for analyzing pre-merger innovation performance. These worksheets identify the type of merger (consolidation versus acquisition), integration strategy (horizontal, vertical, or mixed), and industrial classification (SIC codes). For a subset of cases, Nelson also recorded incorporation details, capitalization values, notes on assets and output. To collect this information Nelson consulted a wide variety of historical sources, and extensively double checked and validated his data.

R&D laboratories. I digitize surveys by the National Research Council conducted between 1920 and 1946. This data contains information on the firm operating the lab and its locations (some firms operated more than one laboratory), together with personnel and other details. The 1940 and 1946 surveys also collected establishment dates for each lab location, so dates are available for only a subset of the recorded locations; I track undated labs separately. Improving on previous efforts, I carefully disaggregate information at the level of each lab location and use their establishment dates to better track lab activity over time. The definition of an R&D lab in these surveys sometimes includes modest facilities like workshops and testing sites, which nonetheless contributed to product development, patent evaluation, and incremental innovation, making them meaningful indicators of corporate inventive capacity.

Patent data. I rely on the CUSP dataset (Berkes 2018), which includes the name of both inventors and assignees, their location, and each patent's technological classification and citations. I complement this information with measures of breakthrough quality developed by Kelly et al. (2021). These are text-based metrics that capture how original (different from previous patents) and impactful (similar to future ones) each patent is. Being in the top 10 percent of this measure, after residualizing cohort effects, defines a patent as a breakthrough. This method has key advantages: (1) because it does not require external information, it can be computed for all patents; (2) it extends to historical patents for which citation data is very limited and unreliable. Kelly et al. (2021) extensively validate their measure, using lists of historical great inventions, modern citation data, and market valuation, among others.

Additional data. For the firm linkage described in the next subsection, I rely on subsidiary data (1926-1950) from Kandel et al. (2019) and on extensive manual collection of new information. I derived additional information from a variety of primary and secondary sources, see Appendix A.

Paper mergers. To assess selection into consolidation, I assemble a control group of consolidations that were announced and negotiated during the GMW era but ultimately abandoned—what Moody (1904) called "paper" mergers. Candidate deals are drawn from contemporary compilations of trusts and from a systematic search for abandonment language over the full run of the *Commercial & Financial Chronicle*, 1890–1908; each deal's fate is then verified across digitized primary sources, and constituent firms are linked to the patent record with the same procedure used for treated firms. The resulting pool comprises 30 abandoned deals, of which 24 carry a reconstructed constituent roster and enter the deal-level analysis. Appendices A.6 and B.3 detail sources, construction, and validation.

Announcement dates. Nelson (1959) dated each merger by the transfer of control

and deliberately discarded earlier press reports of prospective deals. For event-study timing, I recover that discarded layer: the earliest credible public announcement of every treated consolidation and every abandoned deal, collected from the *Chronicle* and allied period sources under a single symmetric protocol, with Nelson's dating retained where no announcement is found. Appendix B describes the collection and its validation.

3.2 Firm Linking

A fundamental challenge for studying historical enterprises is disambiguating and tracking firms over time when working with unstructured data that lacks unique identifiers. Companies frequently change names, use aliases, acquire other businesses, establish subsidiaries, and reorganize—all creating discontinuities in the historical record. To address this challenge, I implement a multi-stage algorithm that identifies and links firms across disparate sources. Appendix B covers in greater details the steps outlined here.

First, I process the patent record's unique 370,000 assignee strings (1840-1960) to distinguish between firms (78 percent of assigned patents) and individuals (22 percent). I do this using a combination of rule-based classification and a machine learning algorithm for named entity recognition. For firms, I standardize names using context-specific cleaning rules and employ locality-sensitive hashing to create computationally efficient comparison blocks. Within each block, deterministic matching rules based on string similarity and temporal proximity resolve assignee identities. At this step, I identify approximately 137,000 unique firm assignees. However, this is not yet a fully complete or reliable measure of a firm's patent record, especially for the largest and most long-lived businesses. The same firm may radically change its name over time or have some patents assigned to a distinct subsidiary that it fully controls, both of which might result in distinct assignees.

Second, I harmonize firm names across non-patent sources, creating an unique identifier and a list of alternative names for each firm. To start, I manually collect a dictionary of aliases, abbreviations and name changes for firms in all my sources. For instance, I record that "American Car & Foundry" was also known as "ACF", a link that would be missed by most reasonable fuzzy matching rules. Next, I standardize name strings in all my non-patent sources and use the hand-collected dictionary to perform both exact and fuzzy pairwise matching between all sources. Afterward, I impose that my pairwise matches be transitive and resolve any conflicts by further manual collection of data.[18] I iterated these steps until there were no conflicts. This process creates a master list of approximately 12,800 disambiguated firms with harmonized identifiers across all sources.

Third, I match assignees to this master list of disambiguated firms using the hand-collected firm name dictionary and fuzzy matching techniques. About one-third of my

18. Concretely, if 'American Car & Foundry' in source 1 is linked with 'ACF' in source 2 and 'ACF' is linked to 'ACF i' (clearly a typo) in source 3, then 'American Car & Foundry' should also be linked to 'ACF i' (linking sources 1 and 3). Imposing transitivity of pairwise links can result in conflicts, where different firms in the one list are linked to the same firm in another list. When this happened, I would review all the firms involved and collect new information that would explain the conflict (e.g. a missing firm name change) or correct any mistakes in the primary and secondary sources used.

12,800 disambiguated firms appear in the patent record. Assignees (classified as firms) that are not matched to my disambiguated firms are assumed to be independent enterprises.[19] I also assign each firm in the analysis samples to one of twelve broad product-market categories inferred from its name, since for most patent assignees no industry information exists outside the patent record itself; Appendix B describes the classification and its validation against the SIC codes and Moody's sections available for a subset of firms.

Finally, I dynamically map ownership by integrating Nelson's M&A data, manually collected information, and subsidiary data from Kandel et al. (2019). The year-by-year ownership resolution algorithm begins with a baseline year (1869) where all firms with no information to the contrary own themselves, then sequentially processes ownership changes through 1960 and propagates them through the ownership chain.[20] At each stage, the algorithm traces ownership chains to identify ultimate owners. As for firm disambiguation, I manually investigated and resolved ownership loops and conflicts, iterating data construction until there were none. Overall, I account for about 23,000 ownership changes. This dynamic mapping allows me to consolidate patent activity at the enterprise level throughout the sample period.

3.3 Inventor Disambiguation

I address the challenge of identifying unique inventors across the patent record by implementing a probabilistic record linkage approach within the Fellegi and Sunter (1969) framework. I build on previous work by Akcigit et al. (2022) by extending this effort to the 1875-1955 period and including richer information. I implement a probabilistic record linkage approach with multiple comparison dimensions (name similarity, technological overlap, geographical proximity, temporal distance, co-authorship and assignee overlap). The algorithm, which incorporates Expectation-Maximization training to generate match probabilities, identifies 1.012 million unique inventors responsible for 2.273 million patents between 1875-1955. The disambiguation process and its validation exercises are described in Appendix B.

From these disambiguated inventors, I construct a longitudinal panel at the inventor-year level. For each inventor, I create balanced yearly observations spanning their first to last patent. Patent information is aggregated annually, capturing counts, breakthrough innovations, fractional contributions (adjusting for co-inventorship), technological diversity, and collaboration patterns. To fill in information gaps in non-patenting years, I assigning characteristics from the nearest patenting year.

To my knowledge, this is the first historical individual-firm panel large enough to allow the estimation of a two-way fixed effect model *à la* Abowd, Kramarz, and Margolis (1999).

19. To the extent that there are radical name changes or ownership changes not reported in my data, this might be over estimating the number of distinct assignees in the patent record.

20. That is, if firm A buys firm B which was recorded to own firm C, firm A also buys firm C. This ensures that repeated consolidations and subsidiaries are correctly carried over.

4 Firm-Level Innovation Effects of the Great Merger Wave

This section estimates the firm-level innovation effects of consolidation. My treatment group is composed of enterprises that have gone through at least one consolidation between 1895 and 1904. In the years before the merger, the outcome (e.g., patents) for the consolidation is the sum of the outcomes for the individual constituent firms that will merge—similarly to the International Harvester example (Figure 4). Additional details on selection and construction of treatment units can be found in Appendix B.

In the short run, a consolidation reorganizes existing firms under unified ownership, replacing the profit-maximizing choices of many competitors with those of one larger and more dominant enterprise. As such, the relevant counterfactual that identifies the effects of this sudden increase in "bigness"—firm size and market concentration—is not a large incumbent firm, but the innovative behavior of the merging firms had they remained separate. A limitation of this approach is that it does not separately identify the effect of size and market dominance, but the compound effect of bigness.[21]

There is no perfect control group for the GMW consolidations so far back in time. I therefore present the firm-level evidence against three complementary comparisons, in increasing order of structure. The first is the consolidating firms' own pre-merger trajectory, which establishes the raw magnitude any explanation must account for. The second is a newly assembled control group of *paper mergers*—consolidations that were announced and negotiated but ultimately abandoned—which holds fixed selection into *attempting* a merger, if not into completing one. The third, and the workhorse for the remainder of the paper, compares consolidations to non-merging firms in the same broad technological area, netting out technology-specific trends. While each comparison has its limits, their agreement is informative.

The comparisons also differ in how they treat time. Where the analysis is conducted in time relative to the deal, I anchor event time at each consolidation's earliest public announcement—newly collected from the contemporary financial press for all treated and abandoned deals (Section 3)—since negotiations, options, and expectations begin at announcement rather than at the transfer of control, and abandoned deals never reach the latter stage at all. In practice, most announcements fall in the calendar year of Nelson's dating, and the results below are robust to using Nelson's dates instead. The matched-firm comparison, by contrast, is estimated in calendar time; the reason, discussed there, is that its sample conditions on pre-1895 patenting outcomes.

Table 1 provides descriptive statistics for the 265 consolidations I reconstructed from the Nelson worksheets.[22] The majority (65 percent) had no patenting activity between 1885 and 1894, thus forming the basis for the extensive margin analysis. Consolidations involved an average of 9.8 firms each, with intensive margin consolidations being notably larger (14.1 firms on average) and more likely to achieve substantial market dominance.

21. See Section 4.5 for suggestive evidence on the role of market concentration in driving innovation effects.

22. Notice that regression analyses use a balanced panel that excludes consolidations which were themselves bought by another firm during the sample period. Unbalanced panel results are entirely similar and reported in Appendix C. There are 33 such consolidations, 15 on the intensive margin and 18 on the extensive margin.

The sample spans all major manufacturing sectors, with the largest representation in food and consumer goods (27 percent) and primary metals (22 percent). Consolidations on the intensive margin were most active in engineering and industrial technologies (50 percent of intensive margin firms).

4.1 Identification

The identifying assumption required for a causal interpretation of my empirical strategy is that, conditional on technology-specific trends, selection into consolidation is not correlated with unobserved determinants of future innovation.[23] Therefore, the key threat to identification is that merging firms were likely to diverge technologically around the GMW because of latent innovation potential or contemporaneous shocks. For example, if mergers were organized in order to control a firm that was about to produce a technological breakthrough; or if concurrent shocks, like financial fluctuations, both correlated with merger decisions and affected innovation.

Five historical and empirical considerations support the validity of the identification assumption.

First, extensive quantitative and historical evidence shows that consolidations targeted price competition, not expectations of future innovation. The 1893-97 depression triggered severe deflation, with wholesale prices falling 15% before recovering as consolidations accelerated (Figure C1), creating strong competitive pressures (Parsons and Ray 1975; Kramer 1964).[24] Lamoreaux (1985) shows that industries with lower profit margins, higher fixed costs, and larger plants were significantly more likely to consolidate, while growth rates had weak predictive power (Table C2). GMW mergers happened to the exceptionally large and sharply timed extent they did (Figure 2) because of the additional pull factors coming from antitrust incentives (Bittlingmayer 1985), Wall Street activism (Du Boff and Herman 1989; Markham 1955) and tariff policy (Ahumada 2025). More generally, there is no historical or empirical evidence that the GMW was driven by attempts to control R&D assets nor potential.

Second, neither individual firms nor broader technological sectors selected into the GMW based on differential innovative trends. Event study results later in Figures 7 and 17 show flat pre-trends both in the firm-level and technology-level analyses. Though data are limited, in Appendix C.1 I show that individual firms that partook in GMW consolidations are not very different from other firms, with goodness of fit of probit specifications explaining participation in the GMW not exceeding 6 percent.

Third, the timing and magnitude of effects rule out mean reversion due to the Depression of 1893. One could worry that the economic downturn both increased merger activity and temporarily depressed firm innovation, making post-merger increases appear larger than they really are due to mean reversion. However, two patterns in the data argue against this interpretation. First, the sample extends back to 1885, providing

23. For the extensive margin, where pre-1895 technology cannot be assigned to firms without prior patents, the analogous assumption conditions on sector-by-year rather than technology-by-year trends, as discussed in the description of the extensive-margin design below.

24. For instance, the consolidation of International Harvester was pursued to end a period of intense competition that industry members dubbed the "Harvester War" (Kramer 1964).

Table 1: Key Characteristics of GMW Consolidations

Sample:	All Firms (1)	Intensive Margin (2)	Extensive Margin (3)
N	265 (100%)	94 (35%)	171 (65%)
Consolidation Characteristics			
Year (median)	1900	1899	1900
Firms per consolidation	9.8	14.1	7.3
Total patents in 1885-1894	6.5	18.4	0.0
Same-technology firms (share)	–	85%	–
Initial market share			
>70%	36 (14%)	22 (23%)	14 (8%)
Between 40% and 70%	29 (11%)	22 (23%)	7 (4%)
<40%	17 (6%)	2 (2%)	15 (9%)
Not available	183 (69%)	48 (51%)	135 (79%)
Integration type			
Horizontal	132 (50%)	56 (60%)	76 (44%)
Vertical	24 (9%)	3 (3%)	21 (12%)
Mixed	10 (4%)	5 (5%)	5 (3%)
Not available	99 (37%)	30 (32%)	69 (40%)
Sector			
Mining & Natural Resources	34 (13%)	3 (3%)	31 (18%)
Food & Consumer goods	71 (27%)	18 (19%)	53 (31%)
Chemicals & Materials	37 (14%)	14 (15%)	23 (13%)
Primary & Fabricated Metals	57 (22%)	25 (27%)	32 (19%)
Machinery & Electrical Equipment	31 (12%)	20 (21%)	11 (6%)
Transportation & Instruments	23 (9%)	12 (13%)	11 (6%)
Not available	12 (5%)	2 (2%)	10 (6%)
Technological area (if ever patent)			
Science-based	47 (18%)	25 (27%)	22 (13%)
Engineering & Industrial	99 (37%)	47 (50%)	52 (30%)
Infrastructure & Consumer	40 (15%)	22 (23%)	18 (11%)
Not available	79 (30%)	0 (0%)	79 (46%)

Note: This table presents descriptive statistics for the GMW consolidations I reconstructed from Nelson's handwritten worksheets. The intensive margin comprises firms with at least one patent between 1885-1894, while the extensive margin includes firms with no activity in that window. Market share data comes from Lamoreaux (1985).

sufficient pre-depression data to detect such patterns, yet pre-trends remain flat with no evidence of a pre-merger innovation decline. Second, the magnitude of post-merger innovation far exceeds any pre-merger peak—averaging more than 300 percent above pre-merger patenting levels—making simple mean reversion implausible as an explanation.

Fourth, one might worry that observed innovation gains stem from technological complementarities between merging firms rather than the act of consolidation itself. This would be a threat to identification if consolidations were selectively formed to combine firms with latent synergies. However, two pieces of evidence suggest otherwise: first, 80 percent of consolidations with available information were horizontal mergers of close competitors, and on average 85 percent of merging firms on the intensive margin were active in the same technological area (Table 1); second, restricting the sample to strictly horizontal consolidations yields similar results (Appendix C.3). Thus, selection on *cross-technology* complementarities is unlikely to explain the findings. The evidence bears less directly on complementarities among close horizontal competitors, who may have shared know-how or engineering teams; such within-industry recombination, however, is best understood as a mechanism through which "bigness"—the joint increase in scale and market power—operates once firms combine, rather than as selection into consolidation. At least at this coarse level, recombination across diverse technologies does not appear to be the main driver of my results.

Fifth, one might worry that consolidations selected firms with superior managerial talent, and that better management—rather than consolidation—drove innovation gains. Yet the evidence makes this explanation unlikely. First, only one-twelfth of major consolidations were organized by industrialists themselves (Nelson 1959), indicating that prior managers played little role.[25] Mergers also fundamentally reorganized management structures, as documented in case studies of International Harvester, U.S. Steel, and more broadly (Kramer 1964; Parsons and Ray 1975; Lamoreaux 1985; Du Boff and Herman 1989; Markham 1955). Such reorganizations created major discontinuities by combining leadership teams, severing the link between (family) ownership and control, and accommodating the demands of creditors and promoters. Second, the timing of effects argues against pre-existing managerial superiority: if better management were already present in merging firms, we would expect to observe differential innovation trends before consolidation. Instead, pre-trends are flat. Third, my estimates substantially exceed the magnitude and persistence of known management effects on innovation and performance (Bertrand and Schoar 2003; Benmelech and Frydman 2015; Acemoglu, Akcigit, and Celik 2022).

To the extent consolidation improved access to capital, enabled managerial reorganization, or facilitated technological recombination, these represent mechanisms through which increased bigness affected innovation, not threats to identification.

25. Two-thirds were promoted by independent financiers and one-quarter by investment banks.

4.2 Consolidated Firms' Own Trajectories

The natural starting point is what happened to the consolidating firms themselves. Figure 5 pools all consolidations—across the intensive and extensive margins of patenting—with no control group. In calendar time, mean patenting per consolidation is low and flat through the early 1890s and rises steadily after the wave (panel a), and breakthrough patents, though far rarer, follow the same path (panel b).

To summarize this shift within firms, I estimate an event study in years relative to each consolidation's earliest public announcement:

$$y_{it} = \alpha_i + \sum_{\tau=-10,\,\tau\neq-1}^{25} \beta_\tau \cdot \mathbb{1}[t - a_i = \tau] + \varepsilon_{it}, \tag{1}$$

where i indexes consolidations, t indexes calendar years, a_i is consolidation i's announcement year, and $\tau = t - a_i$ is the resulting event time; y_{it} is the outcome—patents or breakthrough patents—α_i are consolidation fixed effects, and the coefficients β_τ trace the outcome relative to the year before announcement ($\tau = -1$, the omitted category). Standard errors are clustered at the consolidation level. Because the specification includes only consolidation fixed effects, the β_τ net out each unit's own pre-announcement level but nothing else.

Panels (c) and (d) plot the β_τ. Averaging over the post-announcement years, patenting rises by about 2.0 patents per consolidation-year, roughly tripling the pre-announcement mean of 0.83 patents per consolidation-year; breakthroughs rise by about 0.2, an order of magnitude above their pre-announcement mean of 0.02. Appendix Table C6 reports the corresponding estimates.

This comparison remains deliberately descriptive: the within-firm event study removes each consolidation's own pre-announcement level, but no control group absorbs the secular growth of the patent system or technology-specific trends, so I attach no causal interpretation to it. Its role is to fix the magnitudes at stake. Whatever the counterfactual, it must be weighed against a near tripling of patenting among consolidating firms, and the two comparisons that follow discipline this raw pattern in complementary ways.

4.3 Paper Mergers as a Control Group

A first concern with any comparison group is selection into consolidation itself: firms that merged may have differed, in ways that predicted innovation, from firms that never contemplated merging. To hold this margin fixed, I assemble a control group of consolidations that were announced and negotiated but ultimately abandoned—what Moody (1904) called "paper" mergers. These deals arguably faced the same selection pressures as completed consolidations: promoters organized them, the financial press reported them, and in many cases options on the constituent plants had been signed. The pool comprises 30 abandoned deals announced between 1891 and 1907, clustering at the same 1899 and 1902–03 peaks as completed consolidations; 24 deals carry a reconstructed constituent roster and enter the deal-level analysis. Constituent firms are linked to the

patent record with the same matching procedure used for treated firms. A firm on either side—completed or abandoned—exits the panel in the years after it loses independence through a subsequent acquisition, so that a unit stops serving as a control once it has itself received the consolidation treatment; its earlier years and the rest of its deal are retained. This symmetric sample restriction keeps the two groups on the same footing, and the estimate is stable across variations that instead retain the post-acquisition years (Appendix Table C8). Appendix B details construction, sources, and validation.

I compare the two groups in a dynamic difference-in-differences at the deal level, in event time relative to each deal's earliest public announcement:

$$y_{dt} = \alpha_d + \delta_\tau + \sum_{\tau=-10,\,\tau\neq-1}^{25} \beta_\tau \cdot \mathbb{1}[t - a_d = \tau] \cdot \mathbb{1}[\text{Completed}_d] + \varepsilon_{dt}, \tag{2}$$

where d indexes deals, $\tau = t - a_d$ is time relative to deal d's announcement year a_d, α_d are deal fixed effects, δ_τ are event-time fixed effects, and Completed_d indicates a consolidation that closed rather than one abandoned; the remaining notation follows Equation 1. The coefficients β_τ capture the completed-minus-abandoned difference relative to the year before announcement, and standard errors are clustered at the deal level. For patents I estimate the β_τ with the interaction-weighted estimator of Sun and Abraham (2021); for the sparse breakthrough outcome I use two-way fixed effects, whose static estimate is stable where the interaction-weighted standard errors collapse.

Figure 6 presents the comparison. In calendar time, mean patenting per deal is low and similar across the two groups before 1895 and diverges afterward, with completed consolidations rising while abandoned deals stay flat (panel a); breakthroughs follow the same pattern from a lower base (panel b). The event studies in panels (c) and (d) quantify the divergence. Averaging over the post-announcement years, completed consolidations add about 2.2 patents per deal-year relative to abandoned ones, and about 0.2 breakthroughs per deal-year—set against pre-announcement means of 0.77 patents and 0.02 breakthroughs per deal-year. The magnitudes are smaller than in the intensive-margin analysis below because the panel pools deals on both margins, with and without prior patenting, which lowers the baseline. Appendix Table C7 reports the corresponding estimates.

The pool is small, so the comparison offers limited statistical power. Reassuringly, the estimate is positive in every variation examined: dropping influential deals one at a time, retaining absorbed constituents' later years, excluding constituent firms that publicly refused their deals, using Nelson's dating, and splitting by the documented cause of abandonment (Appendix Tables C8 and C9). If anything, the comparison understates the effect: deals often collapsed because strong constituent firms declined to sell, so the pool tilts toward firms whose independent prospects were good.

4.4 Matched Non-Merging Firms

My third comparison keeps the event-study design of the previous two but adds a matched control group of non-merging firms in the same technological area, in a standard difference-in-differences framework.

Data limitations make the selection of control firms more challenging. Ideally, one would observe a wide range of comparable firms active before the GMW (regardless of their patenting status), with rich information on their product market sector, main technological area and asset size. Unfortunately, such wealth of data is unavailable for most firms so far back in time.

To address limitations in data availability, I distinguish between the intensive and extensive margins of patenting and adopt tailored strategies for each. The intensive margin—focusing on firms that had patented prior to the GMW—forms the core of my analysis, as the patent record offers richer pre-period data and a natural comparison pool of non-merging firms with established innovation histories in similar technological areas. For the extensive margin—where firms had no prior patents—I employ a complementary but more limited analysis. Because we cannot condition on being in the patent record, I rely on firms that appeared in the 1900 Moody's Manual of Industrial Securities. Restricting to comparable firms in this source that were not patenting before the GMW and were not involved in the GMW, I obtain a control group for consolidations on the extensive margin. For both intensive and extensive margins, I restrict the control group to firms with asset sizes comparable to those of individual firms that merged.[26]

On the intensive margin, my preferred approach compares merging to non-merging firms in the same technological area. This comparison identifies the effect of consolidation by contrasting merged entities against what other firms achieved in the same technological and macroeconomic conditions. To ensure that comparisons reflect differences driven by consolidation rather than baseline innovative capacity, I stratify firms into six groups based on their pre-1895 patenting levels and reweight the control group to match the distribution observed among treated firms.[27] In my preferred specification, the panel is balanced and years with no patents are coded as zeroes.

The dynamic specification carries over the design of Equation 1, now with the matched control group and technology-by-year fixed effects, and with event time measured in calendar years around 1895 rather than around each deal:

$$y_{itc} = \alpha_i + \delta_{tc} + \sum_{m=1885, m\neq 1894}^{1930} \beta_m \cdot \mathbb{1}[t = m] \cdot \mathbb{1}[GMW\ Firm_i] + \varepsilon_{itc}, \tag{3}$$

where c indexes the technological area (one of nine CPC sections) in which a firm is most active, δ_{tc} are technology-by-year fixed effects, and $GMW\ Firm_i$ indicates a merging firm; the remaining notation follows Equation 1. The β_m capture the treatment effects, the outcomes are patents and breakthrough patents, the sample runs from 1885 to 1930, and standard errors are clustered at the firm level. The static specification is analogous:

$$y_{itc} = \alpha_i + \delta_{tc} + \beta \cdot \mathbb{1}[t \geq 1895] \cdot \mathbb{1}[GMW\ Firm_i] + \varepsilon_{itc}, \tag{4}$$

26. Approximate range of \$0.5–10 million, averaging around \$3 million. This asset band is a one-sided screen: I exclude only firms positively attested in Moody's Manual as larger than the range, and retain firms absent from Moody's on the presumption that they are small. Finer comparability comes not from this coarse screen but from the reweighting of controls to the pre-1895 patenting distribution described above.

27. The six strata correspond to the following pre-1895 patenting levels: 1, 2–3, 4–9, 10–25, 26–75, 76+.

where the difference-in-differences parameter of interest is β.

Unlike the previous two comparisons, I define treatment effects for this design in calendar time, with 1895 marking the start of the GMW period, rather than in time relative to each deal. This choice reflects both the historical nature of the merger wave—a sharp, concentrated episode driven by common shocks—and an econometric consideration specific to this design: both margins condition on pre-1895 patenting outcomes, a selection on the dependent variable. Treated units therefore approach the wave with a mechanical relative-time structure that control firms, which have no deal date, cannot share; calendar time harmonizes the margins across the two groups, so that the conditioning cancels out. Calendar time also sidesteps known TWFE bias for staggered adoption designs (Callaway and Sant'Anna 2021; Sun and Abraham 2021). Nevertheless, Appendix C.2 shows the main results to be robust to a relative-time specification anchored at each deal's earliest public announcement and estimated with the interaction-weighted estimator of Sun and Abraham (2021), and confirms flat pre-trends even when examining GMW constituent firms.

For the extensive margin, while methodologically similar, data limitations restrict the scope of analysis. I use the 1900 Moody's Manual of Industrial Securities to identify manufacturing firms that neither patented before nor participated in the GMW, and thus create a control group of firms at risk of patenting. Since firms on the extensive margin had no pre-GMW patenting history by construction and not all of them did patent after 1895, I cannot define technological areas c from the patent record. Instead, in this analysis the time fixed effects δ_{tc} in Equations 3 and 4 vary by one of eight economic sectors.[28] The primary outcome is a time-varying indicator that switches from 0 to 1 in the first year a firm is granted a patent after the GMW and remains 1 in all subsequent years, capturing the cumulative entry of firms into patenting. This analysis, while more limited, provides complementary evidence on how mergers affected the entry of new innovators.

4.4.1 Main effects

Figure 7 presents the results from the main event study specification (Equation 3), displaying β_m estimates and their 95 percent confidence intervals. Across both patenting and breakthrough measures, we observe flat pre-trends, followed by a sustained and large increase in the aftermath of the merger wave. These effects build gradually and peak after around 15–20 years—consistent with a long-run ramp-up in innovation output rather than short-lived bump.

The magnitudes are substantial. As shown in columns 1 and 2 of Table C1 and reported on the Figures, consolidating firms already patenting before 1895 gained, on average, approximately 6 additional patents and 0.56 additional breakthroughs per year. In relative terms, because of relatively low initial levels, these intensive-margin effects correspond to a 310 percent increase in the number of patents and a 536 percent increase in breakthroughs relative to the treated units' pre-wave mean.

28. I use sectors c reported in Moody's, matched to 1949 SIC codes reported in Nelson's worksheets.

4.4.2 Robustness

In Appendix C.3, I show the robustness of these results. Figure C8 shows that my baseline results are not driven by the long strings of zero-patent years that thinly documented firms contribute: I re-estimate Equation 3 in an unbalanced panel that retains only observations from three years before a firm's first patent through three years after its last patent. Because the window is defined by realized patenting rather than by an independent measure of firm survival, it does not by itself rule out differential exit; it does show that the results are not an artifact of periods when firms were unlikely to be at risk of patenting, and it yields quantitatively similar results. Figure C7 shows that my results hold when I trim the sample by excluding the top and bottom 5 percent of firms. Figure C6 shows that restricting to (known) horizontal consolidations yields similar results, suggesting the effect of consolidation is not driven by complementarities. Appendix C.2 shows the main results in a relative-time specification.

Two further checks address the concern that the control group is matched on technology rather than on product market. Control firms operating in the very niches the consolidations absorbed may themselves have been affected by the consolidation—through competition, the prospect of eventually being acquired, or the same demand conditions that prompted the merger—and technology-by-year fixed effects do not absorb industry-level shocks such as tariffs or demand that moved whole product markets after 1895. A comparison within the narrow product market is not an available response, since in the narrowest markets the consolidation often *was* the market; I therefore proceed from the two opposite directions. First, I exclude from the control group every firm whose 1885–94 patenting falls, by majority, in the largest pre-1895 CPC main group of any consolidation—about 7 percent of control firms; Appendix B details the construction and a menu of looser and stricter rules. Second, I add industry-by-year fixed effects based on the name-based classification of firms into twelve broad product-market categories described in Section 3. Table 2 reports the resulting ladder. Excluding same-niche controls leaves the estimates essentially unchanged, at 5.98 patents and 0.56 breakthroughs per year, and adding the industry dimension moves them to 5.69 and 0.57, with confidence intervals well away from zero.[29] Appendix Table C4 reports stricter and looser niche exclusions—down to removing every control firm whose modal CPC subclass coincides with a consolidation's largest subclass, a third of the control group—together with the fixed-effects variants taken separately; the estimates move little under every feasible rule.

Finally, I address a measurement choice: all results in this section use the grant year of patents rather than the filing year. While filing dates more closely approximate when inventions were created, and are used elsewhere in the paper, grant dates better capture when intellectual property rights were assigned to specific firms—a critical consideration when studying a period of intense firm-level transformations and when many patents were bought from independent inventors, rather than produced in-house. Nevertheless, Figure C10 and Table C5 confirms that results using filing dates yield substantively

29. The industry specification is estimated on the roughly half of control firms whose names support a classification. Re-estimating the specification of column 2 on that same subsample yields 5.88 patents, so the modest movement reflects the added fixed effects rather than the change in sample. The corresponding event studies preserve flat pre-trends (Appendix Figures C5a and C5b).

Table 2: Effect of Consolidation on Innovation

	(1) Baseline	(2) Excluding same-niche controls	(3) Adding industry-by-year FE
Panel A: Patents			
GMW Firm	5.989	5.977	5.690
	(1.910)	(1.910)	(1.938)
	[2.245, 9.732]	[2.234, 9.720]	[1.891, 9.488]
% of treated pre-mean	310%	310%	295%
Treated pre-mean	1.929	1.929	1.929
Panel B: Breakthroughs			
GMW Firm	0.563	0.562	0.569
	(0.181)	(0.181)	(0.183)
	[0.208, 0.919]	[0.207, 0.917]	[0.211, 0.928]
% of treated pre-mean	536%	535%	542%
Treated pre-mean	0.105	0.105	0.105
Same-niche controls excluded	N	Y	Y
CPC section × year FE	Y	Y	Y
Industry × year FE	N	N	Y
# Treated	79	79	79
# Controls	11,722	10,858	5,369
N	542,846	503,102	250,608

Note: This table re-estimates the difference-in-differences specification in Equation (4) under product-market restrictions. Column 1 reproduces the baseline. Column 2 excludes control firms whose 1885–94 patenting falls, by majority, in the largest pre-1895 CPC main group of any consolidation. Column 3 additionally replaces the sample with firms carrying a name-based industry classification and adds industry-by-year fixed effects to the technology-by-year effects. Stratification weights are recomputed within each sample. Standard errors clustered at the firm level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

similar conclusions.

4.4.3 Extensive margin

Beyond the intensive margin of firms with pre-merger innovation activity, I also examine whether consolidation induced non-patenting firms to begin obtaining patents. Figure 8 presents event study estimates for the probability that a firm obtained at least one patent after the merger wave, while Table C1, column 4, reports the average effects.

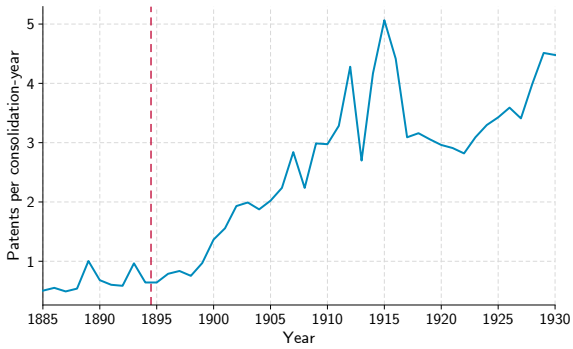
Recall that, because firms on the extensive margin had no pre-GMW patenting history, I use the 1900 *Moody's Manual* to identify comparable non-merging firms as controls, with time fixed effects defined at the industry level rather than by technological areas.

The results indicate that consolidations were 23.1 percentage points more likely to begin patenting compared to non-merging firms. This corresponds to a near tripling of

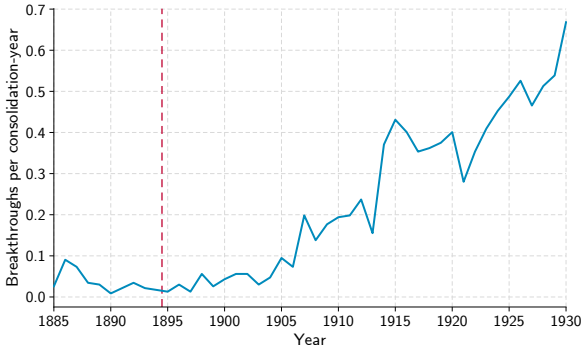
entry relative to the 13.8 percent post-1895 average in the control group. By 1930, 52.4 percent of GMW firms had patented at least once, versus 19.8 percent of controls. Figure 8 shows that this gap emerged rapidly in the first decade after 1895, suggesting a swift innovation response.

Appendix C provides complementary evidence: robustness to defining outcomes by filing rather than issue year, and to an unbalanced panel including consolidations later acquired by other firms.

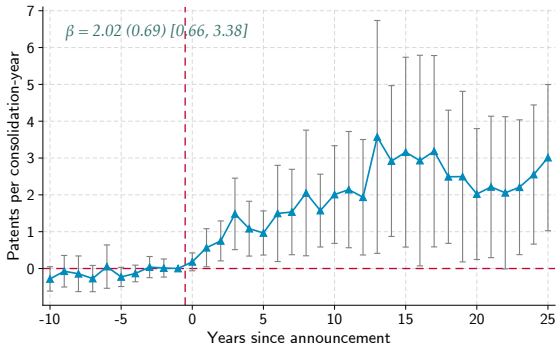
Figure 5: Patenting of Consolidating Firms Around the GMW



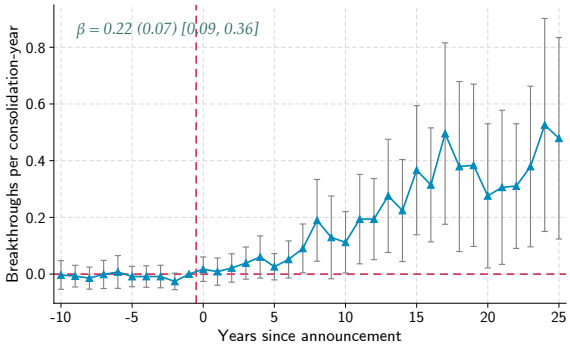
(a) Patents, calendar time



(b) Breakthroughs, calendar time



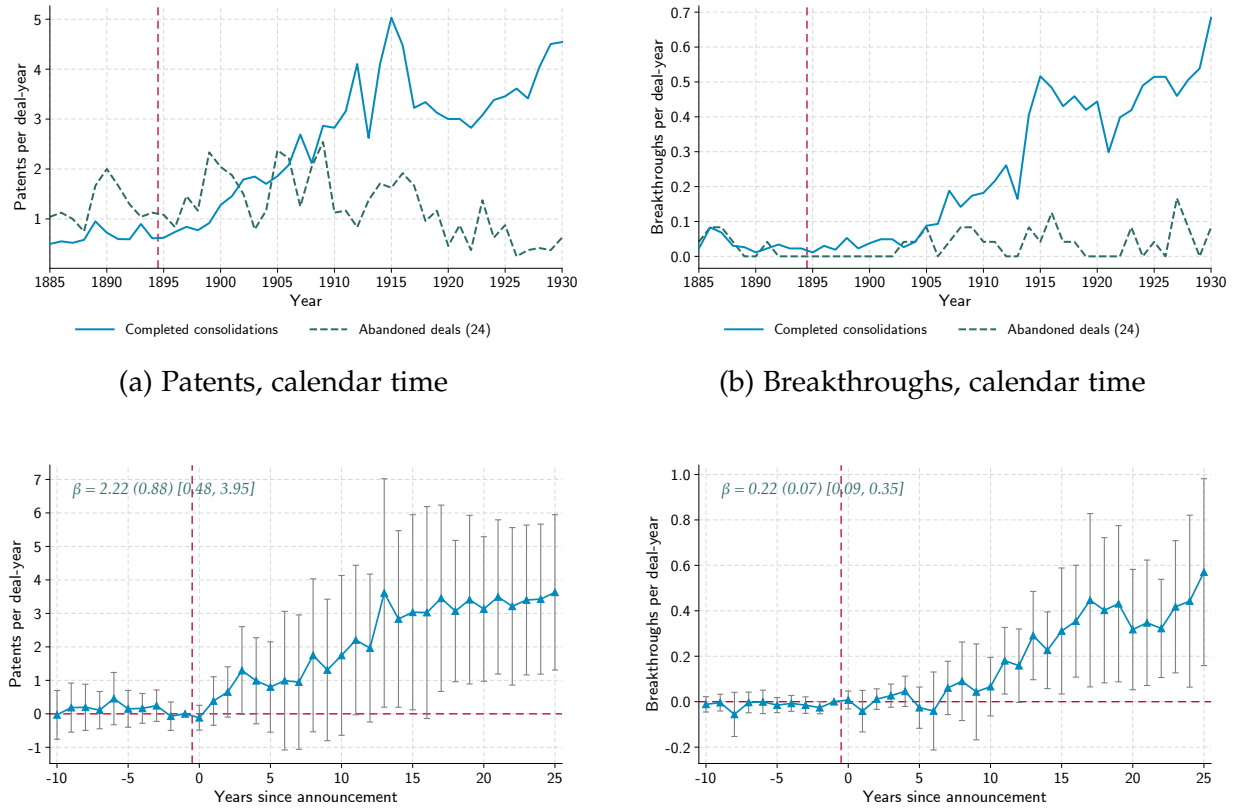
(c) Patents, event study



(d) Breakthroughs, event study

Note: This figure pools all GMW consolidations, with no control group. The top row plots means in calendar time: mean patents (panel a) and mean breakthrough patents (panel b) per consolidation-year, with the vertical line marking 1895. The bottom row plots within-firm event studies of patents (panel c) and breakthrough patents (panel d) in years relative to each deal's earliest public announcement (Equation 1), with 95 percent confidence intervals and standard errors clustered at the consolidation level. The annotated boxes report the average of the post-announcement event-time coefficients, with the standard error in parentheses and 95 percent confidence interval in brackets. Because no control group absorbs secular or technology-specific trends, the comparison is descriptive.

Figure 6: Completed Consolidations versus Paper Mergers



(a) Patents, calendar time

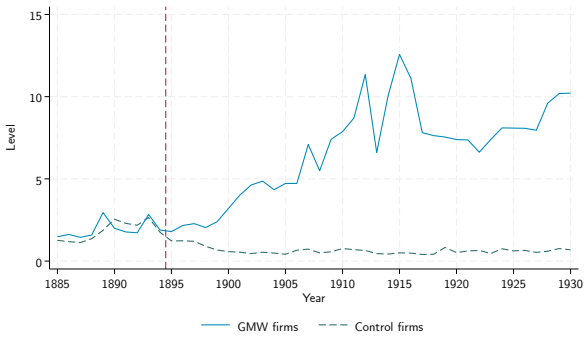
(b) Breakthroughs, calendar time

(c) Patents, event study

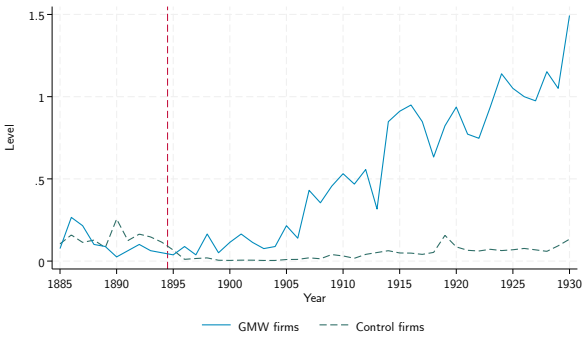
(d) Breakthroughs, event study

Note: This figure compares completed consolidations to abandoned "paper" mergers at the deal level. The top row plots group means in calendar time: mean patents (panel a) and mean breakthrough patents (panel b) per deal-year, for completed consolidations (solid) and the 24 abandoned deals (dashed), with the vertical line marking 1895. A firm on either side exits the panel in the years after it is absorbed into a later consolidation, a symmetric sample restriction that retains its earlier years and the rest of its deal. The bottom row plots event studies of the completed-minus-abandoned difference in years relative to each deal's earliest public announcement (Equation 2), with deal fixed effects and standard errors clustered at the deal level. Panel (c), patents, uses the interaction-weighted estimator of Sun and Abraham (2021); panel (d), breakthroughs, reports the two-way fixed-effects event study, since for this sparse outcome the interaction-weighted standard errors collapse while the static estimate does not. The annotated boxes report the average of the post-announcement event-time coefficients, with the standard error in parentheses and 95 percent confidence interval in brackets. Bands are 95 percent confidence intervals.

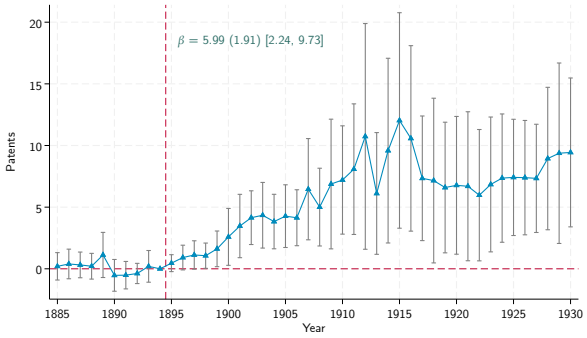
Figure 7: Effect of Consolidation on Innovation—Intensive Margin Firms



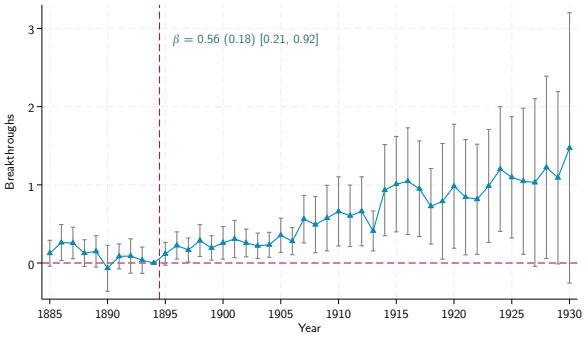
(a) Patents, outcome levels



(b) Breakthroughs, outcome levels



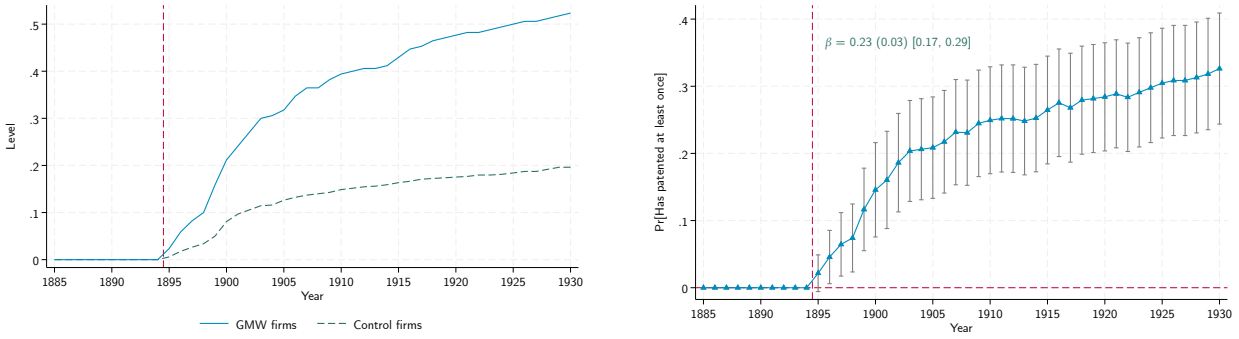
(c) Patents, event study estimates



(d) Breakthroughs, event study estimates

Note: This figure presents the firm-level effect of consolidation on innovation for intensive-margin firms, from Equation (3). The left column reports patents and the right column breakthrough patents. The top row (panels a and b) shows IPW-weighted outcome levels; the bottom row (panels c and d) shows the β_m estimates and their 95 percent confidence intervals, together with the static estimate β from Equation (4), with its standard error in parentheses and confidence interval in brackets. Standard errors are clustered at the firm level.

Figure 8: Effect of Consolidation on Probability of Patenting—Extensive Margin Firms



(a) Outcome levels

(b) Event study estimates

Note: This figure presents event study estimates from Equation (3) showing the firm-level effect of consolidation on the probability of patenting at least once, for firms that had not patented prior to 1895. Panel (a) shows outcome levels; panel (b) shows the β_m estimates and their 95 percent confidence intervals, together with the static estimate β from Equation (4), with its standard error in parentheses and confidence interval in brackets. Standard errors are clustered at the firm level.

4.5 Heterogeneity

To better understand how innovation effects vary systematically across different firm characteristics, I adapt the difference-in-differences specification in Equation (4) by replacing the single treatment indicator with separate post-1895 indicators for each category, allowing the consolidation effect to differ across dimensions of interest.[30]

I begin by examining heterogeneity by market concentration. Figure 9 presents heterogeneity by the initial market share captured by consolidations. I use data from Lamoreaux (1985), who compiled information from archival sources, including industry publications, newspaper accounts, and Nelson's worksheets. Since this data is not available for all consolidations, I partition firms into three categories: (i) consolidations without reported market share information (likely less successful or smaller mergers), (ii) those achieving substantial but less than 70 percent market share, and (iii) those capturing over 70 percent of their markets.[31]

The results reveal a consistent pattern: while innovation increases are positive across all categories, they are substantially larger for consolidations that achieved the highest level of market concentration. The pattern holds for both patents and breakthroughs. Firms with the highest initial market share added 1.57 breakthroughs per year, compared to an average effect of 0.56.

Next, Figure 10 shows that the increase in breakthrough innovations concentrates in more science-based technologies, like chemistry, metallurgy, and electronics. I group CPC sections in three categories based on widespread assessment of their historical R&D intensity and closeness to the scientific frontier (Mowery and Rosenberg 1998; Chandler 1990; Arora et al. 2024). The science-based technology group encompasses chemistry, metallurgy, scientific instruments, computing, electronics, and telecommunications (CPC sections C, G and H). The engineering and industrial technology group spans mechanical engineering, manufacturing processes, vehicles, weapons, heating systems and cross-cutting technologies (CPC sections F, B and Y). The infrastructure and consumer-oriented technology group includes agriculture, food processing, medical devices, construction, textiles and apparel (CPC sections A, E and D).[32]

The results reveal a striking pattern in the distribution of innovation gains across technology categories. While consolidations increased patenting activity across all three groups, the quality composition of these innovations differs markedly. Science-based

30. Formally, I estimate:

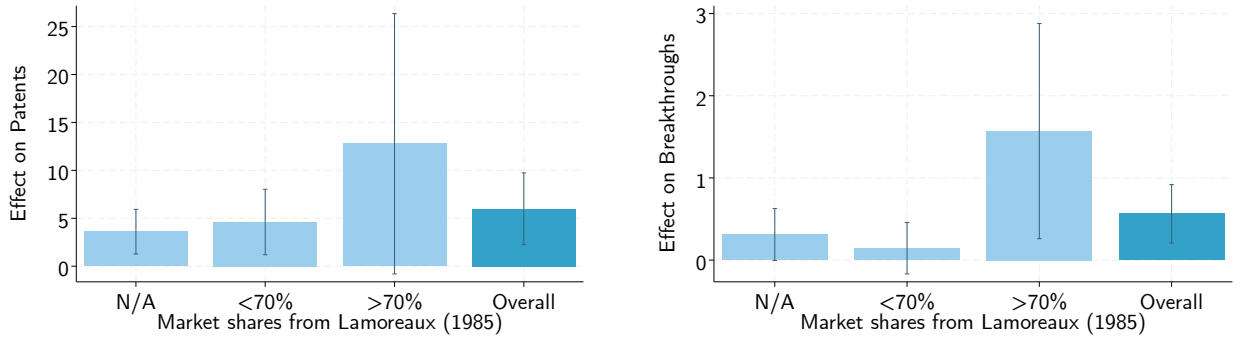
$$y_{itc} = \alpha_i + \delta_{tc} + \sum_{h=1}^{H} \beta_h \cdot \mathbb{1}[t \geq 1895] \cdot \mathbb{1}[GMW\ Firm_i] \cdot \mathbb{1}[C_i = h] + \varepsilon_{itc},$$

where $h \in 1, \ldots, H$ indexes treatment categories defined by the variable C_i (e.g., market share groups or technology types), and all other notation follows Equation (4).

31. On the intensive margin, roughly half of consolidations fall into the first group, with the remaining split evenly between the second and third (see Table 1).

32. Firms are grouped by the patent-weighted modal CPC section of their full 1885–1940 portfolio, so this classification is not strictly pre-determined and could shift if consolidation altered a firm's technological mix; the caveat applies to this firm-level heterogeneity only, as the technology-level exposure in Section 6 is defined from pre-1895 patents alone.

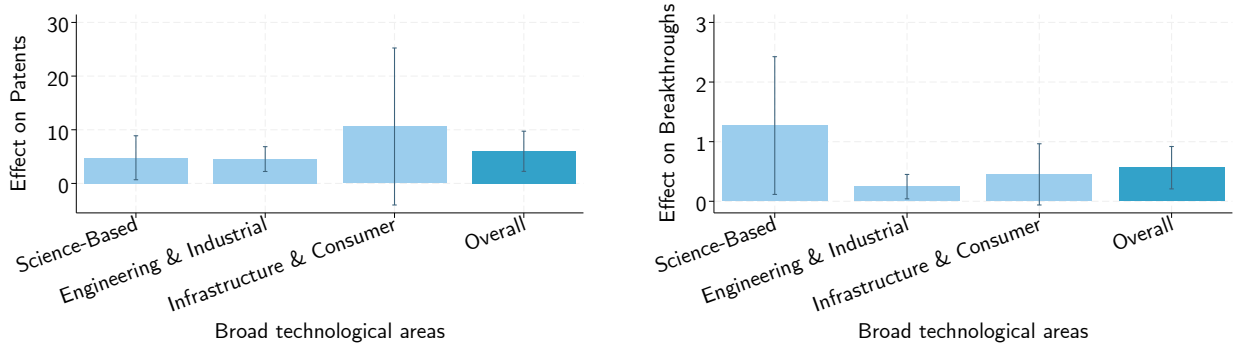
Figure 9: Heterogeneity in Consolidation Effects by Market Concentration



(a) All patents (b) Breakthrough patents

Note: This figure shows how the effect of consolidation on innovation varies with market concentration. Firms are categorized into three groups: those without reported market share information (likely smaller and less successful mergers), those achieving substantial but less than 70 percent market share, and those capturing over 70 percent market share. Market share data comes from Lamoreaux (1985). Each bar shows the estimated effect of consolidation on patents (panel a) and breakthroughs (panel b). Error bars represent 95 percent confidence intervals, computed from SEs clustered at the firm level.

Figure 10: Heterogeneity in Consolidation Effects by Broad Technological Area



(a) All patents (b) Breakthrough patents

Note: This figure shows how the effect of consolidation on innovation varies with broad technological area. Firms are categorized into three groups according to the CPC section where they patent the most: (i) sections C, G and H for the more science-based technology group, (ii) sections B, F and Y for engineering and industrial technologies, (iii) sections A, D and E for infrastructure and consumer-oriented technologies. Each bar shows the estimated effect of consolidation on patents (panel a) and breakthroughs (panel b). Error bars represent 95 percent confidence intervals, computed from SEs clustered at the firm level.

technologies exhibit the strongest response in breakthrough innovations (1.27), with over a quarter of additional patents representing high-impact discoveries (given the effect on patents is 4.78). In contrast, infrastructure and consumer-oriented technologies show much larger increases in total patenting but substantially lower breakthrough rates, with only a small fraction of new patents achieving breakthrough status (about 4 percent).

This pattern suggests that the innovation benefits of consolidation were particularly pronounced in research-intensive fields where systematic R&D and laboratory-based discovery processes could yield greater innovation gains.

In Appendix C.6, I present additional heterogeneity analyses across sectors, integration types, and the degree of "business success". Splitting by 1949 SIC industry codes, although results are less precisely estimated, I find qualitatively stronger effects in the machinery and chemicals industries. When examining integration type, I find that horizontal consolidations—which comprised the majority of the GMW—drive the main results. Vertical and mixed integration types show larger but substantially more imprecise effects, likely reflecting their limited representation in the sample. More successful consolidations, as reported in Livermore (1935), saw the largest increases in innovation.

4.6 Alternative Inference Approaches

I implement three complementary strategies that more closely resemble the construction of treatment units in that they aggregate several control units into one counterfactual merger. The key limitations of these alternative strategies are that they are only feasible when richer pre-treatment data is available, making them unsuitable for the extensive margin and for the sparsest outcomes such as lab formation, and less precise for breakthroughs, and that they offer less flexibility for conducting analyses that rely on linearity like the spatial decomposition by R&D lab proximity (Equation 5). Nevertheless, each of these approaches replicates the core results, both qualitatively and quantitatively

First, I conduct a placebo merger analysis. Within each technological field, I randomly select sets of control firms in order to match the pre-period innovation levels of treatment units. I then compare real to placebo consolidations in a standard difference-in-difference framework. Standard errors are obtained by boostrapping the procedure 1,000 times.

Second, I use a synthetic control (SC) approach. For each consolidating firm, an optimal weighted combination of control firms is constructed to match the treatment unit's pre-merger innovation trajectory. These weights are selected using standard optimization techniques, following Abadie (2021).

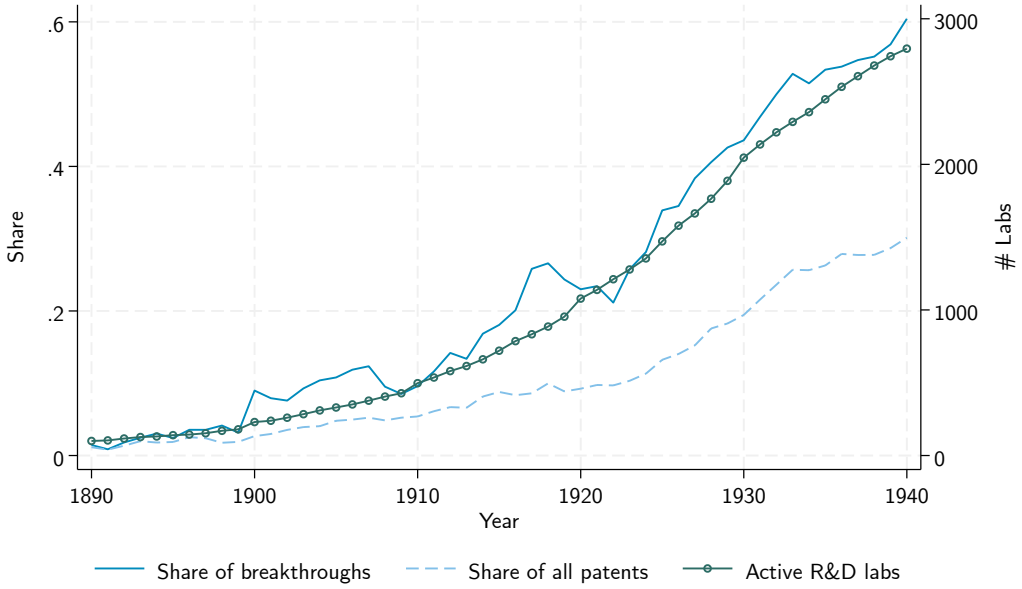
Third, I implement synthetic difference-in-differences (SDID), following the method developed by Arkhangelsky et al. (2021). This approach combines elements of difference-in-differences and synthetic control by assigning weights to both units and time periods, aligning pre-treatment trends while accounting for latent time-varying confounders. By embedding these weights in a two-way fixed effects regression, SDID improves robustness to violations of parallel trends and yields valid inference under weaker conditions.

Each of these strategies yields comparable estimates. Approximate Bayesian model averaging (BMA) across specifications, including my preferred one, suggests that the increase in patents per firm per year is around 6.35, and the increase in breakthroughs is around 0.58—closely aligned with the effects from the baseline analysis. Appendix D provides full results across robustness designs and details the BMA methodology.

5 Corporate R&D Labs and Innovative Productivity

Why did consolidation lead to substantial firm-level innovation gains? The previous section established that firms already patenting before the wave experienced approximately four-fold increases in patents and six-fold increases in breakthrough patents after merging during the GMW, with particularly pronounced gains for firms achieving greater market concentration and operating in science-based technologies. Several non-mutually exclusives explanations are possible, as big firms may: (i) face stronger incentives to innovate when they can better appropriate returns from R&D (Schumpeter 1942; Spulber 2013), (ii) benefit from improved access to capital markets, enabling longer-term, riskier projects (Atkinson and Lind 2019), (iii) adopt distinct organizational practices (Chandler 1977), such as setting up dedicated R&D facilities like industrial research laboratories (Mowery 1990).

Figure 11: Lab-owning firms' role in U.S. innovation expanded dramatically



Note: The figure shows the expanding role of lab-owning firms in U.S. innovative activity between 1890 and 1940. The solid line shows the share of all U.S. breakthrough innovations produced by firms with R&D laboratories; the dashed line reports their share of all patents. The connected line with circular markers shows the number of active R&D labs over time. Lab-owning firms are identified using National Research Council surveys. Breakthroughs are defined using the Kelly et al. (2021) measure.

Influential narratives of American technological and economic development before World War II emphasize the role of lab-based corporate R&D (Chandler 1977; Mowery and Rosenberg 1998; Gertner 2013; DeLong 2022). Descriptive evidence on the U.S. innovation ecosystem also suggests a strong relationship between labs and firm-level innovation. Figure 11 shows that, by 1940, lab-owning firms—though a small minority of innovative firms—accounted for about 30 percent of all new patents. More strikingly, they captured an outsized share of breakthrough patents, roughly 60 percent.

This section argues that R&D lab establishment is indeed a key mechanism in the causal chain from "bigness"—large firm size and market dominance—to innovation. I proceed in two steps. First, I show that consolidation directly leads to laboratory establishment (the bigness to labs link) using the same quasi-experimental framework from Section 4. Most consolidation-driven innovation gains concentrate spatially around laboratory facilities. Second, I investigate whether laboratories genuinely enhance firm innovative productivity (the labs to innovation link) by expanding the analysis beyond GMW firms. I build a matched inventor-firm panel and use it to measure a lab-specific productivity premium from inventor sorting and other confounding factors.

5.1 Consolidation and Laboratory Establishment

To assess whether bigness leads to lab establishment, I estimate the GMW treatment effect on firms' probability of having at least one (dated) R&D lab. This outcome is based on labs that have a non-missing establishment date in the NRC surveys. The analysis thus requires that the distribution of missing establishment dates does not correlate with differential trends in lab openings across GMW and non-GMW firms. Because the NRC surveys begin in 1920, moreover, a lab enters the data only if it survived and was documented with a usable establishment date; the outcome is therefore best read as durable, later-observed lab ownership, and its interpretation also relies on survey inclusion and lab survival not differing systematically across GMW and non-GMW firms. The specifications are the same as in Section 4, Equations 3 and 4. These estimates use the baseline specification—without the same-niche exclusion or industry-by-year fixed effects of the product-market robustness checks—because lab establishment is a rarer outcome estimated on a smaller sample, so that they are directly comparable to column 1 of Table 2.

Results show consolidation significantly increased laboratory adoption. Among firms with pre-merger patents (intensive margin), treatment firms were 16 percentage points more likely to establish at least one R&D lab during the post-wave period (Figure 12 and Table C1, column 3). Notice that, because this outcome features limited variation, observations are grouped in 3-year bins. By 1930, 35.4 percent of GMW firms on the intensive margin had at least one (dated) R&D lab, compared to only 2.6 percent of control firms.[33]

For firms with no prior patenting activity (extensive margin) consolidation similarly increased laboratory establishment. Consolidated firms were 4.5 percentage points more likely to establish at least one R&D lab during the post-wave period (Figure 12 and Table C1, column 5). By 1930, 11.8 percent of GMW firms on the extensive margin had at least one (dated and active) R&D lab, while only 2.2 percent of control firms did.[34]

33. If we consider labs without an establishment date and extend the horizon to the last NRC survey in my data, 1946, 54 percent of GMW firms on the intensive margin had at least one lab, compared to less than 5 percent in the control group. See Table C11 in the Appendix.

34. If we consider labs without an establishment date and extend the horizon to the last NRC survey in my data, 1946, 20 percent of GMW firms on the extensive margin had at least one lab, compared to about 4 percent in the control group. See Table C11 in the Appendix.

Appendix C presents results on heterogeneity by initial market share and technological area that mirror those for patenting outcomes. It also includes additional event studies for the number of active labs.

5.1.1 Spatial Decomposition: Linking Labs to Innovation Gains

In both modern and historical patent data, it is hard to discern if an inventor who assigns a patent to a firm works for that firm directly or not.[35] One way to address this problem is to rely on physical distance (Nicholas 2009). If laboratories represent the primary mechanism linking consolidation to innovation, then we should observe that innovation gains concentrate among lab-owning firms, and especially in proximity to their laboratory facilities. To test this prediction, I decompose the consolidation effects on patents and breakthroughs from Section 4 by the geographical proximity of inventors to laboratory locations.

I partition patent and breakthrough counts into mutually exclusive categories that, together with the small residual group noted below, are jointly exhaustive, based on the relative locations of a firm's R&D lab locations (if any) and its inventors. Each patent is classified as: (i) within 50km of a dated, active lab; (ii) within 50km of an undated lab; (iii) beyond 50km of any lab operated by the firm; or (iv) assigned to firms without any labs.[36] By linearity of the difference-in-differences estimator in Equation 4, running separate regressions of the form:

$$y^j_{itc} = \alpha^j_i + \delta^j_{tc} + \beta^j \cdot \mathbb{1}[t \geq 1895] \cdot \mathbb{1}[GMW\ Firm_i] + \varepsilon^j_{itc}, \tag{5}$$

with partitions j such that $y_{itc} = \sum_j y^j_{itc}$, yields a set of coefficients β^j that sum to the total effect β reported in Table C1. Table 3 presents these decomposed effects, with each coefficient's share of the total effect indicating the relative contribution of different categories to the overall innovation response.

The spatial decomposition in Table 3 is consistent with the laboratory mechanism, especially for higher-impact patents. It allocates the causal total across locations—an accounting exercise rather than a separate identification of the labs' own contribution—and shows that more than half of the breakthrough effect comes from patents filed within 50km of an R&D lab. An additional 30 percent comes from patents assigned to lab-owning firms but filed by inventors beyond this 50km radius, consistent with a key function of early laboratories: testing and evaluating outside inventions before purchasing them. Overall, 4 in 5 breakthrough innovations causally attributed to the Great Merger Wave came from lab-owning firms.

35. For instance, an independent inventor might sell their patent rights to a firm before issuance.

36. A residual category, not shown but included in the analysis, captures patents with missing inventor location or those assigned by inventors residing abroad, accounting for about 2–3 percent of the total effect. See Appendix B for details on data construction.

Table 3: Effect of Consolidation on Firm-Level Innovation—Decomposition by R&D Lab Proximity

Outcome:	Total effect (1)	Within 50km active lab (2)	Within 50km undated lab (3)	Outside 50km any lab (4)	No lab firm (5)
Panel A: Patents					
GMW Firm	5.989	2.088	0.391	1.922	1.375
	(1.910)	(1.311)	(0.179)	(0.612)	(0.444)
	[2.245, 9.732]	[-0.482, 4.658]	[0.040, 0.741]	[0.723, 3.122]	[0.504, 2.245]
Share of total (%)	—	34.9	6.5	32.1	23.0
Panel B: Breakthroughs					
GMW Firm	0.563	0.265	0.021	0.173	0.090
	(0.181)	(0.135)	(0.010)	(0.058)	(0.049)
	[0.208, 0.919]	[0.000, 0.530]	[0.000, 0.041]	[0.059, 0.288]	[-0.005, 0.185]
Share of total (%)	—	47.1	3.7	30.7	16.0
Controls	Y	Y	Y	Y	Y
# Firms	11,801	11,801	11,801	11,801	11,801
N	542,846	542,846	542,846	542,846	542,846

Note: This table decomposes the effect of consolidation on firm-level innovation by geographical proximity to R&D laboratories. Each column represents a mutually exclusive category based on inventor location relative to a firm's labs: within 50km of an active (dated) lab, within 50km of an undated lab, outside 50km of any lab, and firms without labs. Residual categories covering patents with missing inventor location or from inventors residing abroad account for two to three percent of the total effects and are omitted. The decomposition follows Equation (5) where coefficients sum to the total effects reported in Table C1. "Share of total (%)" indicates each category's contribution to the overall innovation response. Standard errors clustered at the firm level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

5.2 The R&D Lab Innovation Premium

The evidence presented so far shows that bigness led to greater firm innovation and investments in industrial research labs. However, the observed association between laboratory adoption and firm-level innovation could reflect sorting of talented inventors or selection of inherently innovative firms into lab ownership rather than genuine productivity advantages from laboratory organization.

To fix ideas, define the total innovation output of firm j as Y_j, given by the sum of inventor-level productivity:

$$Y_j = \sum_{i \in \mathcal{E}_j} y_{ij}, \qquad y_{ij} = f\left(\alpha_i, \phi_j, \lambda_j\right), \tag{6}$$

where $\mathcal{E}_j$ is the set of inventors active at firm j, α_i is inventor ability, ϕ_j captures all firm-specific characteristics affecting productivity (such as size, market position, management quality) apart from laboratory adoption, $\lambda_j \in \{0, 1\}$ indicates whether the firm organizes R&D through a laboratory. Assume that $f(\cdot)$ is increasing in its first two arguments.

The proposed mechanism linking labs to greater firm-level innovation would imply systematic innovative productivity differences by firm lab status. Lab-owning firms are more innovative on average: $\mathbb{E}\left[Y_j | \lambda_j = 1\right] > \mathbb{E}\left[Y_j | \lambda_j = 0\right]$, as implied by Figure 11. Because part of this pattern is mechanical (larger firms both employ more inventors and are more likely to operate labs), let us focus on productivity. The key empirical question is whether an observational productivity premium for lab-owning firms π^{lab} exists at the inventor level:

$$\pi^{\text{lab}} = \mathbb{E}\left[y_{ij} | \lambda_j = 1\right] - \mathbb{E}\left[y_{ij} | \lambda_j = 0\right] > 0, \tag{7}$$

and, if so, why.

Three alternative mechanisms could underlie an observed inventor-level productivity premium π^{lab}. Each condition below generates an observed premium only when the remaining determinants—inventor ability or other firm attributes—are held fixed or balanced, so the implications should be read as *ceteris paribus*:

1. **Inventor sorting**. Labs attract the most talented inventors (but do not raise productivity):

$$\mathbb{E}\left[\alpha_i | \lambda_j = 1\right] > \mathbb{E}\left[\alpha_i | \lambda_j = 0\right] \Rightarrow \pi^{\text{lab}} > 0 \qquad \text{but} \qquad f\left(\alpha_i, \phi_j, 1\right) = f\left(\alpha_i, \phi_j, 0\right) \quad \forall i, j.$$

Alternatively, once inventor sorting is accounted for, lab-owning firms could still display an innovative productivity premium because of either:

2. **Selection on firm characteristics**. Inherently more innovative firms simply choose to establish labs (but do not raise productivity):

$$Corr(\phi_j, \lambda_j) > 0 \Rightarrow \pi^{\text{lab}} > 0 \quad \text{but} \quad f\left(\alpha_i, \phi_j, 1\right) = f\left(\alpha_i, \phi_j, 0\right) \quad \forall i, j,$$

3. or a **Lab-specific productivity premium**. Laboratories genuinely enhance firm-level innovation production:

$$f\left(\alpha_i, \phi_j, 1\right) > f\left(\alpha_i, \phi_j, 0\right) \quad \forall i, j \quad \Rightarrow \pi^{\text{lab}} > 0$$

To assess the roles of these explanations, I expand the analysis beyond GMW firms to consider all patenting firms, building a matched inventor-firm panel spanning 1875 to 1950.[37] This dataset provides sufficient variation in inventor mobility and lab adoption to both: separate firm-level productivity effects from inventor sorting; and investigate labs' contribution to firm-specific innovative productivity. To do this, I employ a two-way fixed effects decomposition following Abowd, Kramarz, and Margolis (1999).

5.2.1 AKM Framework

The standard AKM model specification is:

$$\ln y_{it} = \alpha_i + \psi_{j(it)} + X'_{it}\beta + \varepsilon_{it} \tag{8}$$

where y_{it} denotes the inventive output of inventor i in year t. The term α_i captures time-invariant inventor ability, while $\psi_{j(it)}$ represents the productivity effect associated with the firm j where i is in year t. The vector X_{it} includes time-varying controls (inventor experience and field-by-year fixed effects), and ε_{it} is the residual. My preferred measure of inventive output y_{it} is quality-weighted patents, where each patent's weight depends on the continuous version of the Kelly et al. (2021) breakthrough score.[38]

The AKM firm effect ψ_j captures all firm-level productivity determinants and nests the lab-specific premium:

$$\psi_j \;=\; h(\phi_j, \lambda_j), \qquad \frac{\partial h}{\partial \phi_j} > 0, \quad \frac{\partial h}{\partial \lambda_j} \gtrless 0,$$

where ϕ_j represents other firm attributes, and $\lambda_j \in \{0, 1\}$ indicates whether the firm ever operated an R&D laboratory during the sample. Because ψ_j is a single, time-invariant firm effect, the decomposition speaks to differences between lab-owning and other firms rather than to a contemporaneous lab treatment; the lab-adoption and inventor-mover analyses below address the latter.

For identification, I assume the following mean independence condition:

$$\mathbb{E}\left(\varepsilon_{it} \;\middle|\; X_{11}, \ldots, X_{NT},\, j(1,1), \ldots, j(N,T),\, \alpha_1, \ldots, \alpha_N,\, \psi_1, \ldots, \psi_J\right) \;=\; 0, \tag{9}$$

37. See Appendix B for details on inventor disambiguation and variable construction.

38. Because the score can be negative, I employ a simple exponential transformation calibrated to roughly match the patent value distribution in Kogan et al. (2017). Kelly et al. (2021) demonstrate positive correlation between their quality measure and Kogan et al. (2017) dollar values. See Appendix B for details. Quality-weighted patents offer two key advantages over simple counts: they are strictly positive over a spell (enabling the logarithmic transformation that is standard in the AKM framework), and they capture both quantity and significance of inventive output. Variance decomposition for the additional outcomes is in Table E2.

where N, J, T denote the total number of inventors, firms and time periods respectively. This assumption underlies all AKM applications (Abowd, Kramarz, and Margolis 1999; Card, Heining, and Kline 2013; Bonhomme et al. 2023) and has two main implications. First, the model assumes additive separability: firm and inventor effects enter linearly without interactions or complementarities. Second, inventor mobility across firms must be exogenous to transitory shocks in productivity. Conditional on observed characteristics, inventors do not systematically move to higher- or lower-performing firms in response to short-term fluctuations in output.

In practice, an additional requirement for robust identification is sufficient mobility of inventors across firms. The incidental parameter bias resulting from limited mobility has been shown to severely affect variance decomposition and in particular quantification of sorting (Bonhomme et al. 2023; Kline, Saggio, and Sølvsten 2020). Short spells also contribute to biasing estimates. I address potential biases arising from limited variation in three ways. First, I employ the heterogeneity-robust bias correction developed by Kline, Saggio, and Sølvsten (2020). The KSS method requires a stronger leave-one-out connectedness condition: the estimation set remains connected when any observation is taken out. Secondly, I follow the best practices outlined in Bonhomme et al. (2023) and collapse observations at the spell level, which ensures unbiasedness in the presence of serial correlation within spell.[39] Third, to exploit as much variation as possible, I retain spells where an inventor did not assign their patents to a firm, by creating pseudo-firm categories: $j = -1$ for unassigned patents and $j = 0$ for patents assigned to individuals. Thus, I let these observations increase mobility and contribute to the estimation of the fixed effects.

The final estimation sample comprises approximately 94,000 inventors and 18,000 firms; the underlying 1.3 million inventor-year observations collapse to roughly 227,000 inventor-firm spells, the unit at which the fixed effects are estimated. Because most inventors patented only once and the leave-one-out connectedness condition is demanding, the estimation sample is considerably smaller than the total data at our disposal (see Table B7). However, in this application limited mobility bias is strong and thus the bias-correction is indispensable (see Appendix Table E2). Moreover, because in practice we are interested in career inventors working within corporate R&D labs, discarding the many sporadic or loosely connected patenters is not concerning.

I conduct two key tests to validate the core assumptions of the AKM framework for the main outcome of quality-weighted patent output. Appendix Figure E2 presents the results. If the model in Equation 8 is valid and its assumptions satisfied, then, conditional on controls, changes in inventor output between spells should align with changes in firm effects. Panel (a) in Figure E2 tests this by plotting output changes against firm effect changes for inventors moving between firms of different quality quartiles. The close alignment between these series supports the model validity. If firm and inventor effects enter additively without complementarities, model residuals should be flat across

39. A spell (i, s) is the contiguous run of years inventor i is associated with firm $j = j(i, s)$. Let $\bar{y}_{is}$ be the average quality-weighted output over the spell. I estimate the spell-level counterpart of Equation (8): $\ln \bar{y}_{is} - X'_{is}\beta = \alpha_i + \psi_{j(is)} + u_{is}$, where X_{is} contains dummies for the calendar year in which the spell starts and inventor cumulative experience (linear and squared) measured at the end of the spell. I fit β first, and estimate AKM on the residualized log output.

the joint distribution of fixed effects. Panel (b) shows average residuals by deciles of both inventor and firm effects. Consistently with the model assumptions, residuals are generally flat.[40]

Additional validation tests are reported in Appendix E. As a robustness check, Appendix E.1 implements the semi-structural approach of Bonhomme, Lamadon, and Manresa (2019), which also allows for complementarities between worker and firm types and confirms the main findings.

5.2.2 From Firm Effects to a Lab-specific Premium

Table 4: Innovative Productivity, Inventor Ability and Firm Effects by R&D Lab Status

	Lab = 0 (1)	Lab = 1 (2)	Difference (3)
Innovative productivity $\ln y_{ijt}$	0.028	0.225	0.197
	(1.310)	(1.401)	[0.000]
Inventor ability α_i	0.091	0.041	-0.050
	(1.013)	(1.112)	[0.000]
Firm effect ψ_j	-0.141	0.081	0.222
	(0.965)	(0.796)	[0.000]

Note: This table shows average values and differences for key variables in the AKM sample. Quality-weighted patents (logged) measures inventor-level innovative productivity; inventor fixed effects (α_i) capture time-invariant individual ability; firm fixed effects (ψ_j) measure firm-level productivity effects. Both fixed effects are derived from estimating the AKM model in Equation 8. Lab status is based on NRC surveys. The difference column shows lab-owning minus non-lab-owning firms. Standard deviations in parentheses, p-values in brackets for differences.

Evidence from the AKM sample confirms the existence of sizable productivity premium for lab-owning firms. Table 4 shows that, as hypothesized in Equation 7, inventor productivity as measured by quality-weighted patents is 0.2 log points higher in firms operating R&D laboratories.

The variance decomposition in Table 5 further reveals that firm-level effects (which nest lab adoption) account for approximately 33 percent of total explained variation in productivity ($\text{Var}(\psi)/\text{Var}(\psi+\alpha)$). Inventor effects account for 75 percent, and the sorting term contributes negatively at -8 percent. Crucially, the firm share rises from 26 percent in 1875-1904 to 32 percent in 1905-1950, coinciding with the proliferation of corporate R&D laboratories.

40. The largest deviations appear among the lowest deciles, which is similar to the finding in Card, Heining, and Kline (2013).

Table 5: Bias-Corrected Variance Decomposition of Innovative Productivity

Sample:	Full sample (1)	1875–1904 (2)	1905–1950 (3)
$Var(y - X'\beta)$	1.784	1.791	2.028
R^2	0.250	0.312	0.236
$Var(\psi)/Var(\psi+\alpha)$	0.328	0.264	0.321
$Var(\alpha)/Var(\psi+\alpha)$	0.755	0.757	0.713
$2\,Cov(\psi,\alpha)/Var(\psi+\alpha)$	-0.083	-0.021	-0.034
$Corr(\psi,\alpha)$	-0.084	-0.024	-0.035
Observations	1,310,550	125,741	1,085,024
Spells	227,284	30,587	190,690
Firms	18,286	3,451	15,183
Inventors	94,040	10,618	85,544
Movers (%)	61.38	82.93	58.15

Note: This table reports bias-corrected variance decomposition results from the AKM model (Equation 8) using quality-weighted patent output as the dependent variable. The decomposition separates the contributions of inventor ability (α), firm productivity (ψ), and sorting between inventors and firms. Column 1 shows results for the full sample (1875-1950); columns 2-3 show results for early (1875-1904) and later (1905-1950) sub-samples. Bias correction follows the KSS method. The sample includes only inventors and firms in the largest connected set satisfying the leave-one-out connectedness condition. Movers are inventors matched to multiple firms during their careers.

AKM estimates reject inventor sorting (explanation 1) as the primary driver of the observed association between firm innovation and lab-ownership. Table 4 shows that average inventor fixed effects—capturing time-invariant inventor ability net of firm-specific productivity—are actually slightly lower in lab-owning firms. Conversely, lab-owning firms exhibit substantially higher firm-level productivity effects. Therefore, the innovative productivity gap π^{lab} in Equation (7) does not arise from sorting of superior inventors to lab-owning firms, but might be explained by a lab-specific productivity premium. Appendix Figure E1 shows the overall distribution of inventor and firm effects by lab status.

5.2.3 Lab Premium Persists within Size and Technology Classes

A firm's choice to operate an R&D lab is non-random, and likely to correlate with unobserved determinant of innovative productivity. Thus, without exogenous lab adoption, we cannot definitely distinguish between selection (explanation 2) and a lab-specific premium (explanation 3). However, suggestive evidence supports a direct, independent, and positive effect of organizing R&D through laboratories on firm-level innovation.

To start, greater innovative productivity may be spuriously correlated with lab adoption through its association with firm size or technological domain. Large firms may

both more easily afford dedicated research infrastructure and also innovate more for reasons entirely unrelated to their laboratories. Under this view, the apparent lab effect would disappear when we properly account for scale. Alternatively, laboratories might simply reflect involvement in inherently more innovative, high-reward technological domains. In this scenario the lab effect would be an artifact of underlying technological opportunity rather than organizational capability.

Figure 13 shows that lab-owning status is associated with greater firm-level productivity even within large firms or science-based technologies. Panel (a) shows average firm effects (ψ) by bigness and lab status, where the baseline category comprises firms that neither are large nor have a lab. For this exercise, a firm is big if it ever appeared in a list of top 100 firms by asset size between 1900 and 1948. The figure shows that having an R&D lab is associated with about 0.25 SDs higher firm effects, regardless of size. Similarly, Panel (b) shows average firm effects by lab status and type of technology in which they specialize. Here too the figure shows higher firm productivity associated with research laboratories, even within technology type. Within science-based domains, firms without a lab have 0.17 SDs higher firm effects than baseline firms, but this premium grows to 0.43 SDs for firms with a lab. This evidence indicates that correlation with size and technological domain does not fully explain the observed productivity premium associated with operating R&D infrastructure. It does not, by itself, separate a lab-specific premium (explanation 3) from selection on other persistent firm characteristics (explanation 2); the adoption and mover analyses that follow speak more directly to that distinction.

5.2.4 Evidence from Firms and Inventors Gaining Labs

Two complementary sources of variation provide additional evidence supporting an independent effect of R&D labs on innovative productivity: changes in firm-level effects following lab establishment, and changes in inventor productivity when moving between firms with different lab status.

First, I test whether gaining an R&D lab is associated with higher firm-specific productivity net of sorting changes. To do so, I focus on firms present in both the 1875-1904 and 1905-1950 AKM sub-samples (see Table 5, columns 2 and 3), for which a change in their fixed effect can be computed, $\Delta\psi_j$. I run the following cross-sectional regression:

$$\Delta\psi_{jc} = \delta_c + \beta GainedLab_j + \gamma X_j + \varepsilon_{jc}, \tag{10}$$

where j indexes the firm, δ_c are fixed effects by broad technological area c (CPC Section), $GainedLab_j$ is a binary indicator for firms that did not have a lab before 1904, but did after. X_j is either an indicator for making it in the list of top 100 firms by asset size after 1904 or for having consolidated during the GMW. For comparability, firm effects ψ_j are standardized. ε_j are heteroskedasticity-robust errors.

Table 6 presents results from the regression specification in Equation 10. Firms that gained laboratories between sub-samples experienced large increases in firm-specific productivity compared to firms that did not, controlling for technological area and increases in size and consolidation activity. This pattern is consistent with a lab-specific

productivity premium rather than with time-invariant firm characteristics that merely correlate with lab adoption, though the broad definition of the two sub-periods means it cannot pin down exactly when, after adoption, the gain accrued.

Table 6: Effect of Lab Establishment on Changes in Firm Productivity

	Outcome: Change in Firm FE $\Delta\psi$		
	(1)	(2)	(3)
Opened Lab	0.576	0.580	0.461
	(0.114)	(0.117)	(0.117)
	[0.352, 0.800]	[0.351, 0.809]	[0.231, 0.690]
Became Big		-0.067	
		(0.395)	
		[-0.842, 0.708]	
GMW Firm			0.818
			(0.254)
			[0.319, 1.317]
CPC Section FE	Y	Y	Y
N	1,091	1,091	1,091

Note: This table reports cross-sectional regression results examining changes in firm productivity following lab establishment (Equation 10). The dependent variable is the change in firm effects between the 1875-1904 and 1905-1950 AKM sub-samples. "Opened Lab" indicates firms that gained a laboratory between periods. "Became Big" indicates firms entering the top 100 by asset size after 1904 (Collins and Preston 1961). "GMW Firm" indicates participation in Great Merger Wave consolidations. Robust standard errors are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Second, I exploit inventor mobility to test whether being in a firm with an R&D lab improves individual productivity. To do so, within the AKM sample, I identify about 14,000 inventors who have been in the same non-lab-owning firm for at least three years, then transition to another firm and stay there for at least four years. Next, I compare firm movers who joined a firm with a lab to those who joined another non-lab-owning firm in the following specification:

$$\ln(1+y_{itmc}) = \alpha_i + \gamma_t + \delta_m + \kappa_c + \sum_{m=-3, m\neq -1}^{3} \beta_m \cdot \mathbb{1}[\text{RelTime}_{im} = m] \cdot \text{ToLab}_i + \varepsilon_{itmc} \quad (11)$$

where i indexes the inventor, t calendar year, m time relative to the move, c the technological area (CPC Section), and α_i, γ_t, δ_m, κ_c are the respective fixed effects; y_{itmc} is the

quality-weighted patent measure used for AKM,[41] ToLab_i is an indicator for inventors who moved from non-lab to lab firms, ε_{itmc} are the errors, clustered at the inventor level. The coefficients of interest are captured by β_m.

Figure 14 suggests inventor gain significantly from joining an R&D lab-owning firm. Panel (a) shows that, even in the raw data, levels and trends are perfectly aligned before the move and then diverge sharply after the move.[42] Panel (b) reports the event study estimates from Equation 11. Even accounting for calendar time, relative time, technological composition, and individual effects, inventors that move to firms with labs experience a significant increase in their innovative productivity, of about 0.12 log points in the years after the move. This pattern is consistent with a lab-specific premium that benefits individual inventors' productivity.

5.3 Innovation Processes in Industrial Research Laboratories

The productivity advantages associated with industrial research laboratories reflected systematic organizational innovations in the research and development process itself. Historical evidence reveals that labs fundamentally transformed innovation from individual invention to collaborative, team-based research and development. Lab-based inventors were 20 percent more likely to work in teams compared to inventors in firms without dedicated research facilities, with repeat collaborations occurring significantly more frequently within laboratory settings (Hartog et al. 2024). These collaborative structures enabled systematic exploration of radical technological possibilities: lab-based teams were 3.3 percentage points more likely to patent novel combinations of technology classes, representing more than a 60 percent increase above baseline rates (Hartog et al. 2024).

Industrial research laboratories also attracted and concentrated scientific talent in ways that enhanced firms' innovative capabilities. Large corporations operating research laboratories were significantly more likely to employ prominent scientists and engage in scientific publication, bridging the gap between academic research and commercial application (Arora et al. 2024). This scientific orientation enabled labs to systematically monitor technological frontiers, evaluate external innovations, and pursue longer-term research projects that individual inventors or smaller firms could not sustain (Hounshell and Smith 1988). The combination of collaborative research structures, scientific expertise, and systematic R&D processes helps explain why laboratory adoption generated genuine firm-level productivity premiums.[43]

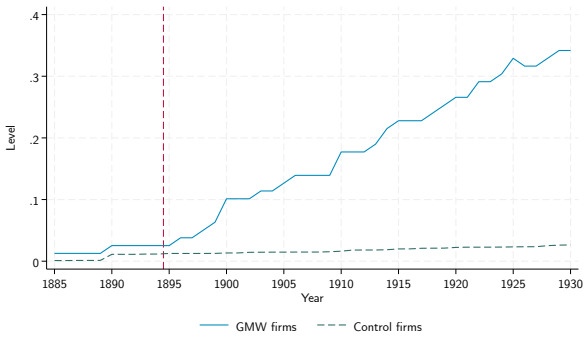
41. Note that the AKM estimation is done at the spell level, where output is never zero, so that the outcome can be log-transformed. Yearly observations can contain zeros, so I use $ln(1+x)$ to transform the outcome.

42. Notice that the spike at time 0 is an artifact of panel construction where firms are assigned on the bases of patent assignment, so that a firm move is mechanically associated with a patent. This is not an issue for the AKM estimation as observations are collapsed at the spell level, thus taking no stance on within-spell timing of productivity shocks. Comparing moving inventors nets out that mechanical effect.

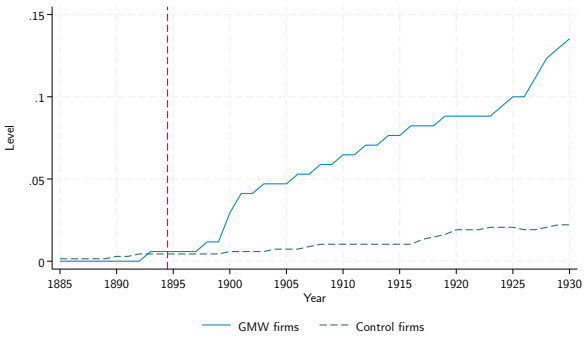
43. Using modern citations to historical patents as a proxy for quality, Nicholas (2009) finds patents originating within 30 miles of corporate labs to be lower quality; however, Kelly et al. (2021) discuss the potential biases and limitations in using long-lagged citations. Using the share of breakthroughs out of total patents issued as a measure of quality, I find that for lab-owning firms between 1905-1940 patents

originating from within 50km of a lab had a breakthrough rate above 20 percent, compared to about 17 percent outside that range. Patents assigned to firms with no labs or not assigned to firms at all had a breakthrough rates of 11 and 6 percent, respectively. See Table E1 in the Appendix.

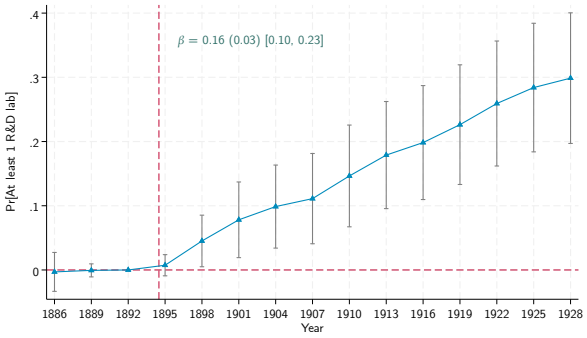
Figure 12: Effect of Consolidation on Having any R&D Lab



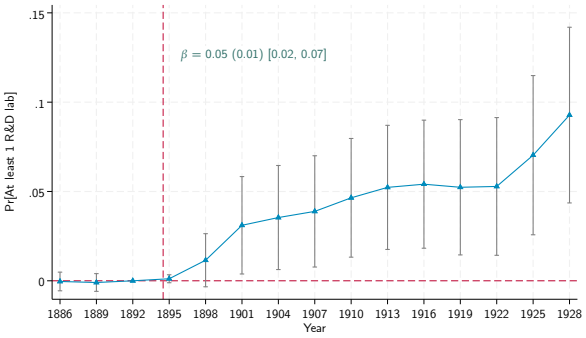
(a) Intensive margin, outcome levels



(b) Extensive margin, outcome levels



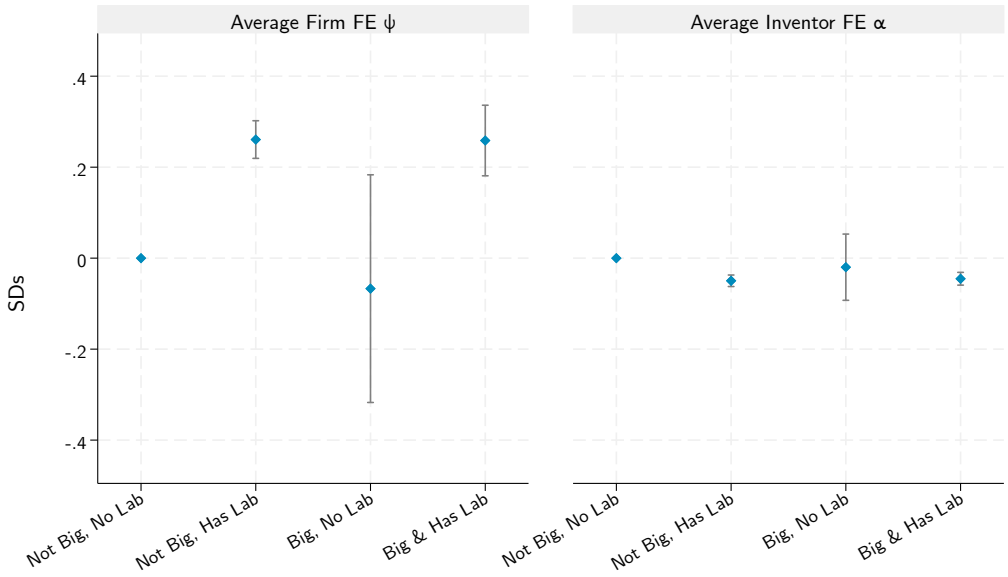
(c) Intensive margin, event study estimates



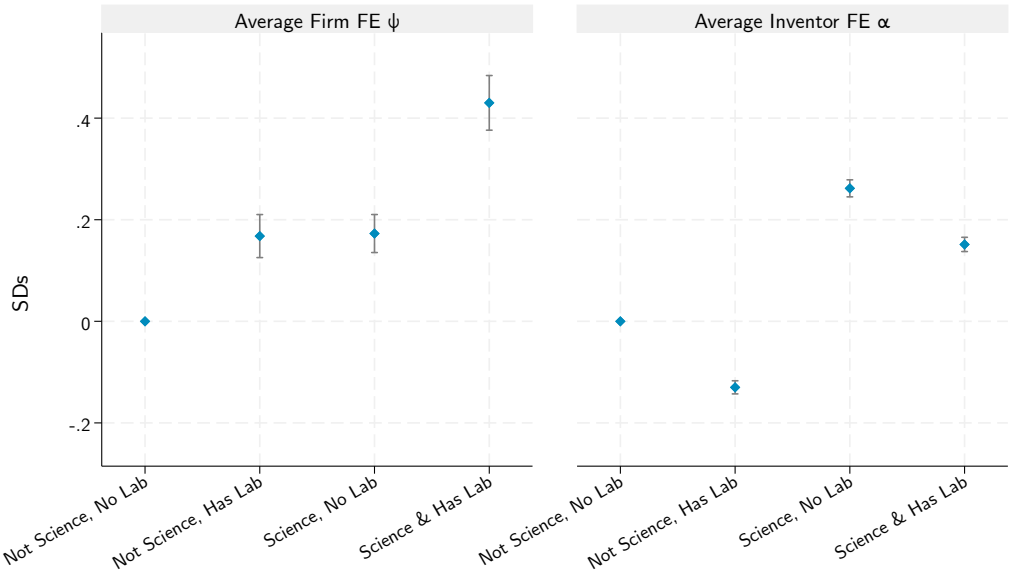
(d) Extensive margin, event study estimates

Note: This figure presents event study estimates showing the firm-level effect of consolidation on having any R&D lab, for firms with pre-merger patents (intensive margin, left column) and firms with no pre-merger patents (extensive margin, right column). The specification is analogous to Equation (3), but observations are grouped in 3-year bins. The top row (panels a and b) shows IPW-weighted outcome levels; the bottom row (panels c and d) shows the β_m estimates and their 95 percent confidence intervals, together with the static estimate β from Equation (4), with its standard error in parentheses and confidence interval in brackets. In the event-study panels the x-axis reports the first year in each bin, so that the values shown for 1895 correspond to the 1895–1897 bin. Standard errors are clustered at the firm level.

Figure 13: Firm and Inventor Effects by Lab Status and Firm Characteristics



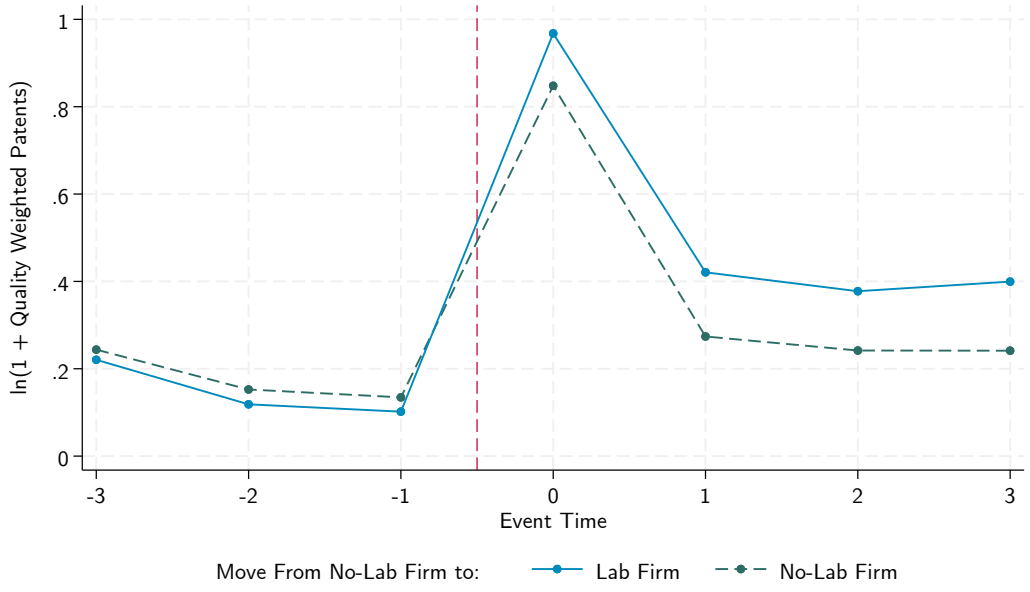
(a) By firm size and lab status



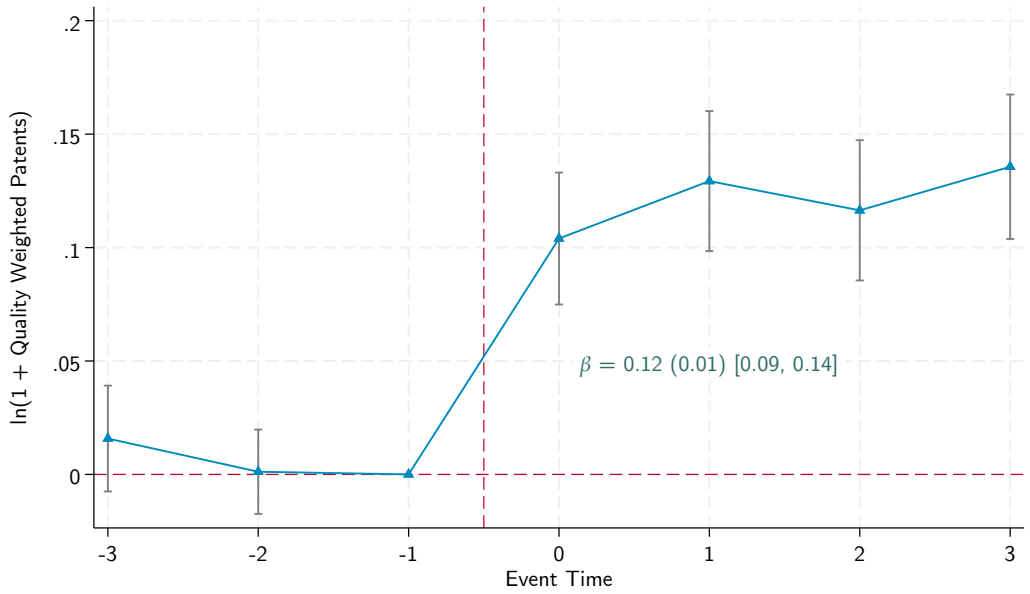
(b) By technological domain and lab status

Note: This figure shows average firm and inventor fixed effects from Equation 8. Panel (a) compares effects by firm size (big firms are those appearing in the top 100 by asset size, 1900-1948) and lab ownership. Panel (b) compares effects by technological specialization (science-based vs. other domains) and lab ownership. Both effects (ψ and α) are measured in standard deviations. Average effects are relative to the baseline category, first from the left.

Figure 14: Effect of Joining Lab-Owning Firms on Inventor Productivity



(a) Outcome levels



(b) Event study

Note: This figure shows the effect of joining a lab-owning firm on inventor productivity. The sample includes inventors who worked at non-lab firms for at least three years before moving to another firm for at least four years. Panel (a) shows raw outcome levels comparing moves to lab-owning versus non-lab firms. Panel (b) shows event study estimates from Equation 11 controlling for inventor, calendar year, event time, and technological area fixed effects. The outcome is quality-weighted patent output. Standard errors are clustered at the inventor level. 95 percent confidence intervals are shown in gray.

6 Spillover Effects of the Great Merger Wave

Did the innovation gains of consolidating firms produce a net positive effect on U.S. technological development? The previous sections showed that firms undergoing consolidation during the Great Merger Wave experienced substantial increases in patenting, breakthrough innovations and R&D lab establishment. However, market consolidation could trigger additional responses. On the one hand, product market and technological barriers to entry (possibly from preemptive patent strategies) could reduce innovative efforts by new firms, independent inventors and other competitors in affected technologies. More broadly, decreased competition could harm overall innovation, depending on a number of structural factors (Aghion, Akcigit, and Howitt 2014; Akcigit and Ates 2023; Bryan and Williams 2021). On the other hand, investments in R&D infrastructure may spark technological races, produce knowledge spillovers and open new breakthrough opportunities (Bloom, Schankerman, and Van Reenen 2013).

To assess the GMW's broader impact on U.S. innovation, I shift focus from individual firms to patenting activity across entire technological domains, capturing the response of all actors rather than the consolidating firms alone. I examine this question separately for established and emerging technologies. For technologies that were already active before the merger wave (at least one patent before 1895), I measure how consolidation exposure affected subsequent innovation levels. For technological domains that emerged after 1895 (no patent before 1895), I analyze whether greater exposure to the merger wave delayed or accelerated their initial development.

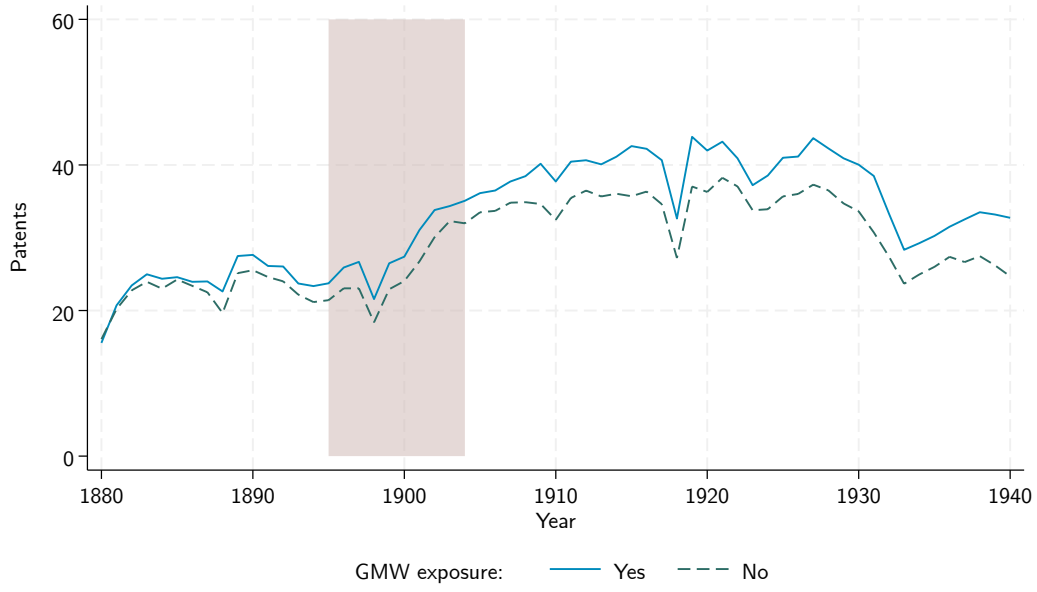
Figures 15 and 16 visualize how innovation patterns evolved for technologies with different exposure to the Great Merger Wave. In established technologies (Figure 15), the number of patents rises steadily for both exposed and unexposed technologies, with exposed domains eventually exhibiting slightly higher levels. The gap is more pronounced for breakthrough innovations, where GMW-exposed technologies show a visibly steeper increase starting around 1903-1905. In contrast, Figure 16 shows the cumulative share of emerging technologies that reach at least one patent or breakthrough after 1895. Here, the emergence of breakthroughs appears delayed in technologies more exposed to the GMW, even though the timing of first patents is roughly similar across groups.

These patterns suggest that the broader innovation effects of the Great Merger Wave varied by technological maturity: among established domains, those more exposed to consolidation exhibit higher rates of breakthrough innovation, while among emerging domains, exposure is associated with slower breakthrough development. In contrast, total patenting shows only modest differences by exposure status, highlighting that the most meaningful effects operate on the quality—not the quantity—of innovation.

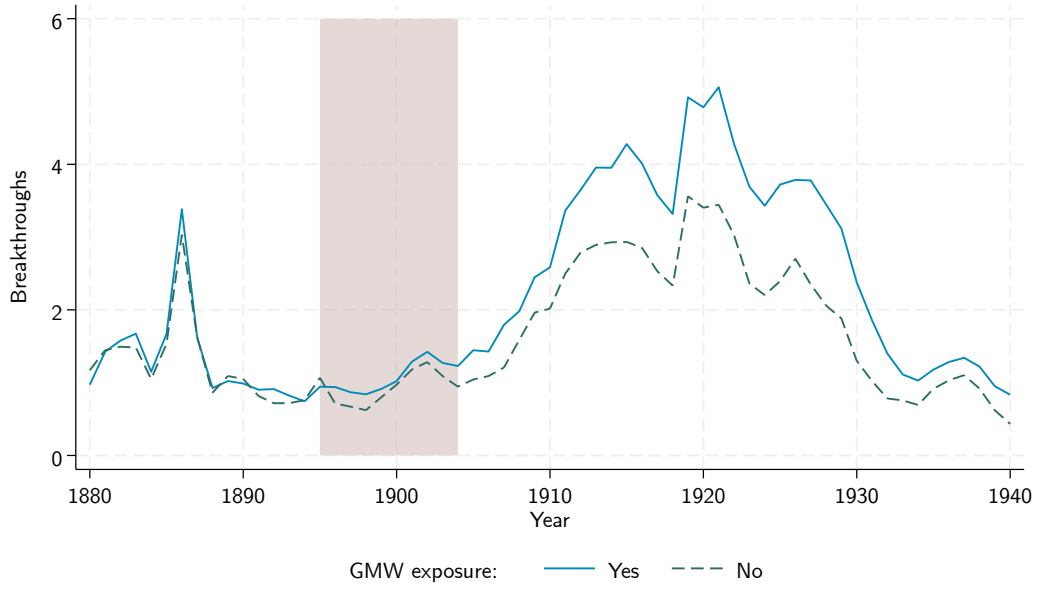
The formal analysis below reveals a second key source of heterogeneity: whether a technology is science-based or not.[44] The interaction between technological maturity and scientific intensity produces a clear pattern in the results. In science-based fields, consolidation exposure is associated with consistently positive or neutral effects on breakthrough innovation—both in established and emerging domains. By contrast,

44. As in Section 4, science-based technologies are defined as CPC sections C, G and H, encompassing chemistry, metallurgy, scientific instruments, computing, electronics, and telecommunications.

Figure 15: Innovation levels by GMW exposure status for established technologies



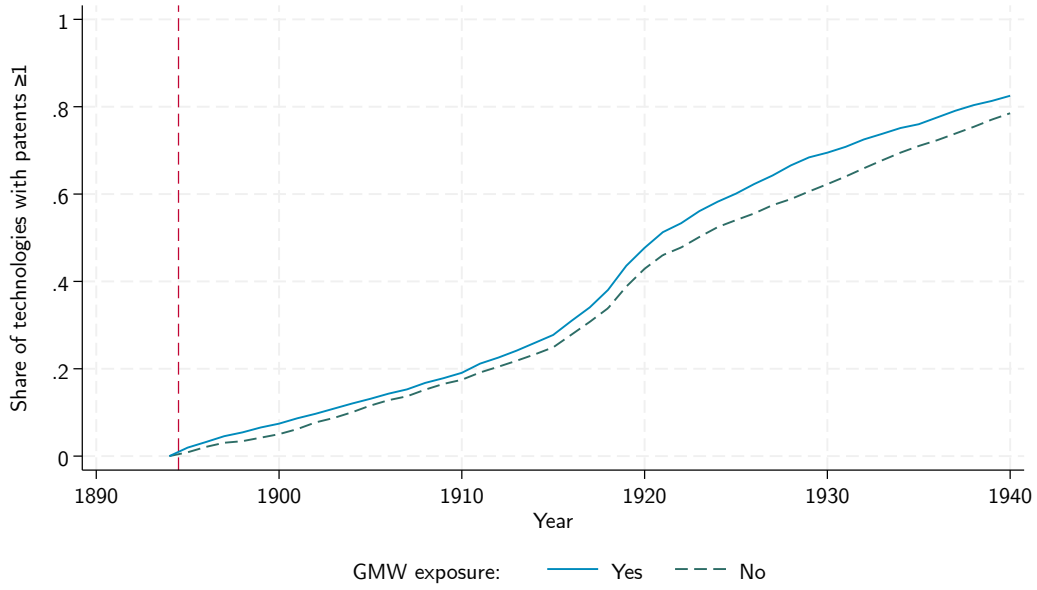
(a) Patents



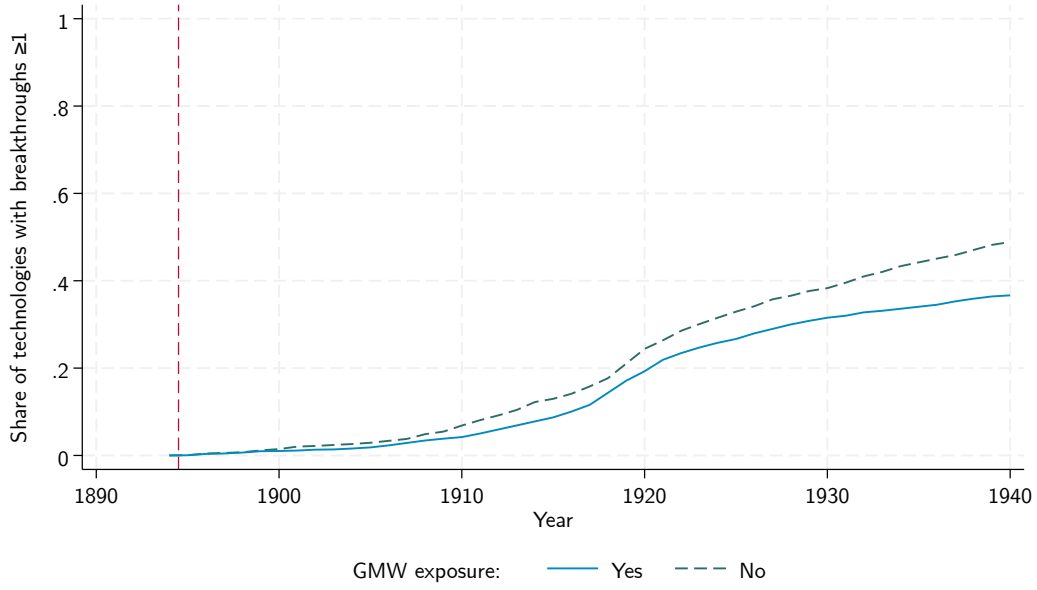
(b) Breakthroughs

Note: This figure shows innovation levels in established technologies (those with at least one patent before 1895) by GMW exposure status, between 1880-1940. Units of observation are clusters of CPC subgroups aggregated into 977 size-balanced technology domains. GMW exposure is defined as having any pre-1895 patents held by firms that subsequently participated in the Great Merger Wave. Breakthroughs are defined using the Kelly et al. (2021) measure. Both patents and breakthroughs are adjusted to account for multiple CPC classifications.

Figure 16: Development of Emerging Technologies by GMW Exposure Status



(a) Emerging technologies with at least one patent



(b) Emerging technologies with at least one breakthrough

Note: This figure shows the cumulative share of emerging technologies (CPC groups with no patents before 1895) that have emerged by GMW exposure status, between 1895–1940. Panel (a) shows the share with at least one patent; Panel (b) shows the share with at least one breakthrough. GMW exposure is defined by proximity in the CPC classification system: emerging technology groups inherit the exposure status of their broader subclass based on whether related, pre-existing technologies had GMW firm activity before 1895.

non-science-based technologies exhibit limited effects in established domains and moderately negative effects in emerging ones.

6.1 Innovation Responses in Established Technologies

How did exposure to the Great Merger Wave affect subsequent innovation outcomes in established technological domains? To answer this question, I examine whether technologies that had greater exposure to consolidating firms before 1895 experienced different innovation trajectories afterward. While the descriptive patterns in Figure 15 are suggestive, they may reflect systematic differences in the types or maturity of technologies where GMW firms were active, rather than causal effects of consolidation. To estimate causal effects, I employ a difference-in-differences design that compares exposed and unexposed technologies over time, within specific technological areas and vintage categories.

I focus the analysis on patent classes with at least one patent before 1895. To define a technological domain, I aggregate CPC subgroups into approximately 1,000 size-balanced technology domains. In particular, I use an agglomerative algorithm that groups nearby CPC subgroups within nested CPC levels to reach approximately equal pre-1895 sizes.[45] Next, I measure consolidation exposure at the technology level using the pre-1895 patents of firms that subsequently participated in the Great Merger Wave. Specifically, I define a technology as exposed to consolidation if GMW firms held any patents in that technology before 1895. See Appendix B for details on the data construction.

The empirical framework employs a difference-in-differences design comparing exposed to unexposed technologies over time. The static specification is:

$$y_{itcm} = \alpha_i + \delta_{tc} + \nu_{tm} + \beta_1(\text{During}_t \times \text{GMW}_i) + \beta_2(\text{Post}_t \times \text{GMW}_i) + \varepsilon_{itcm}, \tag{12}$$

where y_{itcm} represents outcomes for technology i during year t in CPC section c and vintage m ; α_i are technology fixed effects, δ_{tc} are year fixed effects by nine CPC sections, ν_{tm} are year fixed effects by three vintage groups,[46] GMW_i indicates binary exposure before 1895, During_t captures the merger wave period (1895–1904), and Post_t covers the subsequent period (1905-1940). The key outcomes are total patents and breakthrough innovations within each technology, as well as measures that exclude direct contributions from GMW firms themselves. I also examine the number of firms and inventors active

45. This approach addresses two key issues. First, CPC subgroups differ dramatically in pre-1895 size. When exposure is defined as having at least one GMW-related patent, larger groups would mechanically be more likely to meet that threshold, creating a spurious correlation between exposure and technology size. Grouping CPC subgroups into similarly sized domains addresses this issue. Second, size-balancing ensures comparable effects across technologies without requiring log transformations or normalizations that would be problematic given the prevalence of zeros in the data. Technologies are constructed by clustering CPC subgroups based on pre-1895 patent counts, while respecting the classification system's hierarchical structure. Results are robust to alternative numbers of clusters (e.g., $k = 750$ and $k = 1,500$), as shown in Appendix F.

46. To capture technological maturity (vintage), I define Earliest Year as the first year a patent is recorded in a given CPC group and then take the median Earliest Year across all groups in a technology. I then compute the terciles of the technology-level median Earliest Year to obtain my three vintages m.

in a given technology-year, and the share of patents assigned to firms (versus other assignees like independent inventors).

The identification strategy relies on the assumption that, conditional on the fixed effects, technologies exposed and unexposed to GMW firms would have followed parallel innovation trends absent the consolidation shock. This assumption is supported by the flat pre-trends already evident in Figure 15 and formally tested in the event study reported in Figure 17. Further, the parallel trend is plausible given that consolidation decisions were driven by factors unrelated to innovation potential—primarily market competition, overcapacity, and financial distress as discussed extensively in Section 2 and 4.1. The design effectively compares more-exposed to less-exposed technologies rather than truly unexposed ones, since after 1895 GMW firms also patented in established technologies where they had held no patents before the wave, so that even nominally unexposed domains received some indirect exposure. The estimated effects therefore represent the differential impact of greater consolidation exposure.

Panel A of Table 7 shows that technologies directly exposed to the Great Merger Wave experienced small and statistically insignificant increases in both total patents (1.14 additional patents per year) and breakthrough patents (0.37 additional breakthroughs per year) during the post-merger period. However, this aggregate picture masks striking differential effects across technological domains. The heterogeneity analysis in Panel B reveals that science-based technologies (CPC sections C, G and H)[47] experienced large increases in both patenting and breakthrough innovations, though only the latter effect is significant at the 5 percent level. These gains represent a significant acceleration in high-impact innovation within domains that would prove central to twentieth-century technological progress. In contrast, other technological areas showed small, insignificant and directionally negative effects.

Gains for science-based technologies appear to be driven more by the quality of innovation than by sheer volume. The estimated increase in breakthrough patents (3.47 per year) represents nearly half of the total increase in patents (7.39 per year), implying a breakthrough rate of roughly 47 percent. This is substantially higher than the baseline rate of 10 percent that would be expected if breakthroughs were distributed uniformly over time and across technologies. These effects persist when excluding direct contributions from GMW firms themselves (Table 7, columns 3-4). This persistence indicates that the estimated effects capture spillovers to other innovators rather than mechanical increases from the consolidating firms alone.

The event study in Figure 17 shows flat pre-trends for both science-based (blue) and non-science-based (green) technologies, supporting the identifying assumption.[48] Dy-

47. This definition is analogous to that in Section 4. CPC sections C, G and H encompass: chemistry, metallurgy; scientific instruments, computing; electronics, telecommunications. Science-based technologies represent about 13 percent of the total number of technologies in both exposed and unexposed domains.

48. The specification is the dynamic equivalent Equation 12 with year-specific coefficients:

$$y_{itcm} = \alpha_i + \delta_{tc} + \nu_{tm} + \sum_{s=1885, s\neq 1894}^{1940} \mathbb{1}[t = s] \cdot \text{GMW}_i \cdot \left(\beta_s^{\text{Sci}} \, \text{Sci}_c + \beta_s^{\text{Oth}} (1 - \text{Sci}_c)\right) + \varepsilon_{itcm}$$

, where Sci_c indicates science-based CPC sections and the paths β_s^{Sci} and β_s^{Oth} are plotted separately.

namic estimates also confirm differential effects over time. Post-merger, the figure reveals a clear divergence: science-based technologies experience large and positive effects on both patents and breakthroughs, while non-science-based technologies remain flat.

To shed more light on potential effects on the structure of the innovation ecosystem, I consider the number of active firms and inventors and the overall share of patents obtained by firms. Table F1 reveals that GMW-exposed technologies attracted significantly more firms (about 2.2) and inventors (3.2). The share of patents assigned to firms also increased substantially (2.6 percentage points), consistently with a shift toward more organized, corporate-led innovation. In science-based technologies specifically, this pattern intensified with larger magnitudes across the board.

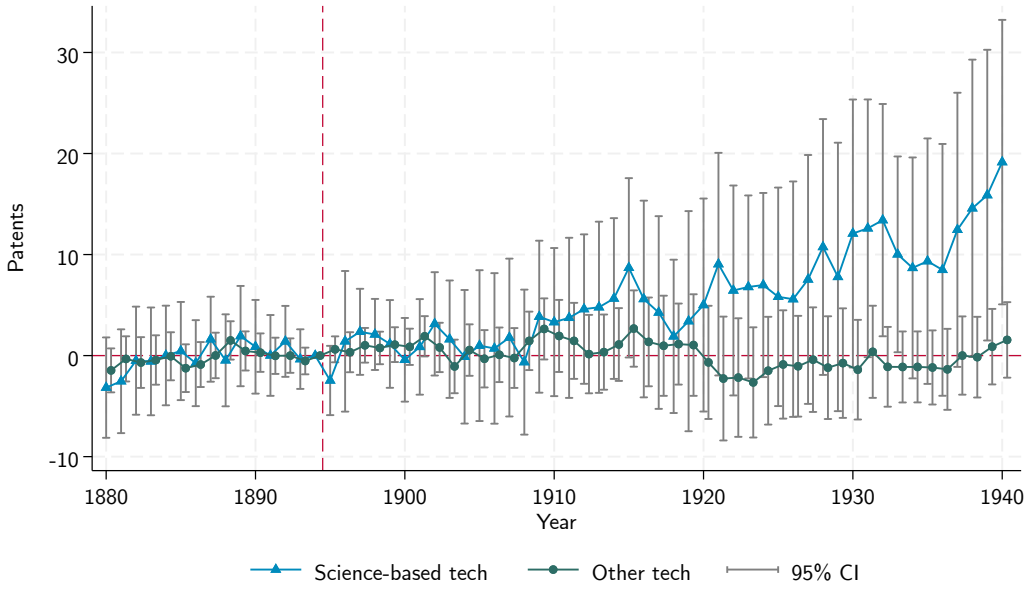
Robustness exercises are reported in Appendix F. In particular, I show that results are consistent when changing the number of technologies (clusters) and the set of controls included. I also show results by terciles of GMW exposure, which do not reveal significant differences.

Table 7: Effect of Consolidation Exposure on Technology-Level Innovation

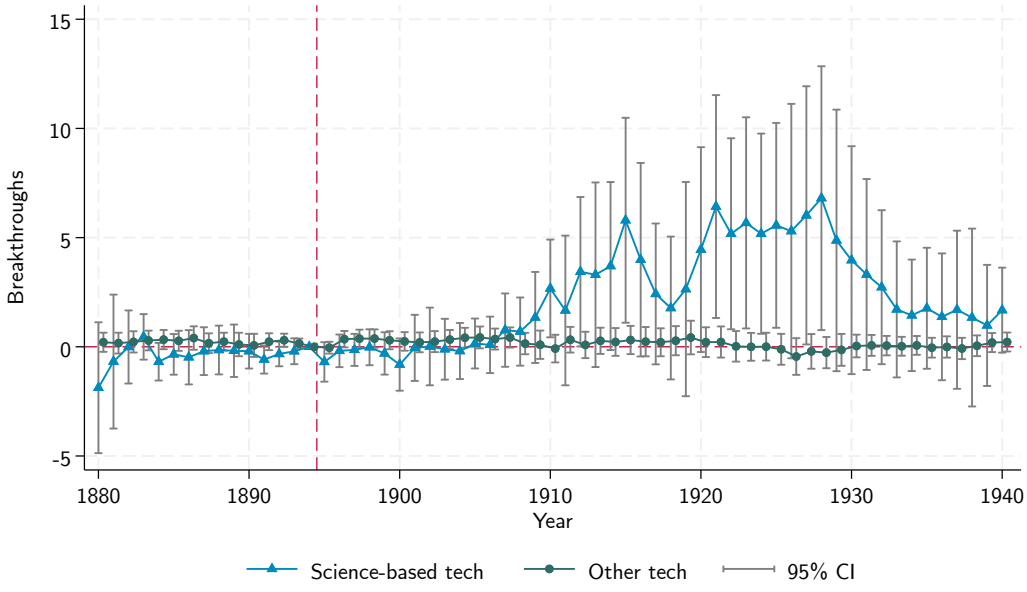
Outcome:	Patents	Breakthroughs	Patents	Breakthroughs
Panel A: Overall Effect				
Post × GMW	1.139	0.374	0.881	0.360
	(1.664)	(0.267)	(1.649)	(0.260)
	[-2.126, 4.403]	[-0.151, 0.899]	[-2.354, 4.116]	[-0.149, 0.869]
Panel B: Heterogeneity				
Post × GMW × Science-Based	7.388	3.472	7.110	3.254
	(4.740)	(1.563)	(4.683)	(1.517)
	[-1.914, 16.690]	[0.405, 6.539]	[-2.079, 16.300]	[0.277, 6.231]
Post × GMW × Other Tech	0.162	-0.111	-0.093	-0.093
	(1.746)	(0.175)	(1.731)	(0.171)
	[-3.264, 3.587]	[-0.455, 0.234]	[-3.490, 3.304]	[-0.428, 0.242]
Excluding GMW firms	N	N	Y	Y
Year x CPC Section FE	Y	Y	Y	Y
Year x Vintage Tercile FE	Y	Y	Y	Y
Technologies	977	977	977	977
N	59,597	59,597	59,597	59,597

Note: This table reports difference-in-differences estimates of the effect of consolidation exposure on technology-level innovation outcomes after 1905 (β_2 in Equation 12). The unit of observation is a technology-year, where technologies are clusters of CPC subgroups aggregated into 977 size-balanced domains. GMW exposure is defined as having any pre-1895 patents held by firms that subsequently participated in the Great Merger Wave. Panel A shows overall effects; Panel B shows heterogeneity by science-based (CPC sections C, G, H) versus other technologies. Columns 3-4 exclude direct contributions from GMW firms themselves. Standard errors clustered at the technology level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Figure 17: Effect of Consolidation Exposure on Technology-Level Innovation—Dynamic Effects



(a) Event Study for Patents



(b) Event Study for Breakthroughs

Note: This figure presents dynamic event study estimates showing the technology-level effect of consolidation exposure on innovation outcomes. Blue triangles show effects for science-based technologies; green circles show effects for other technologies. Panel (a) shows effects on patents; Panel (b) shows effects on breakthroughs. Standard errors are clustered at the technology level. 95 percent confidence intervals are shown in gray.

6.2 Innovation Responses in Emerging Technologies

How did exposure to the Great Merger Wave affect the development of emerging technological domains? Newly emerging technologies offer a distinct margin along which consolidation could shape the trajectory of U.S. innovation—by altering the incentives or conditions under which new related technologies are discovered. Dominant firms might delay and discourage entry into nascent fields that could disrupt their existing products, or conversely, their R&D investments might accelerate the discovery of related breakthrough opportunities.

To investigate this question, I conduct a survival analysis focusing on CPC groups that recorded no patents before 1895. The outcome of interest is the timing of their first appearance in the patent record, measured either by the first patent of any kind or, in a separate estimation, by the first breakthrough patent. This framework assumes that these technologies would eventually emerge in the patent record, but that GMW exposure might accelerate or delay their development. The estimates therefore speak to the timing of emergence among classified domains, not to whether some potential technologies never appeared at all.

To construct a measure of exposure for these emerging domains, I exploit the hierarchical structure of the CPC system. Its structure is such that: Group $\subset$ Subclass $\subset$ Class $\subset$ Section. Thus, within each CPC subclass, I flag exposure if any of its CPC groups active before 1895 had at least one patent assigned to a GMW firm. All emerging groups born into that subclass inherit its exposure status. This approach reflects the idea that technologies within the same subclass are closely related in terms of knowledge inputs, scientific foundations, or end-use applications. If consolidation influenced innovation dynamics in exposed domains, these effects could spill over to related technologies emerging nearby.

The empirical framework employs a Cox proportional hazards model to estimate the likelihood that an emerging technology becomes active in a given year. The specification includes CPC section-specific non-parametric baseline hazards, fixed effects for CPC class, and controls for vintage:

$$h_{gcs}(t|X_{gcs}) = h_s(t)\exp(\beta \cdot \text{GMW Exposed}_g + \delta \cdot \text{EarliestYear}_g + \theta_c) \tag{13}$$

where g indexes groups, c classes, and s sections; $h_{gcs}(t|X_{gcs})$ is group g's hazard of activating for the first time at time t, given observables. The variable GMW_g indicates whether the group belongs to an exposed subclass, while EarliestYear_g captures the year in which any pre-1895 patent first appeared in g's subclass.[49] Standard errors are clustered at the subclass level to account for shared treatment and unobserved correlations. Because class fixed effects absorb level differences in first activation rates across broader technological families, the coefficient β is identified from within-class comparisons between groups belong to exposed and unexposed subclasses. I estimate this model separately for first patent and first breakthrough as the defining event.

49. Subclasses with no pre-1895 patents have no defined EarliestYear_g and contain no exposed group; they enter only as unexposed controls within classes that hold some pre-1895 variation, since identification is within-class. The group and subclass counts reported in the results tables refer to this estimation risk

Table 8: Effect of Consolidation on Emerging Technologies

Outcome:	First Patent	First Breakthrough
Panel A: Overall Effect		
GMW	1.016	0.865
	(0.063)	(0.090)
	[0.900, 1.148]	[0.704, 1.062]
Panel B: Heterogeneity		
GMW × Science-Based	1.029	0.965
	(0.097)	(0.128)
	[0.856, 1.237]	[0.744, 1.252]
GMW × Other Tech	1.005	0.682
	(0.080)	(0.099)
	[0.859, 1.175]	[0.513, 0.906]
Controls	Y	Y
Groups	2,898	2,898
Subclasses	474	474
N	81,693	107,085

Note: This table reports hazard ratios from Cox proportional hazards models estimating the effect of consolidation exposure on the development of emerging technologies (Equation 13). The sample includes CPC groups with no patents before 1895. GMW exposure for emerging groups is inherited from their CPC subclass: a group is exposed if any group within its subclass had at least one pre-1895 patent assigned to a GMW firm. Panel A shows overall effects; Panel B shows heterogeneity by science-based (CPC sections C, G, H) versus other technologies. Hazard ratios above one indicate accelerated emergence; ratios below one indicate delayed emergence. Standard errors clustered at the subclass level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

The baseline survival analysis reveals no significant effects of GMW exposure on the development of emerging technologies. Table 8 presents results from the Cox proportional hazards analysis where estimates for β are reported as hazard ratios: a value above one indicates that GMW-exposure increased the likelihood of a emerging technology entering the innovation record in a given year, while a value below one implies a slowdown. Panel A shows that, on average, GMW exposure had no statistically significant effect on the timing of emerging technology development. While the effect on first patent is very close to one (1.02), the hazard ratio for first breakthrough (0.87) implies

set.

a rather sizable slowdown of about -13 percent, though not significant at the 5 percent level.

However, as with established technologies, average effects mask meaningful heterogeneity across types of technologies.[50] Panel B in Table 8 distinguishes between science-based and non-science-based technological domains. In science-based technologies, GMW exposure has insignificant hazard ratios are near one for both patents (1.03) and breakthroughs (0.97). By contrast, non-science-based technologies show delayed breakthrough development following GMW exposure. The estimated hazard ratio for first breakthroughs in these domains is 0.68, implying a 32 percent reduction in the annual likelihood of achieving a breakthrough. The effect on first patents null.

Appendix F reports additional results and robustness checks. In particular, I restrict to exposed CPC groups and explore whether the intensity of GMW exposure (rather than the fact of being exposed) affects emerging technology development rates. I find insignificant effects, with point estimates suggesting a mild acceleration in more science-based technological domains. I also show that the main results are robust to changing the controls.

6.3 Overall Net Effect of Great Merger Wave Exposure

To assess the Great Merger Wave's overall impact on American technological development, I translate the preceding empirical estimates into aggregate breakthrough counts. These back-of-the-envelope calculations quantify the total number of breakthrough innovations attributable to consolidation exposure during 1905-1940. For established technologies, I multiply the difference-in-differences coefficients and the number of exposed technology-year observations. For emerging technologies, the survival analysis represents the most appropriate empirical approach given the absence of a meaningful pre-period and the focus on timing of technological emergence. However, hazard ratios do not translate into absolute breakthrough counts. I therefore estimate a linear specification that captures the average difference in breakthrough production between exposed and unexposed emerging technology groups, controlling for the same sources of variation as the Cox model. While this linear approach sacrifices some of the econometric appeal of the hazard framework, it provides a reasonable approximation of the aggregate effects by measuring observed differences in breakthrough activity across groups with different consolidation exposure.

These aggregate calculations, detailed in Appendix F.1 and summarized in Table F7, suggest that consolidation exposure increased breakthrough innovation by 13.2 percent above counterfactual levels between 1905 and 1940. However, this aggregate effect masks pronounced heterogeneity across technological domains. Among science-based technologies, GMW exposure generated a 30.3 percent increase in breakthrough innovations relative to what would have occurred without consolidation. Conversely, non-science-based technologies experienced a net reduction in breakthrough activity, declining by 6.7 percent below their counterfactual levels. These estimates indicate that while the

50. Science-based groups represented about 57 percent of the total number of groups in unexposed domains, and about 39 percent in exposed domains.

GMW generated substantial innovation gains in aggregate, these benefits were highly concentrated in technological domains that required the organizational capabilities and R&D infrastructure that large-scale consolidation made economically viable.

7 Discussion

How do this paper's findings align with economic theory and historical accounts of American innovation? This section interprets the empirical results and discusses their broader implications.

The substantial positive effects documented in Section 4 indicate that, at the level of consolidated firms, the forces favoring innovation by large enterprises dominated on average. The replacement effect of Arrow (1962) predicts that a monopolist has weaker incentives to invest in a given innovation than a competitive firm would, because new technologies cannibalize existing rents; yet two channels can generate increases in innovation at the firm level. First, larger firms can benefit from scale in R&D—absorbing fixed costs, attracting superior talent, and undertaking riskier long-term projects (Schumpeter 1942; Rosenberg 1990; Atkinson and Lind 2019). Second, greater market concentration can improve the appropriation of returns from innovation (Schumpeter 1942; Spulber 2013).

Moreover, the evidence suggests that GMW firms' patenting gains reflect genuine invention rather than rent-seeking, on average. Some firms may have pursued "preemptive" or "defensive" patenting to deter rivals (Gilbert and Newbery 1982; Igami 2017; Akcigit and Ates 2023), and historical accounts raise this concern for several leading corporations (Noble 1979; Reich 1985; Lamoreaux 2000). While the analysis cannot directly observe intent, increases extend to breakthrough inventions—highly original, influential discoveries that are unlikely to be primarily defensive. In addition, the gains are concentrated among lab-owning firms and patents filed near R&D facilities, consistent with substantial resource commitments to systematic experimentation and research.

The evidence is consistent with a key causal chain linking "bigness"—large firm size and market dominance—to innovation running through the reorganization of inventive activity within laboratories dedicated to testing, experimentation, and discovery. The joint increase in scale and market dominance enabled GMW firms to bear the fixed costs and organizational investments of R&D laboratories, which are in turn associated with higher inventive productivity. Section 5 shows that (i) the observed productivity premium for lab-owning firms is not explained by sorting, size, or broad technological area; (ii) firm productivity improves after gaining a lab; and (iii) inventors become more productive upon joining lab-owning firms. The convergence of these patterns supports a genuine productivity effect of laboratory operation.

These findings add quantitative weight to long-standing historical narratives on the transformative role of industrial laboratories. Hounshell (1996) and Mowery and Rosenberg (1998) document the institutionalization of corporate research and its correlation with the rise of systematic, team-based R&D. Hounshell and Smith (1988), Jenkins (1975), and Wise (1985) highlight how labs at DuPont, Eastman Kodak, and General Electric allowed large corporations to internalize innovation and manage disruptive change. More

recently, Gertner (2013) chronicles Bell Labs' extraordinary success.

Even if consolidating firms became more inventive, increased market dominance could have deterred innovation elsewhere—by raising entry barriers, reducing competition, or crowding out smaller innovators. Section 6 shows a substantial net positive effect on U.S. innovation, but an uneven one: in science-based technologies (chemistry, computing, telecommunications), consolidation is associated with increases in breakthrough innovations, whereas non-science-based domains exhibit negative effects.

The positive effects in science-based fields can be rationalized by the dynamics of fast-evolving scientific frontiers. In domains like chemistry and electronics, rapid scientific progress created opportunities for sustained leadership (Noble 1979; Hounshell and Smith 1988). This environment resembles the "neck-and-neck" technological competition of Aghion et al. (2005), in which similarly positioned rivals race to innovate. While consolidation gave GMW firms the scale and finance for systematic R&D, they still faced pressure from independent inventors and new entrants; fast-moving frontiers made lasting advantages harder to entrench (Hounshell and Smith 1988; Mowery and Rosenberg 1998).

Spillovers are also consistent with consolidated firms encouraging broader inventive effort. Rising Big Business with a greater ability to acquire or license strategic assets may raise external researchers' expected returns (Phillips and Zhdanov 2013). The estimated positive spillovers align with evidence that knowledge and productivity externalities from R&D can dominate business-stealing effects, especially in knowledge-intensive sectors (Bloom, Schankerman, and Van Reenen 2013).

By contrast, non-science-based technologies display patterns consistent with reduced innovation under diminished competition. In more stable technological environments that rely on incremental improvement, GMW firm-level patenting increases are consistent with a reallocation rather than a net expansion of innovative activity. "Preemptive" or "defensive strategies" may be more effective here, reducing both entry by new innovators and patenting by non-GMW incumbents. These domains may feature more technological niches with limited knowledge interconnections, where even large R&D performers generate fewer spillovers and business-stealing effects dominate (Bloom, Schankerman, and Van Reenen 2013).

Arguably, Big Business and its R&D labs filled an institutional gap in early twentieth-century U.S. innovation. Before World War II, the United States had virtually no federal R&D funding (Gross and Sampat 2023; Gruber and Johnson 2019), and universities often lagged European counterparts in training applied scientists and engineers (Graham and Diamond 1997). In this vacuum, corporations built internal research capacity to substitute for weak public science (Arora et al. 2024). Especially in fast-moving scientific fields, large firms became not just users but contributors to science (Hounshell and Smith 1988; Senecal 1980). The large lab productivity premium and the positive spillovers from consolidation documented in science-based sectors suggest that corporate R&D delivered particularly high returns in this environment.

8 Conclusion

Did the rise of Big Business drive American innovation in the first half of the twentieth century? This paper exploits the Great Merger Wave—the most sweeping consolidation episode in U.S. history—to provide new evidence on this longstanding question. Between 1895 and 1904, distress from the 1893 Depression, court rulings inadvertently incentivizing consolidation, and Wall Street activism combined to radically transform American industry. Thousands of medium and large firms disappeared into consolidations that aimed to control supply and prices, not to pursue technological synergies. Using newly constructed data linking corporate structure, innovation outcomes, and individual inventors, I trace how this sudden increase in firm size and market dominance affected American innovation through World War II.

The large, dominant enterprises created during GMW consolidations significantly expanded their innovative output. Among firms already patenting before 1895, annual patenting rose by about six patents and 0.56 breakthrough patents—a roughly four-fold and six-fold increase, respectively. Firms with no patents before 1895 were more likely to begin patenting, while announced consolidations that fell through—paper mergers—showed no comparable gains. Enterprises on both margins showed significantly higher probabilities of establishing R&D labs, an organizational innovation that greatly improved firms' innovative capacity. These firm-level gains spilled over to the wider innovation system, on net raising breakthrough innovations by 13 percent, and 30 percent in science-based technologies closer to the frontier.

Thus, the rise of Big Business during the Great Merger Wave played a pivotal role in shaping the trajectory of American innovation. The benefits were concentrated in science-based technologies where systematic R&D mattered most, while non-science-based domains saw some innovation slowdowns. Corporate laboratories emerged not merely as a byproduct of firm bigness, but as the central organizational mechanism enabling top firms' outsized innovative performance before World War II. In an era of weak public science institutions, Big Business became a central engine of American technological progress.

References

Abadie, Alberto. 2021. "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects." *Journal of Economic Literature* 59, no. 2 (June): 391–425.

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (2): 251–333.

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik. 2022. "Radical and Incremental Innovation: The Roles of Firms, Managers, and Innovators." *American Economic Journal: Macroeconomics* 14, no. 3 (July): 199–249.

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. 2014. "Chapter 1 - What Do We Learn From Schumpeterian Growth Theory?" In *Handbook of Economic Growth,* edited by Philippe Aghion and Steven N. Durlauf, 2:515–563. Handbook of Economic Growth. Elsevier, January.

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

Ahumada, Alonso. 2025. *Can Protective Taris Induce Industrial Consolidations? Theory and Evidence from the Great Merger Movement.*

Akcigit, Ufuk, and Sina T. Ates. 2023. "What Happened to US Business Dynamism?" *Journal of Political Economy* 131, no. 8 (August): 2059–2124.

Akcigit, Ufuk, John Grigsby, and Tom Nicholas. 2017. *The Rise of American Ingenuity: Innovation and Inventors of the Golden Age.* Working Paper, January.

Akcigit, Ufuk, John Grigsby, Tom Nicholas, and Stefanie Stantcheva. 2022. "Taxation and Innovation in the Twentieth Century*." *The Quarterly Journal of Economics* 137, no. 1 (February): 329–385.

Andrews, Michael J., and Alexander Whalley. 2022. "150 years of the geography of innovation." *Regional Science and Urban Economics,* Urban Economics and History, 94 (May): 103627.

Argente, David, Salome Baslandze, Douglas Hanley, and Sara Moreira. 2020. "Patents to Products: Product Innovation and Firm Dynamics." *FRB Atlanta Working Paper 2020-4.*

Argyres, Nicholas S., and Brian S. Silverman. 2004. "R&D, organization structure, and the development of corporate technological knowledge." *Strategic Management Journal* 25 (8-9): 929–958.

Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. 2021. "Synthetic Difference-in-Differences." *American Economic Review* 111, no. 12 (December): 4088–4118.

Arora, Ashish, Sharon Belenzon, Konstantin Kosenko, Jungkyu Suh, and Yishay Yafeh. 2024. "The Rise of Scientific Research in Corporate America." *Organization Science* (December).

Arora, Ashish, Sharon Belenzon, and Andrea Patacconi. 2018. "The decline of science in corporate R&D." *Strategic Management Journal* 39 (1): 3–32.

Arora, Ashish, Sharon Belenzon, and Luis A. Rios. 2011. *The Organization of R&D in American Corporations: The Determinants and Consequences of Decentralization.* Working Paper, May.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella. 2001. *Markets for Technology: The Economics of Innovation and Corporate Strategy.* The MIT Press, December.

Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors,* 609–626. Princeton University Press.

Atkinson, Robert D., and Michael Lind. 2019. *Big Is Beautiful: Debunking the Myth of Small Business.* Illustrated edition. The MIT Press, February.

Babina, Tania, Asaf Bernstein, and Filippo Mezzanotti. 2023. "Financial Disruptions and the Organization of Innovation: Evidence from the Great Depression." *The Review of Financial Studies* (April): hhad023.

Baskin, Jonathan Barron, and Jr Miranti Paul J. 1997. *A History of Corporate Finance.* Cambridge: Cambridge University Press.

Benmelech, Efraim, and Carola Frydman. 2015. "Military CEOs." *Journal of Financial Economics,* NBER Conference on the Causes and Consequences of Corporate Culture, 117, no. 1 (July): 43–59.

Berkes, Enrico. 2018. *Comprehensive Universe of U.S. Patents (CUSP): Data and Facts.*

Bertrand, Marianne, and Antoinette Schoar. 2003. "Managing with Style: The Effect of Managers on Firm Policies." *The Quarterly Journal of Economics* 118 (4): 1169–1208.

Bhaskarabhatla, Ajay, Luis Cabral, Deepak Hegde, and Thomas Peeters. 2021. "Are Inventors or Firms the Engines of Innovation?" *Management Science* 67, no. 6 (June): 3899–3920.

Bikard, Michaël, and Matt Marx. 2020. "Bridging Academia and Industry: How Geographic Hubs Connect University Science and Corporate Technology." *Management Science* 66, no. 8 (August): 3425–3443.

Bittlingmayer, George. 1985. "Did Antitrust Policy Cause the Great Merger Wave?" *The Journal of Law and Economics* 28, no. 1 (April): 77–118.

Blonigen, Bruce A, and Justin R Pierce. 2016. *Evidence for the Effects of Mergers on Market Power and Efficiency,* 16-43.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4): 1347–1393.

Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler. 2023. "How Much Should We Trust Estimates of Firm Effects and Worker Sorting?" *Journal of Labor Economics* 41, no. 2 (April): 291–322.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A Distributional Framework for Matched Employer Employee Data." *Econometrica* 87 (3): 699–739.

Braguinsky, Serguey, Joonkyu Choi, Yuheng Ding, Karam Jo, and Seula Kim. 2023. *Mega Firms and Recent Trends in the U.S. Innovation: Empirical Evidence from the U.S. Patent Data.* Working Paper 31460. National Bureau of Economic Research, July.

Bryan, Kevin A., and Heidi L. Williams. 2021. "Innovation: market failures and public policies." In *Handbook of Industrial Organization,* edited by Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, 5:281–388. Handbook of Industrial Organization, Volume 5. Elsevier, January.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics,* Themed Issue: Treatment Effect 1, 225, no. 2 (December): 200–230.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality*." *The Quarterly Journal of Economics* 128, no. 3 (August): 967–1015.

Chandler, Alfred D. 1959. "The Beginnings of "Big Business" in American Industry." *The Business History Review* 33 (1): 1–31.

———. 1977. *The Visible Hand: The Managerial Revolution in American Business.* Harvard University Press.

———. 1990. *Scale and Scope: The Dynamics of Industrial Capitalism.* Harvard University Press.

Cohen, Wesley M. 2010. "Chapter 4 - Fifty Years of Empirical Studies of Innovative Activity and Performance." In *Handbook of the Economics of Innovation,* edited by Bronwyn H. Hall and Nathan Rosenberg, 1:129–213. Handbook of The Economics of Innovation, Vol. 1. North-Holland, January.

Cohen, Wesley M., and Daniel A. Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." *Administrative Science Quarterly* 35 (1): 128–152.

Collins, Norman R., and Lee E. Preston. 1961. "The Size Structure of the Largest Industrial Firms, 1909-1958." *The American Economic Review* 51 (5): 986–1011.

Comanor, William S., and F. M. Scherer. 2013. "Mergers and innovation in the pharmaceutical industry." *Journal of Health Economics* 32, no. 1 (January): 106–113.

Conant, Luther. 1901. "Industrial Consolidations in the United States." *Publications of the American Statistical Association* 7 (53): 1–20.

Cunningham, Colleen, Florian Ederer, and Song Ma. 2021. "Killer Acquisitions." *Journal of Political Economy* 129, no. 3 (March): 649–702.

DeLong, J. Bradford. 2022. *Slouching Towards Utopia: An Economic History of the Twentieth Century.* 1st edition. New York: Basic Books, September.

Donaldson, Dave, and Richard Hornbeck. 2016. "Railroads and American Economic Growth: A "Market Access" Approach *." *The Quarterly Journal of Economics* 131, no. 2 (May): 799–858.

Du Boff, Richard B., and Edward S. Herman. 1989. "The Promotional-Financial Dynamic of Merger Movements: A Historical Perspective." *Journal of Economic Issues* 23, no. 1 (March): 107–133.

Ellison, Glenn. 1994. "Theories of Cartel Stability and the Joint Executive Committee." *The RAND Journal of Economics* 25 (1): 37–57.

Federico, Giulio, Fiona Scott Morton, and Carl Shapiro. 2020. "Antitrust and Innovation: Welcoming and Protecting Disruption." *Innovation Policy and the Economy* 20 (December): 125–190.

Fellegi, Ivan P., and Alan B. Sunter. 1969. "A Theory for Record Linkage." *Journal of the American Statistical Association* 64, no. 328 (December): 1183–1210.

Fernández-Villaverde, Jesús, Yang Yu, and Francesco Zanetti. 2025. *Defensive Hiring and Creative Destruction.* Working Paper, March.

Field, Alexander J. 2003. "The Most Technologically Progressive Decade of the Century." *American Economic Review* 93, no. 4 (September): 1399–1413.

———. 2011. *A Great Leap Forward: 1930s Depression and U.S. Economic Growth.* Yale University Press.

———. 2013. "Economic growth and recovery in the United States: 1919–1941." In *The Great Depression of the 1930s: Lessons for Today,* edited by Nicholas Crafts and Peter Fearon, 358–94. Oxford: Oxford University Press.

Fons-Rosen, Christian, Pau Roldan-Blanco, and Tom Schmitz. 2021. *The Effects of Startup Acquisitions on Innovation and Economic Growth.* SSRN Scholarly Paper. Rochester, NY, February.

Furman, Jeffrey L., and Megan J. MacGarvie. 2007. "Academic science and the birth of industrial research laboratories in the U.S. pharmaceutical industry." *Journal of Economic Behavior & Organization,* Academic Science and Entrepreneurship: Dual engines of growth, 63, no. 4 (August): 756–776.

Gertner, Jon. 2013. *The Idea Factory: Bell Labs and the Great Age of American Innovation.* Reprint edition. Penguin Books, February.

Gilbert, Richard J., and David M. G. Newbery. 1982. "Preemptive Patenting and the Persistence of Monopoly." *The American Economic Review* 72 (3): 514–526.

Goettler, Ronald L., and Brett R. Gordon. 2011. "Does AMD Spur Intel to Innovate More?" *Journal of Political Economy* 119 (6): 1141–1200.

Golbe, Devra L., and Lawrence J. White. 1988. "A Time-Series Analysis of Mergers and Acquisitions in the U.S. Economy." *NBER Chapters,* 265–310.

Gordon, Robert J. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War.* First Edition. Princeton: Princeton University Press, January.

Graham, Hugh Davis, and Nancy Diamond. 1997. *The Rise of American Research Universities: Elites and Challengers in the Postwar Era.* Baltimore, Md: Johns Hopkins University Press.

Griliches, Zvi. 1992. "The Search for R&D Spillovers." *The Scandinavian Journal of Economics* 94:S29–S47.

Gross, Daniel P., and Bhaven Sampat. 2023. *America, Jump-started: World War II R&D and the Takeoff of the U.S. Innovation System.* Working paper.

Gruber, Jonathan, and Simon Johnson. 2019. *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream.* New York: PublicAffairs, April.

Hannah, Leslie. 1974. "Mergers in British Manufacturing Industry, 1880-1918." *Oxford Economic Papers* 26 (1): 1–20.

Hartog, Matte, Andres Gomez-Lievano, Ricardo Hausmann, and Frank Neffke. 2024. "Inventing modern invention: the professionalization of technological progress in the US." *Papers in Evolutionary Economic Geography (PEEG)* (April).

Haucap, Justus, Alexander Rasch, and Joel Stiebale. 2019. "How mergers affect innovation: Theory and evidence." *International Journal of Industrial Organization* 63 (March): 283–325.

Henderson, Rebecca, and Iain Cockburn. 1996. "Scale, Scope, and Spillovers: The Determinants of Research Productivity in Drug Discovery." *The RAND Journal of Economics* 27 (1): 32–59.

Henderson, Rebecca M., and Kim B. Clark. 1990. "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms." *Administrative Science Quarterly* 35 (1): 9–30.

Hounshell, David. 1996. "The evolution of industrial research." In *Engines of Innovation: U.S. Industrial Research at the End of an Era,* edited by Richard Rosenbloom and William Spencer. Boston, MA: Harvard Business School Press.

Hounshell, David A., and John K. Jr Smith. 1988. *Science and Corporate Strategy: Du Pont R&D, 1902–1980.* Cambridge: Cambridge University Press.

Huang, Angela, Ruoke Yang, and Mengxin Zhao. 2025. *Analysis of Merger & Acquisition Activity.* White Paper.

Hughes, Thomas P. 2004. *American Genesis: A Century of Invention and Technological Enthusiasm, 1870-1970.* 2nd edition. Chicago: University of Chicago Press, June.

Igami, Mitsuru. 2017. "Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998." *Journal of Political Economy* 125, no. 3 (June): 798–847.

Igami, Mitsuru, and Kosuke Uetake. 2020. "Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996–2016." *The Review of Economic Studies* 87, no. 6 (November): 2672–2702.

Jenkins, Reese. 1975. *Images and Enterprise: Technology and the American Photographic Industry, 1839 to 1925.* Johns Hopkins University Press.

Jewkes, John, David Sawers, and Richard Stillerman. 1958. *The Sources of Invention.* London: Macmillan.

Jones, Benjamin F., and Lawrence H. Summers. 2020. "A Calculation of the Social Returns to Innovation." In *Innovation and Public Policy*, 13–59. University of Chicago Press, December.

Kandel, Eugene, Konstantin Kosenko, Randall Morck, and Yishay Yafeh. 2019. "The great pyramids of America: A revised history of U.S. business groups, corporate ownership, and regulation, 1926–1950." *Strategic Management Journal* 40 (5): 781–808.

Kang, Hyo. 2023. *How Does Price Competition Affect Innovation? Evidence from US Antitrust Cases.* SSRN Scholarly Paper. Rochester, NY, January.

Kantor, Shawn, and Alexander T Whalley. 2023. *Moonshot: Public R&D and Growth.* Working Paper 31471. National Bureau of Economic Research, July.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run." *American Economic Review: Insights* 3, no. 3 (September): 303–320.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. "Leave-Out Estimation of Variance Components." *Econometrica* 88 (5): 1859–1898.

Kling, Gerhard. 2006. "The long-term impact of mergers and the emergence of a merger wave in pre-World-War I Germany." *Explorations in Economic History* 43, no. 4 (October): 667–688.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth*." *The Quarterly Journal of Economics* 132, no. 2 (May): 665–712.

Kramer, Helen M. 1964. "Harvesters and High Finance: Formation of the International Harvester Company." *The Business History Review* 38 (3): 283–301.

Kwon, Spencer Y., Yueran Ma, and Kaspar Zimmermann. 2024. "100 Years of Rising Corporate Concentration." *American Economic Review* 114, no. 7 (July): 2111–2140.

Lamoreaux, Naomi R. 1985. *The Great Merger Movement in American Business, 1895–1904.* Cambridge: Cambridge University Press.

———. 2000. "Entrepreneurship, Business Organization, and Economic Concentration." In *The Cambridge Economic History of the United States: Volume 2: The Long Nineteenth Century,* edited by Robert E. Gallman and Stanley L. Engerman, 2:403–434. Cambridge Economic History of the United States. Cambridge: Cambridge University Press.

———. 2019. "The Problem of Bigness: From Standard Oil to Google." *Journal of Economic Perspectives* 33, no. 3 (August): 94–117.

Lamoreaux, Naomi R., and Kenneth L. Sokoloff. 1999. "Inventors, Firms, and the Market for Technology in the Late Nineteenth and Early Twentieth Centuries." In *Learning by Doing in Markets, Firms, and Countries,* 19–60. University of Chicago Press, January.

———. 2001. "Market Trade in Patents and the Rise of a Class of Specialized Inventors in the 19th-Century United States." *American Economic Review* 91, no. 2 (May): 39–44.

Lamoreaux, Naomi R., Kenneth L. Sokoloff, and Dhanoos Sutthiphisal. 2011. "The Reorganization of Inventive Activity in the United States during the Early Twentieth Century." In *Understanding Long-Run Economic Growth: Geography, Institutions, and the Knowledge Economy,* edited by Dora L. Costa and Naomi R. Lamoreaux, 0. University of Chicago Press, October.

Lampe, Ryan, and Petra Moser. 2016. "Patent Pools, Competition, and Innovation—Evidence from 20 US Industries under the New Deal." *The Journal of Law, Economics, and Organization* 32, no. 1 (March): 1–36.

Langlois, Richard N. 2023. *The Corporation and the Twentieth Century: The History of American Business Enterprise.* Princeton University Press, June.

Livermore, Shaw. 1935. "The Success of Industrial Mergers." *The Quarterly Journal of Economics* 50 (1): 68–96.

Ma, Yueran, Benjamin Pugsley, Haomin Qin, and Kaspar Zimmermann. 2025. *Superstar Firms through the Generations.* Working Paper, August.

MacGarvie, Megan, and Jeffrey Furman. 2005. *Early Academic Science and the Birth of Industrial Research Laboratories in the U.S. Pharmaceutical Industry.* Technical report w11470. Cambridge, MA: National Bureau of Economic Research, July.

Markham, Jesse W. 1955. "Survey of the Evidence and Findings on Mergers." In *Business Concentration and Price Policy,* 141–212. Princeton University Press.

Moody, John. 1904. *The Truth About the Trusts: A Description and Analysis of the American Trust Movement.* New York: Moody Publishing Company.

Moser, Petra, and Shmuel San. 2020. *Immigration, Science, and Invention. Lessons from the Quota Acts.* SSRN Scholarly Paper. Rochester, NY, March.

Moser, Petra, and Alessandra Voena. 2012. "Compulsory Licensing: Evidence from the Trading with the Enemy Act." *American Economic Review* 102, no. 1 (February): 396–427.

Moser, Petra, Alessandra Voena, and Fabian Waldinger. 2014. "German Jewish Émigrés and US Invention." *American Economic Review* 104, no. 10 (October): 3222–3255.

Mowery, David C. 1983. "Industrial Research and Firm Size, Survival, and Growth in American Manufacturing, 1921-1946: An Assessment." *The Journal of Economic History* 43 (4): 953–980.

———. 1984. "Firm Structure, Government Policy, and the Organization of Industrial Research: Great Britain and the United States, 1900-1950." *The Business History Review* 58 (4): 504–531.

———. 1990. "The Development of Industrial Research in U.S. Manufacturing." *American Economic Review* 80 (2): 345–349.

Mowery, David C., and Nathan Rosenberg, eds. 1989. "The beginnings of the commercial exploitation of science by U.S. industry." In *Technology and the Pursuit of Economic Growth,* 35–58. Cambridge: Cambridge University Press.

———. 1998. *Paths of Innovation: Technological Change in 20th-Century America.* Cambridge: Cambridge University Press.

Nanda, Ramana, and Tom Nicholas. 2014. "Did bank distress stifle innovation during the Great Depression?" *Journal of Financial Economics* 114, no. 2 (November): 273–292.

Navin, Thomas R. 1970. "The 500 Largest American Industrials in 1917." *The Business History Review* 44 (3): 360–386.

Nelson, Ralph L. 1959. *Merger Movements in American Industry, 1895-1956.* Princeton University Press.

Nicholas, Tom. 2003. "Why Schumpeter Was Right: Innovation, Market Power, and Creative Destruction in 1920s America." *The Journal of Economic History* 63 (4): 1023–1058.

———. 2009. "Spatial diversity in invention: evidence from the early R&D labs." *Journal of Economic Geography* 9, no. 1 (January): 1–31.

———. 2010. "The Role of Independent Invention in U.S. Technological Development, 1880–1930." *The Journal of Economic History* 70, no. 1 (March): 57–82.

———. 2011. "Did R&D Firms Used to Patent? Evidence from the First Innovation Surveys." *The Journal of Economic History* 71 (4): 1032–1059.

Noble, David F. 1979. *America by Design: Science, Technology, and the Rise of Corporate Capitalism.* Oxford University Press.

Parsons, Donald O., and Edward John Ray. 1975. "The United States Steel Consolidation: The Creation of Market Control." *The Journal of Law and Economics* 18, no. 1 (April): 181–219.

Phillips, Gordon M., and Alexei Zhdanov. 2013. "R&D and the Incentives from Merger and Acquisition Activity." *The Review of Financial Studies* 26, no. 1 (January): 34–78.

Porter, Glenn. 2006. *The rise of big business, 1860-1920.* 3rd ed. The American history series. Wheeling, Ill: Harlan Davidson, Inc.

Reich, Leonard S. 1977. "Research, Patents, and the Struggle to Control Radio: A Study of Big Business and the Uses of Industrial Research." *The Business History Review* 51 (2): 208–235.

———. 1980. "Industrial Research and the Pursuit of Corporate Security: The Early Years of Bell Labs." *The Business History Review* 54 (4): 504–529.

———. 1985. *The making of American industrial research : science and business at GE and Bell, 1876-1926.* Studies in economic history and policy. New York: Cambridge University Press.

Rosenberg, Nathan. 1990. "Why do firms do basic research (with their own money)?" *Research Policy* 19, no. 2 (April): 165–174.

———. 1994. *Exploring the Black Box: Technology, Economics, and History.* Cambridge: Cambridge University Press.

Sauermann, Henry, and Wesley M. Cohen. 2010. "What Makes Them Tick? Employee Motives and Firm Innovation." *Management Science* 56 (12): 2134–2153.

Schumpeter, Joseph A. 1942. *Capitalism, socialism, and democracy.* First edition. New York: Harper & Brothers.

Senecal, Vance E. 1980. "Du Pont and Chemical Engineering in the Twentieth Century." In *History of Chemical Engineering*, 190:283–301. Advances in Chemistry 190. AMERICAN CHEMICAL SOCIETY, June.

Shapiro, Carl. 2019. "Protecting Competition in the American Economy: Merger Control, Tech Titans, Labor Markets." *Journal of Economic Perspectives* 33, no. 3 (August): 69–93.

Spulber, Daniel F. 2013. "How Do Competitive Pressures Affect Incentives to Innovate When There Is a Market for Inventions?" *Journal of Political Economy* 121 (6): 1007–1054.

Stigler, George J. 1950. "Monopoly and Oligopoly by Merger." *The American Economic Review* 40 (2): 23–34.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, Themed Issue: Treatment Effect 1, 225, no. 2 (December): 175–199.

Tilly, Richard. 1982. "Mergers, External Growth, and Finance in the Development of Large-Scale Enterprise in Germany, 1880–1913." *The Journal of Economic History* 42, no. 3 (September): 629–658.

University of Michigan. 1919. *Financial Report of the University of Michigan for the Year Ended in 1919.* Ann Arbor: University of Michigan.

Watzinger, Martin, Thomas A. Fackler, Markus Nagler, and Monika Schnitzer. 2020. "How Antitrust Enforcement Can Spur Innovation: Bell Labs and the 1956 Consent Decree." *American Economic Journal: Economic Policy* 12, no. 4 (November): 328–359.

Watzinger, Martin, and Monika Schnitzer. 2022. *The Breakup of the Bell System and its Impact on US Innovation.* Rationality and Competition Discussion Paper Series 341. CRC TRR 190 Rationality and Competition, October.

Wise, George. 1985. *Willis R. Whitney, General Electric and the Origins of U.S. Industrial Research.* New York: Columbia Univ Pr.

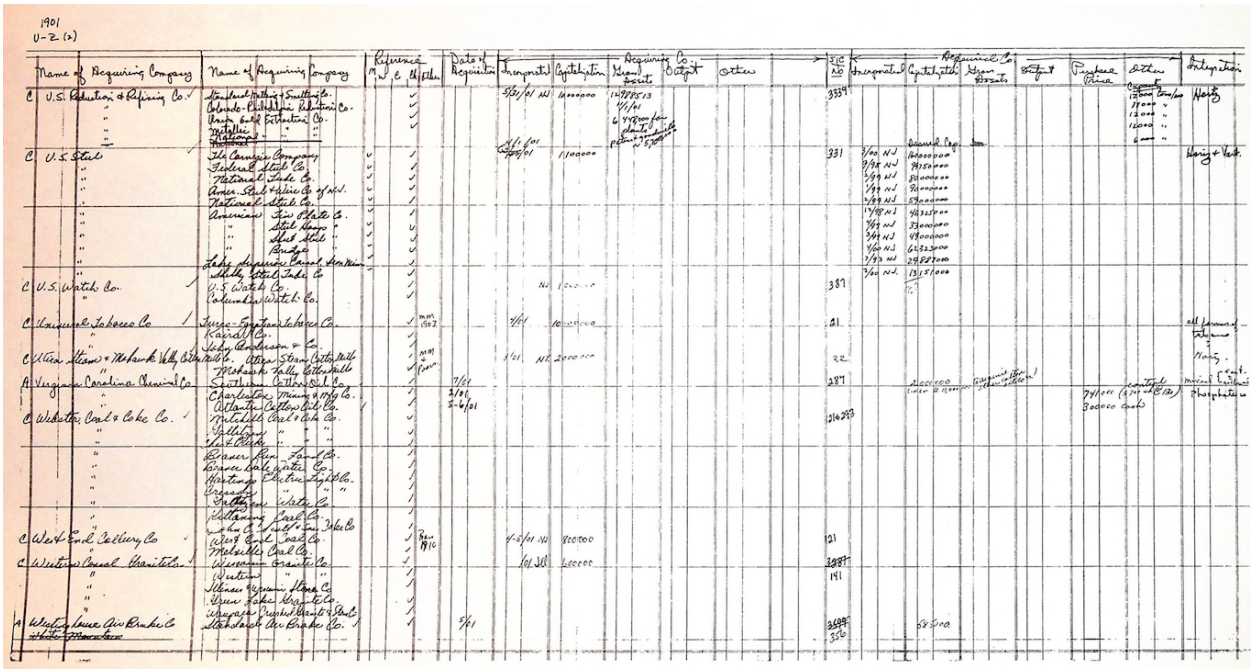
A Data sources

A.1 Nelson (1959) Merger and Acquisition Data

The core merger data derives from detailed handwritten worksheets compiled by Ralph Nelson for his seminal study *Merger Movements in American Industry, 1895–1956* Nelson (1959). See example in Figure A1. These worksheets, graciously shared by Naomi Lamoreaux, provide granular firm-level detail on merger and acquisition activity during the Great Merger Wave. For this paper, the worksheets were extensively digitized and cross-checked by hand.

Nelson assembled his database through systematic examination of the weekly Commercial and Financial Chronicle, supplemented by Moody's Manual, Poor's Manual, and government reports. Every consolidation underwent standardized follow-up verification, with Nelson tracking each firm through financial publications for five years to confirm consummation and continued operations. A secondary verification process checked excluded companies against later editions of financial manuals across multiple years. Nelson conducted additional checks comparing his data to other merger lists and sector-specific sources that had compiled statistics about merger activity. See Chapter 2 in Nelson (1959).

Figure A1: Example of a Nelson (1959) worksheet



A.2 National Research Council (NRC) R&D Laboratory Surveys

The NRC surveys represent the most complete documentation of historical industrial R&D activities in the US. Beginning after World War I as part of efforts to codify lab-

oratory locations and scientific personnel, the NRC conducted direct correspondence surveys with firms operating R&D facilities. The surveys covered publicly traded and private firms, providing nearly universal coverage, particularly starting in 1927. For this paper, data from surveys conducted in 1920, 1921, 1927, 1931, 1933, 1938, 1940, and 1946 were extensively digitized and cross-checked by hand.

The NRC established the population of firms through annual directories, scientific societies' firm lists, and advertising notices in technical journals. Letters were sent to firms doing any research work, with no sharp distinction between scientific and industrial research. The surveys defined research activities broadly, excluding only government-funded laboratories and those tied to educational institutions. Response rates were high as firms were keen to be included in this prestigious directory.

Standardized company names were obtained from Knott and Vinokurova (2023) and extensively cross-checked, modified, and enhanced. Laboratory locations were geocoded and validated against multiple sources. For firms operating multiple laboratories, information was carefully disaggregated at the laboratory location level.

A.3 Patent Data: CUSP Dataset (Berkes 2018)

Patent information derives from the Comprehensive Universe of U.S. Patents (CUSP) dataset (Berkes 2018), kindly shared by Enrico Berkers for the 1840-1960 period, providing broad U.S. patent coverage through systematic integration of five data sources: USPTO website, university/library databases, OCR-digitized patent images, Google Patents, and post-1920 Google-digitized patents.

Filing years are extracted through systematic text parsing starting from patent 137,279 (April 1873), achieving coverage for 93.2% of patents from official sources and 6.1% through text parsing. Technology classes come directly from the USPTO website across all classification schemes (USPC, CPC, IPC), with regular updates ensuring consistency. Name and location extraction employs a three-step approach with frequency-based fuzzy matching for location correction and coordinate assignment through offline databases supplemented by Google Maps queries.

The dataset achieves near-complete coverage with systematic quality indicators for all extracted variables and frequency-based typo correction for geographic locations.

A.4 Breakthrough Patent Measure (Kelly et al. 2021)

The Breakthrough measure used extensively in the paper was designed and computed by Kelly et al. (2021) using textual analysis of patent documents from 1840-2010. The approach modifies traditional TFIDF weighting with "backward-IDF" that calculates inverse document frequency using only prior patents.

The underlying patent importance combines novelty (low backward similarity to prior patents in a 5-year window) and impact (high forward similarity to subsequent patents in a 10- or 20-year window). The importance indicator is the ratio of forward to backward similarity (FS/BS). Breakthrough patents are defined as the top 10 percent after removing patent cohort fixed effects. In this paper, I prefer the 20-year window, given the wide availability of post 1940 data.

Historical validation using 250 historically significant patents shows these patents average in the 74th percentile of importance distribution (vs. 54th percentile for citations). Contemporary validation shows strong correlation with forward citations and significant predictive power beyond measurement horizon. Market validation shows positive correlation with Kogan et al. (2017) patent value estimates.

A.5 Subsidiary and Ownership Data (Kandel et al. 2019)

Corporate ownership data from Kandel et al. (2019) covers U.S. business groups and ownership for 1926-1950, beginning with Berle and Means (1932) list of 200 largest non-financial corporations (about 60% of non-bank corporate assets).

Control trees are mapped using Moody's manuals, tracing both upward (controlling shareholders) and downward (subsidiaries) ownership relationships. Ultimate control determination comes from newspaper archive searches, primarily the Wall Street Journal, supplemented by corporate histories. The dataset tracks relationships across seven time points (1926, 1929, 1932, 1935, 1937, 1940, 1950), creating an unbalanced panel of 15,270 firm-years spanning 2,743 firms.

Control trees are cross-checked against CRSP data and are consistent with original Berle and Means examples. Spot-checking against consolidated annual reports confirms coverage of significant entities.

A.6 Paper Mergers and Announcement Dates: Primary Sources

The paper-merger control group and the announcement dates draw on a common stack of digitized primary sources. The core corpus is the weekly *Commercial & Financial Chronicle*, retrieved through the FRASER archive of the Federal Reserve Bank of St. Louis and cached locally in full for 1890–1912, approximately 1,840 issues.[51] It is complemented by the trade weekly *Iron Age* (1890–1910), Moody's Manuals (1900–1925), R. G. Dun reference books (1890, 1900, and 1905), the New Jersey Secretary of State corporation registry, the *Chronicling America* newspaper corpus, period volumes accessed through Google Books, and CourtListener case law. Candidate paper mergers are identified from the contemporary compilations of Moody (1904)—including his list of "disintegrated, defunct, 'paper,' and inactive" combinations—Conant (1901), and Watkins (1927), net of Nelson's consummated consolidations.

The collection over these corpora was performed by large language model agents executing fixed search protocols. Three safeguards apply throughout: every dated claim rests on a verbatim quotation that is machine-verified against the cached source text; adversarial agents attempt to refute each coding before it is accepted; and all inclusion, fate, and sample-selection judgments were reviewed and decided by the author. Appendix B describes the resulting construction and its validation.

51. This product uses the FRASER® API but is not endorsed or certified by the Federal Reserve Bank of St. Louis.

A.7 Minor Data Sources

1. Moody's Manual of Industrial Securities (1900). Historical firm size, sector and name. For this paper, extensively digitized and cross-checked by hand from original sources.

2. Collins and Preston (1961). Listing of large firms by asset size used for size-based sample restrictions and robustness checks.

3. Conant (1901). Historical merger activity time series.

B Data construction

B.1 Firm Disambiguation and Linking Algorithm

This subsection describes a three-stage process to disambiguate patent assignees, harmonize firm identities across sources, and map corporate ownership over time.

B.1.1 Patent Assignee Processing

The initial challenge is to distinguish between individual and corporate assignees within the patent record, then disambiguate entities within each category. I employ a hybrid classification approach combining rule-based pattern matching with machine learning-based named entity recognition. Deterministic rules using precompiled patterns identify clear company indicators ("corporation," "manufacturing") and person indicators ("jr," "dr"), while transformer-based models classify ambiguous cases. This approach processes approximately 370,000 assignee strings, ultimately classifying 46% as firms and 54% as individuals (firm assignees obtain 78% of all patents).

For computational efficiency across hundreds of thousands of potential pairwise comparisons, I implement locality-sensitive hashing with MinHash signatures to create comparison blocks. Assignees are grouped based on Jaccard similarity thresholds, enabling efficient processing while preserving high-similarity pairs that likely represent the same entity.

Within each block, deterministic matching rules resolve assignee identities using string similarity measures, temporal proximity, and technological overlap. Match thresholds vary by context: exact and near-exact matches require Jaro similarity ≥ 0.975, while more distant matches demand additional constraints including geographic proximity and technological similarity. The algorithm applies more restrictive criteria for individuals than firms to prevent over-clustering of common names across long historical periods.

This process consolidates the roughly 170,000 strings classified as firms—about 46 percent of the 370,000 total—into approximately 137,000 unique firm assignees, a substantial disambiguation that maintains conservative matching criteria to minimize false positives.

B.1.2 Cross-Source Firm Harmonization

The second stage harmonizes firm identities across disparate historical sources: Nelson merger records, National Research Council laboratory surveys, subsidiary databases, manually collected ownership records, and asset rankings. Each source undergoes standardized name cleaning addressing historical spelling variations, abbreviations, and data entry inconsistencies.

Critical to this process is a manually constructed dictionary of over 1,500 alias relationships that captures name variations no algorithmic approach would identify. For example, linking "American Car & Foundry" with "ACF" requires domain knowledge about common industrial abbreviations that automated fuzzy matching would miss.

This manual collection is essential for accurately connecting firms across sources with different naming conventions and temporal coverage.

I employ parallel fuzzy matching between all source pairs. The system expands each source's name space using collected aliases before applying bidirectional matching and ensures symmetric relationships. Next, I impose transitivity across matches. Conflicts arise when transitivity requirements create impossible linkages—for instance, when different firms in one source are both linked to the same firm in another source. I resolve these conflicts through additional manual data collection, investigating historical name changes, corporate reorganizations, or source documentation errors. This iterative process continues until all conflicts are resolved, ensuring logical consistency across the harmonized firm universe.

B.1.3 Dynamic Ownership Mapping

The final stage resolves corporate ownership relationships over time, tracing ownership chains to identify ultimate controlling entities throughout the 1870-1960 period. This process integrates ownership data from multiple sources: Nelson's merger data, manually collected records, systematic subsidiary information (Kandel et al. 2019).

The ownership resolution algorithm begins with a baseline year (1869) where all firms own themselves absent contrary information, then processes ownership changes chronologically. When firm A acquires firm B, which previously owned firm C, the algorithm ensures that A's ownership of C is correctly propagated through the ownership chain.

The system traces ownership chains up to seven levels deep while detecting and handling circular ownership relationships. When chains exceed maximum length or contain circular references, these cases are flagged separately rather than making arbitrary assignments. Most ownership chains resolve in 1-2 steps, few reaching the maximum length and none exceeding it.

Throughout this process, I manually investigate and resolve ownership loops and conflicts using the same iterative approach applied to firm harmonization. The final algorithm accounts for approximately 23,000 ownership changes, enabling consolidation of patent activity at the enterprise level throughout the sample period.

B.2 GMW Panel Construction

The panel construction process creates balanced firm-year datasets spanning 1885-1940 that combine firm disambiguation results with patent outcomes, R&D laboratory data, and merger characteristics.

B.2.1 Treatment and Control Group Definition

The challenge is to define treatment units that capture the full innovative capacity of merging firms both before and after consolidation. For pre-merger periods, I implement retrospective group assignment: each constituent firm is assigned, for all years prior to consolidation, to the final GMW-era enterprise it eventually joined, with the event date set to that enterprise's first consolidation. For combinations built through several

successive deals, later-joining constituents are attributed to the group from that first consolidation date. This approach ensures that the measured pre-treatment innovation baseline reflects the complete innovative potential of the combining entities.

Treatment groups consist of firms participating in consolidations between 1895–1904, identified in the Nelson data. I track but do not use in the analysis acquisition-only enterprises, i.e. firms that expanded only through acquiring individual firms but underwent no major consolidation. Post-1904, the system tracks actual ownership changes through the dynamic mapping system described in Section B.1, handling cases where treatment groups themselves become acquisition targets.

Primary controls are firms with no record of Great Merger Wave participation. They comprise two groups: harmonized firms attested in the cross-source data—Nelson's worksheets, the 1900 Moody's Manual, the NRC laboratory surveys, or the subsidiary records—that never consolidated during the wave, and firm-classified patent assignees that carry no cross-source identifier. Both groups pass the same locality-sensitive hashing and within-block disambiguation applied to all assignees, and control firms with recorded acquisitions receive the same ownership roll-up as treated groups; "no harmonized identifier" therefore denotes the absence of a linked merger record, an operational criterion, rather than a verified history of unbroken self-ownership since 1885. The main residual asymmetry is that the cross-source ownership records, assembled around the GMW firms, are likely more complete for treated groups than for controls; the unbalanced-panel specification of Section C.3 limits the influence of thinly documented assignees by trimming their long zero-patent spells.

B.3 Paper-Merger Control Group

Candidate consolidations announced between 1890 and 1908 come from the compilations described in Appendix A.6, supplemented by an inverted *abandonment sweep* of the full *Chronicle* corpus: a search for abandonment language adjacent to deal language ("consolidation in abeyance," "plan reported off," "options expired") that surfaces short notices for deals absent from every compilation. Each of the resulting 263 candidates is assigned a fate—consummated and durable; operated, then failed later; quickly undone; explicitly or silently abandoned; unresolved. A deal is coded *quickly undone* when a combination was announced and reversed before any durable transfer of control or operation as a going concern, which distinguishes it from a deal that *operated, then failed later*; the latter is excluded from the control pool precisely because it received the consolidation treatment, if only briefly. Coding is deliberately conservative in the direction that matters for a control group: silent abandonment requires that no going-concern trace survive under any alias, and every abandonment coding is attacked by a refutation step that searches forward in time for evidence the combination in fact operated or survived under another name. Table B1 reports the selection funnel. Forty-five candidates are paper mergers; five screens—non-merger promotions, retail rollups, constituents recombined into a treated consolidation, insufficient evidence, and later evidence of an operating combine—reduce the pool to 30 deals, of which 24 carry a verified constituent roster and enter the deal-level analysis; the remaining 6 are documented only descriptively. Table B2 reports descriptive statistics: the pool clusters at the 1899 and 1902–03 peaks

of the wave, with a median authorized capitalization of $10.5 million and a median of 7 named constituent firms, spread across the same sectors as the treated consolidations.

Constituent rosters are reconstructed from the same sources, and only firm names that survive an independent verification step enter the data. Constituent names then pass through the identical standardization and matching code used for treated firms (Appendix B), at the same 0.95 similarity threshold and against the same full assignee reference file. Of the 266 distinct firms rostered at the linkage stage, 84—approximately 32 percent—matched at least one patent assignee; every match was individually adjudicated against the assignee's patenting era, state, and industry, with an adversarial re-check of consequential calls, leaving 55 kept matches, 16 wrong-entity homonyms, 12 firms dropped as treated in a later consolidation, and one flagged unresolved. Firms without a patent match are retained as structural zeros. A control firm exits the panel only upon losing its independence through absorption into another consolidation, dated from the collected evidence; the resulting exit table covers 36 firms across 11 deals. The panel is built as a balanced grid over 1885–1940 at the constituent–deal–year level and accounts for 1,455 patents and 38 breakthroughs after this sample restriction, and roughly 3,000 patents without it. Patenting is concentrated—the three largest deals account for roughly 90 percent of pool patents—which motivates the leave-one-deal-out robustness in Appendix C.4.

Table B1: Paper mergers: from candidate consolidations to the control pool

Selection step	Change	Deals
Consolidation candidates compiled, 1890–1908		263
less: consummated or durable combinations	−169	94
less: operated, then failed only later	−45	49
less: out of scope or not a merger	−4	45
Paper mergers (abandoned consolidations)		45
less: single-firm promotions (greenfield / single-patent / shell)	−7	38
less: retail / distribution rollups (out of scope)	−2	36
less: constituents recombined into a treated consolidation	−2	34
less: consummated or operating combine; duplicate attempt	−3	31
less: insufficient evidence on fate	−1	30
Control pool		**30**
of which: linkage sample (verified roster)		24
of which: descriptive-only (no verified roster)		6

Note: Candidate consolidations are compiled from contemporary lists (Moody 1904, Conant 1901, Watkins 1927) plus an abandonment sweep of the *Commercial & Financial Chronicle*, 1890–1908. The table reproduces the logic of selection; every count is derived from the deal-level classification fields. "Operated, then failed later" deals were consummated and operated before collapsing, and are excluded from the control concept.

B.4 Announcement Dates

Nelson (1959) dated each merger by the transfer of control, falling back to the incorporation date, and discarded "news items phrased in the indefinite terms of rumor and

Table B2: Paper mergers: descriptive statistics

Sample:	All Controls (1)	Linkage sample (2)	Descriptive only (3)
N	30 (100%)	24 (80%)	6 (20%)
Deal characteristics			
Year (median)	1900	1899	1902
Capitalization, $M (median)	10.5	10.0	18.0
Named constituent firms (median)	7	9	2
Cause of abandonment			
Target holdout / refusal	10 (33%)	9 (38%)	1 (17%)
Exogenous shock	3 (10%)	3 (12%)	0 (0%)
Financing / promotion breakdown	12 (40%)	9 (38%)	3 (50%)
Unknown / silent	5 (17%)	3 (12%)	2 (33%)
Type of abandonment			
Explicitly abandoned	15 (50%)	13 (54%)	2 (33%)
Silently abandoned	6 (20%)	5 (21%)	1 (17%)
Quickly undone	9 (30%)	6 (25%)	3 (50%)
Sector			
Mining & Natural Resources	2 (7%)	2 (8%)	0 (0%)
Food & Consumer Goods	9 (30%)	6 (25%)	3 (50%)
Chemicals & Materials	9 (30%)	6 (25%)	3 (50%)
Primary & Fabricated Metals	8 (27%)	8 (33%)	0 (0%)
Machinery & Electrical Equipment	1 (3%)	1 (4%)	0 (0%)
Transportation & Instruments	1 (3%)	1 (4%)	0 (0%)

Note: The pool of 30 abandoned consolidations. The linkage sample (24 deals) carries a verified constituent roster and enters the deal-level analysis; descriptive-only deals (6) are documented paper mergers without one. Cause-of-abandonment categories are descriptive classifications used for robustness cuts. Capitalization is the announced or authorized figure where stated.

prospective consummation." I recover that discarded layer for all 265 treated consolidations and all 30 abandoned deals under a single blind protocol: for abandoned deals, any pre-existing date was withheld from the search and revealed only at reconciliation, so that search effort could not differ across the two groups. Evidence is graded in three tiers—a prospectus, promoter statement, or corroborated report; a single trade-press report stated as fact; and rumor, which can set a date only if it names the vehicle or at least two of the deal's recorded constituents. Incorporation and organization notices are rejected as announcement evidence, since they date Nelson's later event, and serve only as a ceiling. Every unit—including those where no announcement was found—is then attacked by a second, adversarial agent instructed to find an earlier date.

Table B3 summarizes the outcome. Dates were resolved for 246 of the 295 units; 32 units retain Nelson's dating with a flag—no unit is dropped for lacking an announcement, which would select on press coverage—and 17 treated units, typically private family reorganizations, had no public announcement. Resolved dates are precise (239 of 246 to the day) and predominantly rest on the strongest evidence tier. Forty-six units move their annual event-time placement, and for 11 of the 31 abandoned-deal dates recorded at collection—one deal was later dropped from the pool—the blind rederivation showed the recorded "announcement" to be an incorporation notice, which the protocol replaces. Validation proceeds in two layers. First, every quotation in the evidence log is machine-verified against the cached source text; residual flags trace to OCR variants, and no fabricated quotation was found. Second, an independent falsification pass re-searched the 79 highest-risk units with a wider window and a different strategy: it found one annual-placement error, corrected in the current data, and eleven sub-annual refinements, implying accuracy within one quarter for roughly 97 percent of the stressed sample. Residual risk concentrates in the pre-1899 daily press, whose digitization is incomplete.

Table B3: Announcement dates: dispositions, evidence, and precision

	Treated consolidations (1)	Abandoned consolidations (2)	All units (3)
Units	265	30	295
Disposition			
Resolved (dated announcement)	220	26	246
Censored (no announcement found; Nelson dating retained)	28	4	32
No public announcement (private reorganizations)	17	0	17
Evidence tier of resolved dates			
A: prospectus, promoter statement, or corroborated report	128	14	142
B: single trade-press report, stated as fact	86	11	97
C: rumor passing the constituent-continuity rule	6	1	7
Precision of resolved dates			
Day	213	26	239
Month	2	0	2
Year	5	0	5
Event timing			
Movable (announcement year precedes Nelson year)	38	8	46

Note: Units are treated consolidation groups (265) and abandoned deals (30). "Censored" units retain Nelson's dating with a flag. Tier definitions are given in the text. "Movable" units have an announcement year strictly earlier than the Nelson-layer year and shift annual event time.

B.5 Variable construction

R&D Labs. For each firm-year observation, I construct measures counting facilities operational by the given year based on reported start dates, facilities with missing temporal

information, and an indicator for firms operating research facilities at any point during the sample period.

Patent Classification by Geographic Proximity to Labs. Patent counts are partitioned into mutually exclusive categories using a hierarchical classification system based on inventor location relative to firm laboratories. Patents with exclusively foreign inventors are classified separately. For patents with US inventors, I distinguish between firms with and without laboratory facilities, with non-laboratory firms forming a distinct category.

For laboratory-owning firms, patents are subdivided based on inventor location data availability and proximity to research facilities. Patents lacking US inventor coordinates are classified separately, while those with valid locations are assigned to distance-based categories using haversine calculations at a 50km threshold: proximity to operational laboratories, proximity to undated facilities, and distance beyond the threshold.

Firm Technology Classification. Technology categories derive from Cooperative Patent Classification section codes. Firm-level assignment employs patent-weighted modal classification across the entity's complete portfolio during 1885-1940. In cases of tied values, the algorithm defaults to the controlling firm's technology class.

Economic Sector Harmonization. For the extensive margin analysis, I reconcile SIC codes from Nelson merger data with Moody's Manual industrial categories (sections), creating seven harmonized categories: Machinery & Equipment (SIC 35-38, Moody's section 2), Metals & Materials (SIC 32-34, Moody's section 3), Textiles & Apparel (SIC 22–23,31, Moody's section 4), Miscellaneous Manufacturing (SIC 24-30,39, Moody's section 5), Mining & Extraction (SIC 10-15, Moody's section 6), Food & Tobacco (SIC 20-21, Moody's section 7), and Miscellaneous Non-Manufacturing (remaining codes). An eigth category captures missing information.

Name-Based Industry Classification. Because control firms are patent assignees for whom no industry information exists outside the patent record itself—and using patent classes would only reproduce the technology dimension—I classify every firm in the intensive-margin estimation sample into one of twelve broad product-market categories inferred from its name (for instance, food and tobacco; chemicals; fabricated metals and hardware; machinery and instruments; transportation equipment), chosen to be crosswalkable to two-digit SIC codes. The classification is performed by a large language model operating on the firm's name alone, with a conservative protocol: a first pass over all names; a second pass by a stronger model on every name the first could not place, instructed to reclassify only on confident recognition of a specific historical firm or an explicit product word; and a blind re-classification audit of a sample by three independent runs, which agree with the assigned category in 80 percent of cases, with disagreements concentrated in adjacent categories. Names carrying no product information, about half of the control group, remain unclassified by design. Two external checks are available. For the 254 consolidations with a recorded SIC code, the name-based category matches the SIC-implied one exactly in 73 percent of cases, against 29 percent under assignment of every firm to the most common category. For the 39 control firms attested in Moody's 1900 manual, agreement with Moody's coarser section scheme is 79 percent. Table B4 reports the resulting distribution.

Pre-1895 CPC Niches. For the same-niche robustness exercise of Section 4.4, I recon-

struct every sample firm's 1885–94 patent portfolio at three levels of the CPC hierarchy (subclass, main group, and subgroup, excluding Y-section tagging codes), using the same patent-to-firm attribution as the panel construction; the reconstructed pre-period counts match the panel exactly for 99.9 percent of the 11,801 sample firms. A consolidation's niche is defined from the pre-1895 portfolios of all 265 consolidation groups, including those later acquired. Table B5 reports the menu of exclusion rules and the control-group loss under each.

Table B4: Name-based industry classification of the intensive-margin sample

Industry	Treated	Controls	Share of controls
Food & tobacco	7	141	1.2%
Textiles, apparel & leather	2	287	2.4%
Wood, paper & printing	4	303	2.6%
Chemicals	7	259	2.2%
Petroleum & rubber	3	96	0.8%
Stone, clay & glass	7	132	1.1%
Primary metals & mining	7	280	2.4%
Fabricated metals & hardware	11	902	7.7%
Machinery & instruments	17	1,716	14.6%
Electrical equipment	2	674	5.7%
Transportation equipment	10	488	4.2%
Other	2	517	4.4%
Unclassifiable	0	5,927	50.6%
Total	79	11,722	100.0%

Table B5: Same-niche exclusion rules and control-group loss

Niche level	Top-k	Control match	# Niches	Controls dropped	% of controls
CPC subclass	1	modal	77	3,934	33.6%
CPC main group	1	$\geq$50% of patents	94	864	7.4%
CPC subgroup	1	any patent	93	478	4.1%
CPC subgroup	2	any patent	163	770	6.6%
CPC subgroup	3	any patent	219	938	8.0%

Note: A control firm is excluded when its own 1885–94 patenting overlaps a treated niche under the stated rule. Rules based on the top two or three subgroups with any-patent overlap are infeasible under the stratified reweighting: the nine control firms with more than 75 pre-period patents all patent in some consolidation's top-two subgroups, emptying the largest reweighting stratum.

B.6 Inventor Disambiguation and Panel Construction

Identifying unique inventors across the patent record presents a fundamental challenge for studying individual innovative careers and inventor-firm relationships over long historical periods. Patent records contain only inventor names and locations without unique identifiers, while individuals may appear inconsistently across patents due to name variations, geographic mobility, and other changes.

B.6.1 Probabilistic Record Linkage Framework

I implement inventor disambiguation using the Fellegi and Sunter (1969) probabilistic record linkage framework in python using the splink package, which estimates match probabilities based on agreement patterns across multiple comparison dimensions.

The system employs 28 distinct blocking rules that create candidate pairs for detailed comparison, balancing computational efficiency with broad coverage. Core blocking strategies include exact year matches with name token overlap, geographic proximity with name similarity, technological domain overlap within temporal windows, and assignee relationship overlap.

The comparison framework incorporates six key dimensions: name similarity (using exact matches, component-wise matching, and fuzzy string similarity), technological overlap (via CPC hierarchy at multiple levels), geographic proximity (coordinate-based distances and administrative boundaries), temporal distance, co-authorship patterns, and assignee relationships. Name comparisons receive term frequency adjustments to account for commonality.

Model training uses expectation-maximization algorithms applied to candidate pairs, iteratively refining match probability estimates until convergence. The system processes approximately 800 million candidate pairs to generate final match probabilities, applying threshold-based clustering at 0.9 probability cutoffs for conservative cluster assignment. This approach identifies 1.012 million unique inventors responsible for 2.273 million patents between 1875-1955.

B.6.2 Panel Construction

Creating analysis-ready inventor panels requires addressing both career tracking challenges and the specific requirements of Abowd-Kramarz-Margolis (AKM) two-way fixed effects estimation. For each disambiguated inventor, I construct balanced yearly observations spanning their first to last patent, creating annual observations for all intervening years regardless of patenting activity. This approach enables analysis of both productive and unproductive periods within inventor careers while providing the temporal continuity necessary for fixed effects identification.

Firm assignment follows hierarchical rules prioritizing corporate assignees over individual assignments, using the harmonized firm mapping from Section B.1. For inventors with multiple firm affiliations within a year, assignment uses patent-count-based weighting, assigning inventors to the firm receiving the plurality of their patents. Missing location and firm information for non-patenting years is carried over from the nearest patenting year. This rule assigns a definite firm and location to zero-output years rather than interpolating, and so imposes a specific timing on inventor moves between observed patents.

Patent quality measurement for the AKM framework requires positive-valued outcomes suitable for logarithmic transformation, addressed through careful processing of the Kelly et al. (2021) breakthrough measures. I apply an exponential transformation using the formula $q = 10 \cdot exp(4y - 0.8y^2)$, where y represents the original breakthrough score, ensuring strictly positive values while preserving meaningful quality variation

across patents. This transformation roughly matches mean and 99th percentile from the Kogan et al. (2017) patent value distribution. Notice that it is strictly monotone in the observed range of y.

Table B6: Inventor Panel Summary Statistics

	Mean (1)	SD (2)	P10 (3)	P25 (4)	P50 (5)	P75 (6)	P90 (7)
Panel A: All Inventors							
Career Length (years)	4.03	7.55	1.00	1.00	1.00	2.00	12.00
Patents	2.52	7.61	1.00	1.00	1.00	2.00	4.00
Patents (fractional)	2.25	6.97	0.50	1.00	1.00	2.00	4.00
Breakthroughs	0.22	1.98	0.00	0.00	0.00	0.00	0.00
Breakthroughs (fractional)	0.18	1.66	0.00	0.00	0.00	0.00	0.00
Quality-weighted Patents	3.85	21.39	0.39	0.63	1.10	2.30	5.93
Quality-weighted Patents (fractional)	3.35	18.93	0.31	0.53	0.97	2.03	5.19
Technology Sections	1.57	0.98	1.00	1.00	1.00	2.00	3.00
States	1.11	0.46	1.00	1.00	1.00	1.00	1.00
Distinct Firms	0.32	0.67	0.00	0.00	0.00	1.00	1.00
N				1,011,606			
Panel B: Inventors in ≥ 2 years							
Career Length (years)	12.01	10.93	2.00	4.00	8.00	17.00	27.00
Patents	6.37	13.76	2.00	2.00	3.00	6.00	12.00
Patents (fractional)	5.70	12.63	1.50	2.00	3.00	5.00	11.00
Breakthroughs	0.64	3.72	0.00	0.00	0.00	0.00	1.00
Breakthroughs (fractional)	0.53	3.11	0.00	0.00	0.00	0.00	1.00
Quality-weighted Patents	10.39	39.59	1.22	1.88	3.40	7.66	19.82
Quality-weighted Patents (fractional)	9.03	35.04	1.03	1.65	3.01	6.71	17.10
Technology Sections	2.46	1.32	1.00	1.00	2.00	3.00	4.00
States	1.40	0.80	1.00	1.00	1.00	2.00	2.00
Distinct Firms	0.72	1.02	0.00	0.00	0.00	1.00	2.00
N				278,407			

Note: This table presents summary statistics for the disambiguated inventor panel spanning 1875-1955. Panel A includes all inventors, while Panel B restricts to inventors active for multiple years. Career length measures years from first to last patent. Patents and breakthroughs use both standard and fractional (co-invention adjusted) measures. Quality-weighted patents apply exponential transformation to Kelly et al. (2021) breakthrough scores. Technology sections count distinct CPC sections, states count distinct inventor locations, and distinct firms count harmonized firm affiliations.

To handle extreme outliers that could distort fixed effects estimation, I implement

floor values at two stages. During panel creation, individual patent quality measures below 1% of their median are set to this threshold, preserving the underlying quality distribution while ensuring numerically positive values. Fractional patent measures adjust for co-invention by dividing quality measures by coinventor count, enabling accurate attribution in the AKM decomposition. Fractional patent measures adjust for co-invention by dividing quality measures by coinventor count, enabling accurate attribution in the AKM decomposition: fractional patent measure = patent measure / number of coinventors.

Table B6 presents summary statistics for the inventor-level panel. Panel A covers all 1,011,606 disambiguated inventors, including those with patents assigned to individuals rather than firms. The median inventor has a single patent over a one-year career, reflecting the highly skewed nature of innovative productivity. Panel B focuses on the 278,407 inventors active for multiple years, revealing more substantial careers with median length of 8 years and 3 patents. Quality-weighted patent measures show considerable variation, with the 90th percentile inventor producing nearly 20 quality-weighted patents compared to 3.4 for the median multi-year inventor.

The AKM estimation requires specific sample restrictions to ensure realistic career patterns suitable for two-way fixed effects identification. I remove inventors with implausibly long careers (>67 years), excessive geographic mobility (>20 states), or extreme firm switching behavior (>100 distinct employers). These restrictions prevent spurious disambiguation matches that could bias firm and worker effect estimates while maintaining sufficient inventor mobility for identifying separate firm and individual effects. Recall that I implement the KSS leave-one-out bias-correction, which effectively requires each firm to be linked to at least two inventors and inventors appearing in at least two distinct years. In the AKM sample preparation, aggregated inventor-firm spell outcomes below 1% of the spell-level median are clipped before logarithmic transformation, ensuring again numerically positive values. Table B7 compares the size of the full panel and that of the AKM sample.

Table B7: AKM vs Full panel sample sizes

	Full Panel (1)	Excluding Singletons (2)	LOO-Connected Sample (3)
Observations (person-year)	3,847,349	3,169,090	1,310,550
Inventors	944,795	266,536	94,040
Firms	82,042	59,859	18,286

B.6.3 Validation

Comparison with Akcigit et al. (2022) provides external validation of the disambiguation process. Table B8 shows that for the overlapping 1940-1955 period, my disambiguation identifies 239,026 inventors compared to their 300,077, with higher average patents per inventor (2.5 vs. 2.1) and greater geographic mobility (9.1% vs. 4.7% moving states). These differences reflect: (a) different underlying patent datasets; (b) temporal coverage differences (their data starts in 1940, mine extends to 1955); and (c) richer information used in my disambiguation process, including technological overlap and assignee relationships that enable more accurate linking of inventor records across patents.

Table B8: Comparison with Akcigit et al. (2022)

Period:	1940-1955	
Sample:	This Paper	AGNS
Number of inventors	239,026	300,077
Patents/year	0.71	0.63
2+ patents (%)	34.9	27.5
5+ patents (%)	10.5	8.1
10+ patents (%)	3.9	3.0
Patents/inventor	2.5	2.1
Moved states (%)	9.1	4.7
Years active	3.5	7.2

Note: This table compares inventor disambiguation results with Akcigit et al. (2022) for the overlapping 1940-1955 period. Differences reflect distinct underlying datasets, temporal coverage (AGNS ends in 1940), and richer disambiguation information including technological overlap and assignee relationships.

Additional validation comes from comparing disambiguation results against external benchmarks for highly prolific inventors. Matching patent counts to Wikipedia records for prominent historical inventors yields high coverage of the benchmark counts: Thomas Edison (989/1084 patents, 91%), Francis H. Richards (845/894, 95%), Elihu Thomson (623/696, 90%), John F. O'Connor (999/949, 105%), Carleton Ellis (731/753, 97%), Melvin de Groote (907/925, 98%), and George Albert Lyon (866/993, 87%). These figures are recovered-to-benchmark count ratios rather than accuracy rates—a value above 100 percent, as for O'Connor, reflects incompleteness in the Wikipedia benchmark—so they gauge coverage and do not, on their own, separate benchmark under-counting from over-linking.

B.7 Technology-Level Panel Construction

The challenge is to create panels at the level of technology domains that enable quantitative analysis.

B.7.1 Technology definition

The distinction between established and emerging technologies rests on patent activity patterns during the pre-consolidation baseline period (1880-1894). Established technologies encompass CPC groups with positive patent activity during this period. Emerging technologies comprise CPC groups with zero baseline activity but positive innovation beginning in 1895 and within 1960 (when my patent data ends).

My analysis of established technologies estimates effects in levels, which requires comparable technology sizes for effects to be meaningfully comparable and not counfeded by levels. To create such units while preserving technological coherence, I implement a bottom-up clustering algorithm that groups related CPC categories based on both patent volume and hierarchical proximity. The algorithm calculates target cluster size as total pre-1895 weighted patents divided by desired cluster count (approximately 1,000 clusters).

The clustering process respects the CPC hierarchy's nested structure, processing levels from most specific (subgroup) to broadest (section). At each level, the algorithm assigns large individual units ($\geq$ 75% of target size) their own clusters, groups medium-sized sibling units (collective size 50-150% of target) into single clusters, and splits large sibling groups using cumulative distribution boundaries to create approximately equal-sized clusters.

Post-processing consolidation addresses remaining small clusters through CDF-based redistribution within sections, ensuring no cluster falls below 30% of target size while respecting technological boundaries.

The hierarchical clustering algorithm creates balanced technology domains while preserving meaningful technological distinctions. Table B9 shows that GMW-exposed and unexposed technology clusters exhibit comparable characteristics across key dimensions.

B.7.2 Panel Construction

Technology-level analysis requires measuring each domain's exposure to Great Merger Wave activity. In the sample of established technologies (with non-zero patents before 1895), I define a technology to be exposed if GMW firms had any patents in it before 1895. For emerging technologies with zero pre-1895 patents, I define GMW exposure using the exposure status of related technologies, as measured by CPC hierarchical relationships. In practice, an emerging CPC group is marked as exposed if an established CPC group in the same CPC subclass had any pre-1895 patents by GMW firms.

Panel construction creates annual technology-year observations with multiple innovation measures: total patent counts, breakthrough innovation, and measures tracking firm entry, exit, and inventor mobility patterns. The panels incorporate proportional weighting for patents with multiple CPC classifications, where each patent-CPC combination receives weight equal to 1/N (where N represents total classifications), ensuring aggregate counts sum correctly while preserving technological diversity information.

Table B9: Technology Cluster Balance by GMW Exposure

	Mean	SD	P10	P25	P50	P75	P90
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: No GMW Exposure							
Earliest Year (median)	1851.36	9.43	1840.00	1844.00	1850.00	1856.50	1865.00
Earliest Year (mean)	1852.37	8.96	1840.00	1845.67	1851.67	1857.80	1864.30
Groups	5.56	4.84	1.00	2.00	4.00	7.00	12.00
Subclasses	1.51	0.98	1.00	1.00	1.00	2.00	3.00
Classes	1.07	0.25	1.00	1.00	1.00	1.00	1.00
Sections	1.00	0.00	1.00	1.00	1.00	1.00	1.00
Weighted Patents (std.)	93.69	88.10	30.01	42.83	70.33	111.68	187.06
N				307			
Panel B: GMW Exposure							
Earliest Year (median)	1852.07	9.96	1840.00	1843.00	1851.50	1858.00	1866.00
Earliest Year (mean)	1853.29	9.61	1840.00	1845.00	1853.22	1859.75	1866.34
Groups	6.91	6.39	1.00	3.00	5.00	9.00	14.00
Subclasses	1.64	1.17	1.00	1.00	1.00	2.00	3.00
Classes	1.13	0.39	1.00	1.00	1.00	1.00	2.00
Sections	1.00	0.00	1.00	1.00	1.00	1.00	1.00
Weighted Patents (std.)	80.61	63.02	27.94	42.15	62.65	98.30	151.47
N				670			

Note: This table shows descriptive statistics for established technology clusters by GMW exposure status. Earliest year measures temporal origins using both median and mean patent years within clusters. Groups, subclasses, classes, and sections count distinct CPC categories. Weighted patents (std.) is the within-cluster standard deviation of subgroup-level patent counts; it measures dispersion within clusters rather than total cluster size, which the clustering algorithm equalizes by construction.

C Additional results — Impact of the Great Merger Wave on Merging Firms

This appendix presents additional empirical results supporting the main analysis of how the Great Merger Wave affected innovation among merging firms.

Section C.1 presents additional evidence on the determinants of the GMW and the selection of merging firms. Section C.2 shows relative-time event study and pre-merger innovation trends of individual constituent firms to address concerns about aggregating firm outcomes in calendar time. Section C.3 presents robustness checks including: horizontal consolidations only (Figure C6), trimmed samples excluding outliers (Figure C7), unbalanced panels that retain observations from three years before a firm's first patent through three years after its last (Figures C8 and C9), and patent filing dates rather than grant dates (Figures C10, C11, and Table C5). Section C.3 also reports the product-market robustness battery: alternative same-niche exclusion rules and separate fixed-effects variants (Table C4) and event studies under the combined specification of Table 2, column 3 (Figures C5a and C5b). Section C.6 explores heterogeneity by manufacturing sectors (Figure C13), integration type (Figure C14), and business success classification (Figure C15). Section C.7 analyzes R&D laboratory establishment, including event studies for number of active laboratories (Figures C16 and C17) and heterogeneity by market concentration and technological area (Figures C18 and C19). Table C11 provides descriptive statistics on firm exits and laboratory establishment.

Table C1: Effect of Consolidation on Firm-Level Innovation—Static Diff-in-Diff Estimates

Margin:	Intensive			Extensive	
Outcome:	Patents	Breakthroughs	Has lab	Started patenting	Has lab
	(1)	(2)	(3)	(4)	(5)
GMW Firm	5.989	0.563	0.162	0.231	0.045
	(1.910)	(0.181)	(0.033)	(0.033)	(0.015)
	[2.245, 9.732]	[0.208, 0.919]	[0.098, 0.227]	[0.166, 0.295]	[0.017, 0.074]
Controls	Y	Y	Y	Y	Y
Treated pre-mean	1.929	0.105	0.019	0.000	0.001
Control post-mean	0.644	0.045	0.018	0.138	0.012
# Firms	11,801	11,801	11,801	851	851
N	542,846	542,846	542,846	39,092	39,092

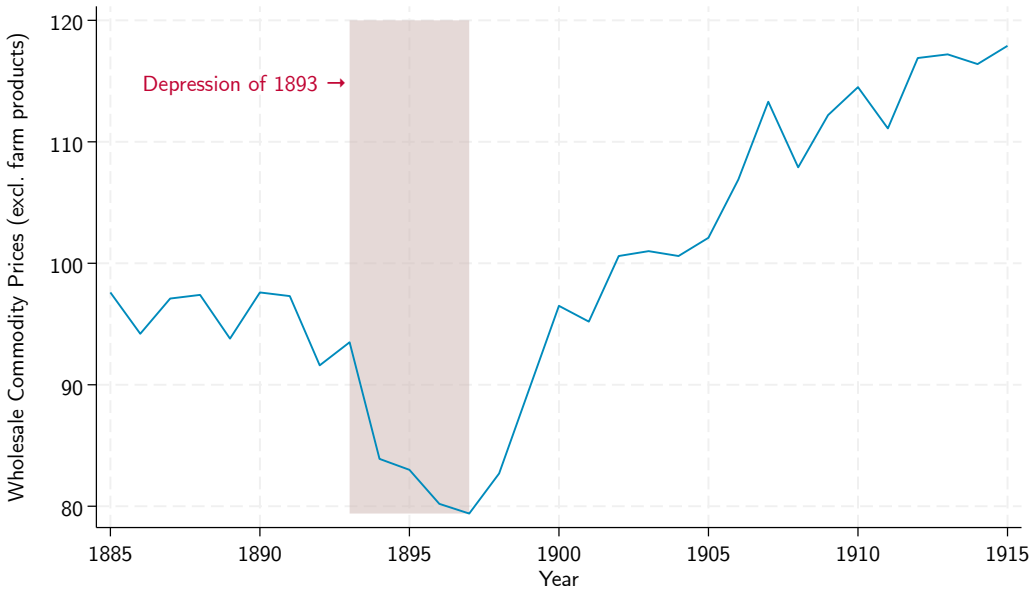
Note: This table reports static difference-in-differences estimates of the effect of consolidation on firm-level innovation outcomes (Equation 4). Columns 1-3 show results for the intensive margin (firms with pre-merger patents), while columns 4-5 show extensive margin effects (firms with no pre-merger patents). The dependent variables are: patent counts (col. 1), breakthrough patent counts (col. 2), probability of having at least one R&D lab (col. 3 and col. 5), probability of patenting at least once by a given year (col. 4). "GMW Firm" denotes the post-1895 interaction $\mathbb{1}[t \geq 1895] \times \text{GMW Firm}_i$ from equation (4); the firm fixed effects absorb the level consolidation indicator. Intensive margin regressions include time fixed effects by technological area, while extensive margin include time fixed effects by economic sector. Standard errors clustered at the firm level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

C.1 Background Evidence

This section provides additional evidence that mergers were triggered primarily by price competition, financial pressure, and legal incentives, rather than expectations of future technological potential. To this end, I (i) document the sharp deflation of wholesale commodity prices during the 1893–1897 depression, (ii) place consolidating firms in the distribution of manufacturing firm size circa 1900, (iii) review Lamoreaux's (1985) industry-level logit regressions, and (iv) present new probit regressions predicting consolidation on firm characteristics.

Figure C1 illustrates the deflationary environment of the 1893–1897 depression. Wholesale commodity prices fell sharply (about 15 percent) and recovered only after 1898, consistent with historical accounts that intensified price competition forced firms to seek consolidation as a survival strategy. Figure C2 compares the capitalization of consolidating firms to the overall distribution of firms in the 1900 *Moody's Manual*. Consolidating firms were somewhat larger on average, but not extreme outliers: they came from the heart of the size distribution, suggesting that the merger wave was not confined to a few giant incumbents but affected a broad swath of medium-to-large manufacturers.

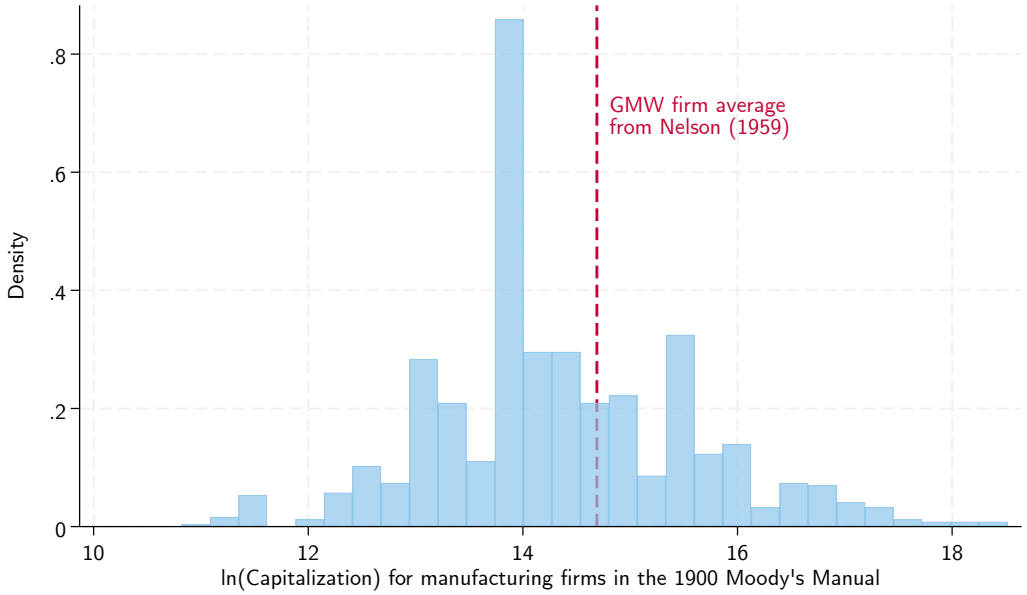
Figure C1: Deflationary Nature of the 1893-1897 Depression



Note: This figure shows wholesale prices for all commodities other than farm products between 1885 and 1915, highlighting the impact of the Depression of 1893-1897. The index base is such that 100 = average between 1890-1914. Data is from Hanes (2006).

Industry-level selection into consolidation. Table C2 reports results from Lamoreaux (1985). The outcome is an industry-level indicator for whether the sector experienced significant horizontal consolidation (1895–1904) from Nelson's list; industries with only marginal activity or missing data are excluded. The final sample has 232 industries, 52

Figure C2: Firm Size (Capitalization) Comparison



Note: This figure contextualizes the average size of individual firms involved in GMW consolidations as computed by Nelson (1959) with the capitalization distribution of manufacturing firms attested in the 1900 Moody's Manual of Industrial Securities.

with consolidation, 180 without. Data on industry characteristics come from the Census of Manufactures. Lamoreaux (1985)'s specification is:

$$\Pr\{Consolidation_i = 1\} = \Lambda\Big(\beta_0 + \beta_1\, Fixed_i + \beta_2\, Growth_i + \beta_3\, Growth_i \times HighFixed_i + \beta_4\, Margin_i + \beta_5\, Size_i\Big), \quad \text{(C1)}$$

where Λ is the logistic function and:

- *Consolidation*: 1 if the industry underwent significant consolidation (1895–1904), 0 otherwise.
- *Fixed*: capital–output ratio (capital / annual output), proxy for fixed-cost burden.
- *Growth*: percent increase in capital, 1889–1899.
- *HighFixed*: 1 if *Fixed* is above the manufacturing mean, else 0.
- *Margin*: profit margin on sales (earnings per dollar of sales).
- *Size*: capital per establishment (thousand $ / establishments), proxy for large, mass-production plants.

Table C2: Industry correlates of consolidation (Lamoreaux, logit)

	Model A (Table 4.3)	Model B (Table 4.4)
Fixed (capital–output)	1.220	1.240
	(0.552)	(0.549)
	[0.138, 2.302]	[0.165, 2.315]
Growth (1889–1899)	-0.000702	
	(0.00439)	
	[-0.0093, 0.0079]	
Growth × HighFixed	0.0109	0.0104
	(0.00568)	(0.00435)
	[-0.0002, 0.0220]	[0.0019, 0.0189]
Margin (profit per $ sales)	-8.640	-8.650
	(2.851)	(2.855)
	[-14.229, -3.051]	[-14.246, -3.054]
Size (capital / establishments)	0.00498	0.00496
	(0.00178)	(0.00179)
	[0.0015, 0.0085]	[0.0015, 0.0085]
Constant	-1.200	-1.220
Observations	232	232
LR χ^2	62.30	62.27

Notes: Coefficients are Lamoreaux's maximum-likelihood logit estimates. Standard errors in parentheses and 95% confidence intervals in square brackets are back-calculated from the published z-statistics (SE $= |\hat{\beta}|/|z|$; CI = $\hat{\beta} \pm 1.96 \times$ SE). Confidence intervals assume asymptotic normality and no finite-sample bias corrections.

Across specifications, industries with thin margin and large, capital-intensive plants are more likely to consolidate: lower earnings per dollar of sales (an immediate proxy of price pressure) strongly predict consolidation even conditioning on fixed-cost intensity and size. The pattern points to consolidation as a response to price competition under high fixed-cost pressure.

Firm-level selection into consolidation. I estimate probit models predicting whether a firm was absorbed into a GMW consolidation. The dependent variable is $GMW_i \in \{0, 1\}$. I conduct two separate analyses in two different samples:

1. **Moody's Manual (1900).** The sample includes firms listed in the *1900 Moody's Manual of Industrial Securities*. Treatment firms are GMW constituent firms that disappeared in a consolidation before 1904 but were still listed independently in

1900. Controls are independent manufacturing firms that appear in Moody's and were not involved in consolidation. Coverage is limited: only a small number of constituent GMW firms are observed (many had already merged by 1900), and the Manual covers only about 1,200 firms. Thus this exercise should be read as an extension of Figure C2—using actual listed constituent firms, rather than Nelson's imputed averages, to proxy the position of GMW firms in the size distribution. The specification is:

$$\Pr\{GMW_i = 1\} = \Phi\Big(\gamma_0 + \gamma_1 \ln(\textit{Capitalization}_i) + \sum_s \gamma_s \mathbb{1}[\textit{Sector}_i = s]\Big). \tag{C2}$$

Sectors are economic sectors defined in Moody's. Robust standard errors are reported.

2. **Patent Record (pre-1895 inventing firms)** The sample includes firms linked to the U.S. patent record with nonzero patenting activity before 1895. Treatment firms are those that subsequently entered a GMW consolidation; controls are other pre-1895 patenting firms that did not consolidate. This sample is broader in coverage than Moody's but still selective, as it excludes firms with no patenting prior to 1895. Here I can also include technology-area dummies and firm-level pre-1895 patent counts as additional regressors. The specification is:

$$\Pr\{GMW_i = 1\} = \Phi\Big(\phi_0 + \sum_t \phi_t \mathbb{1}[\textit{CPC Section}_i = t] + \phi_1 \textit{ Patents 1880-1894}_i\Big). \tag{C3}$$

Sections are defined from the patent record as the modal section of a firm's patent portfolio. Robust standard errors are reported.

Table C3 presents the estimates. Across both samples, the results point to only weak selection into consolidation. Larger firms were modestly more likely to be absorbed, and sectoral dummies absorb some variation. Innovation history—measured by pre-1895 patenting—is a statistically significant but economically small predictor: including it raises the pseudo-R^2 of the patent-record specification from about 0.01 to 0.04, with a positive and significant coefficient, yet the models' overall explanatory power remains minimal, with pseudo-R^2 barely exceeding 0.06 in column (2). These patterns reinforce the historical interpretation that mergers were not systematically organized around firms' innovative potential.

Table C3: Probit models of consolidation participation

	Outcome: GMW Firm$_i$ = 1			
	Sample: 1900 Moody's Manual		*Sample:* Patent Record	
	(1)	(2)	(3)	(4)
ln(Capitalization)		0.102	Patents 1880–1894	0.017
		(0.053)		(0.006)
		[-0.001, 0.205]		[0.006, 0.028]
Moody Sections			**Tech Areas**	
Metals & Materials	0.532	0.516	Operations & Transport	0.207	0.200
	(0.264)	(0.264)		(0.084)	(0.085)
	[0.015, 1.050]	[-0.002, 1.034]		[0.042, 0.373]	[0.034, 0.366]
Textiles & Apparel	-0.519	-0.433	Chemistry & Metallurgy	0.190	0.194
	(0.425)	(0.430)		(0.142)	(0.143)
	[-1.352, 0.314]	[-1.276, 0.411]		[-0.089, 0.469]	[-0.086, 0.474]
Misc. Manufacturing	0.212	0.242	Textiles & Paper	-0.034	-0.039
	(0.257)	(0.259)		(0.170)	(0.171)
	[-0.291, 0.716]	[-0.266, 0.751]		[-0.368, 0.300]	[-0.374, 0.296]
Mining & Extraction	-0.128	-0.119	Construction	0.007	0.011
	(0.277)	(0.279)		(0.147)	(0.147)
	[-0.671, 0.415]	[-0.666, 0.428]		[-0.281, 0.294]	[-0.277, 0.298]
Food & Tobacco	0.609	0.619	Mechanics	0.055	0.052
	(0.344)	(0.348)		(0.107)	(0.108)
	[-0.066, 1.284]	[-0.063, 1.301]		[-0.155, 0.266]	[-0.159, 0.263]
Misc. Non-Manufacturing	-0.178	-0.163	Physics	-0.058	-0.087
	(0.264)	(0.266)		(0.144)	(0.146)
	[-0.695, 0.339]	[-0.685, 0.359]		[-0.341, 0.225]	[-0.373, 0.199]
			Electricity	-0.094	-0.212
				(0.168)	(0.202)
				[-0.423, 0.234]	[-0.607, 0.183]
			Other Technology	0.611	0.459
				(0.206)	(0.226)
				[0.207, 1.014]	[0.016, 0.902]
Constant	-1.796	-3.288	Constant	-2.255	-2.310
	(0.200)	(0.814)		(0.072)	(0.074)
	[-2.189, -1.404]	[-4.883, -1.693]		[-2.397, -2.114]	[-2.455, -2.164]
Pseudo R^2	0.053	0.061	Pseudo R^2	0.010	0.041
GMW Firms	40	40	GMW Firms	188	188
Non-GMW Firms	917	917	Non-GMW Firms	11722	11722

Notes: Dependent variable is an indicator for GMW constituent firm. Columns 1 and 2 use firms listed in *Moody's Manual, 1900*. Columns 3 and 4 use firms in the patent record with nonzero pre-1895 patents. Standard errors in parentheses, 95-percent confidence intervals in brackets.

C.2 Event Study in Relative Time

One might worry that defining treatment effects in calendar time masks pre-trends around specific merger dates. To address this concern, I also estimate a relative-time specification using the bias-correction method by Sun and Abraham (2021). The specification is:

$$y_{itc} = \alpha_i + \delta_{tc} + \sum_{m=-10, m\neq -1}^{30} \beta_m \mathbb{1}[t - ED_i = m] + \kappa_1 \mathbb{1}[t - ED_i < -10] + \kappa_2 \mathbb{1}[t - ED_i > 30] + \varepsilon_{itc}, \tag{C4}$$

where i indexes firms, t years, c technological areas, m relative time, and ED_i is the year of the earliest public announcement of i's consolidation (the first, if multiple consolidations are recorded; Appendix B.4). The sample runs from 1885 to 1930, and standard errors are clustered at the firm level. Never-consolidated control firms have no event date: ED_i is undefined and all event-time indicators are zero for them, so they enter only through the technology-by-year effects δ_{tc}, and the β_m are identified from consolidating firms relative to this never-treated pool. The event study effects of interest are captured by β_m. The κ terms bin the endpoints so that the estimation sample is fixed across specifications; because consolidation dates are staggered within a fixed calendar window, the far post-event coefficients are identified from the earlier cohorts alone. I use the same IPW weights as in the main specification.

Figure C3 plots the resulting event study estimates for the key outcomes of patents (panel a) and breakthroughs (panel b). The results support the validity of my preferred approach.

The two panels rest on different estimators by necessity. For patents, a dense outcome, the interaction-weighted estimator of Sun and Abraham (2021) is well behaved; for breakthroughs it is not. The estimator fits a saturated cohort-by-relative-time regression, and because breakthroughs are rare—only about 13 percent of firm-years are nonzero—most interaction cells are nearly empty, driving their within-cell variances toward zero and collapsing the aggregated event-study standard errors to a small fraction of their pooled counterparts. Panel (b) therefore reports the pooled two-way fixed-effects event study, whose coefficients nearly coincide with the interaction-weighted ones given the large never-treated pool. The static estimates reported in the body are unaffected, since the static aggregation remains dense. Anchoring the event clock at Nelson's consolidation dates rather than the announcement leaves the results essentially unchanged—the static intensive-margin effect is 6.05 (SE 1.96) under announcement timing against 6.09 under Nelson timing—indicating that dating precision does not drive the results.

Another potential concern is that aggregating the pre-merger output of constituent firms into a consolidation-level sum might mask divergent firm-level innovation trends. To address this concern, I implement an additional pre-trends test that examines the innovation trajectories of individual constituent firms leading up to their consolidation dates.

The specification is:

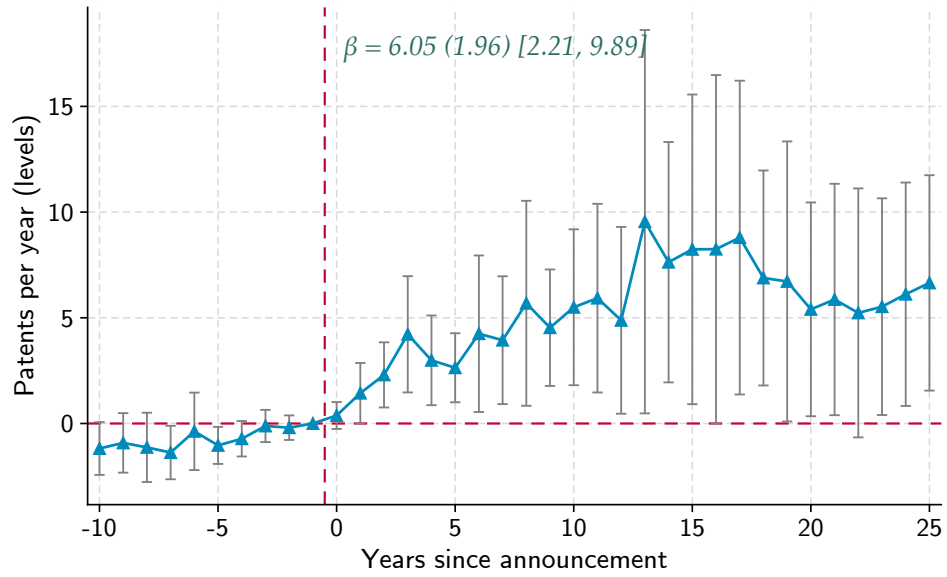
$$y_{itc} = \alpha_i + \delta_{tc} + \sum_{m=-10}^{-2} \beta_m \mathbb{1}[t - ED_i = m] + \varepsilon_{itc}, \tag{C5}$$

where i indexes firms, t years, c technological areas, m relative time, ED_i is the event date. The sample runs from 1885 to 1903, and standard errors are clustered at the firm level.

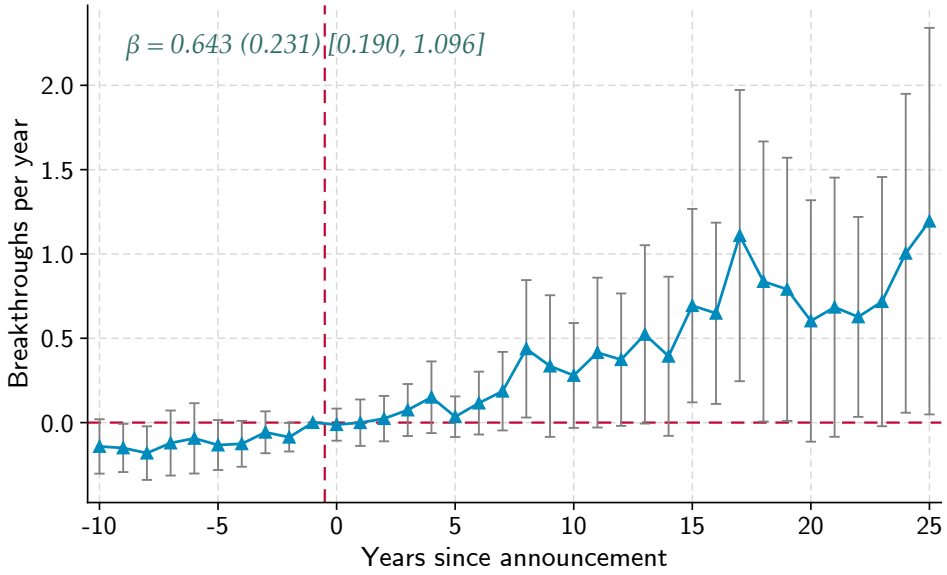
The event study effects of interest are captured by β_m. The sample includes the same control firms as the main analysis. The analysis employs inverse probability weighting based on the pre-1895 patenting distribution of individual constituent firms, rather than the consolidation-level aggregates used in the main specification.

Figure C4 plots the resulting event study estimates for the key outcomes of patents (panel a) and breakthroughs (panel b). Pre-trends are clearly flat.

Figure C3: Event Study Specification in Relative Time



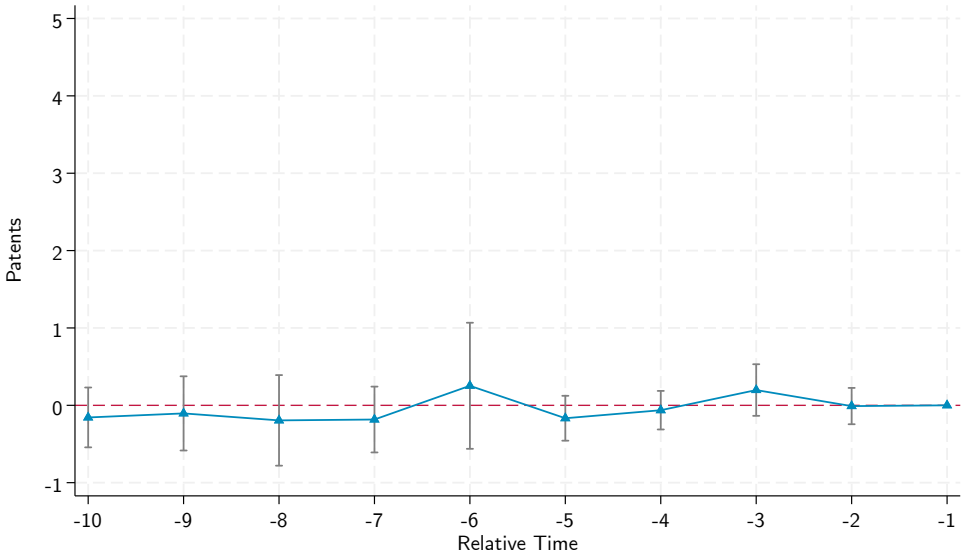
(a) All patents



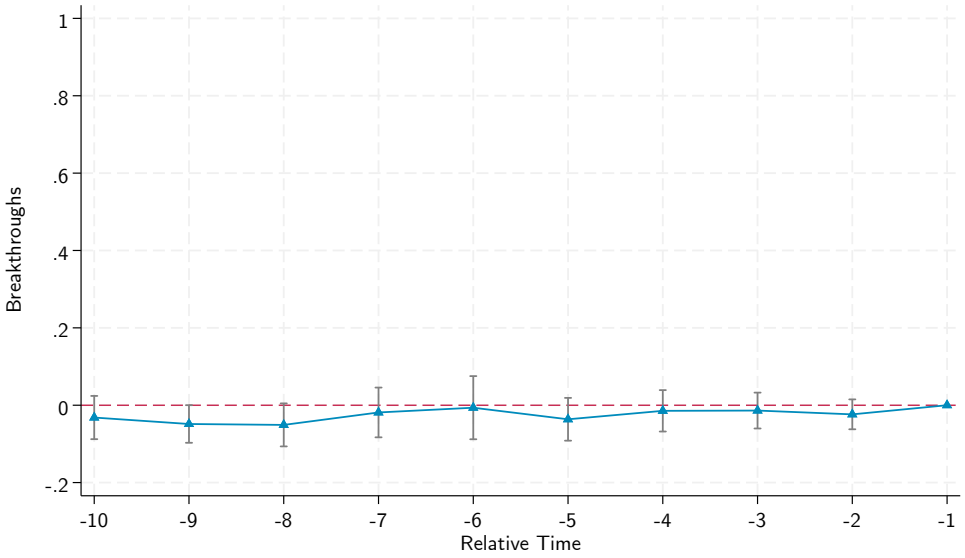
(b) Breakthrough patents

Note: The figure presents event-study estimates in relative time for the two intensive-margin outcomes, with event time anchored at each deal's earliest public announcement (Appendix B.4). Panel (a), total patents, reports coefficients β_m from equation (C4), estimated with the interaction-weighted estimator of Sun and Abraham (2021) to address treatment-time heterogeneity. Panel (b), breakthrough patents, reports a pooled two-way fixed-effects event study; for this sparse outcome the interaction-weighted estimator is unsuitable, so the pooled specification is reported instead (see below). IPW weights, sample, and fixed effects follow the main specification. Each panel displays β_m estimates and their 95 percent confidence intervals, with standard errors clustered at the firm level.

Figure C4: Pre-trends in Relative Time—Constituent Firms



(a) All patents



(b) Breakthrough patents

Note: This figure presents event study estimates from equation (C5), showing pre-merger innovation trends of constituent firms. The analysis disaggregates GMW consolidations to their component firms and examines innovation paths relative to each firm's specific merger date. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. The sample spans from 10 years before merger ($m = -10$) to one year before merger, with $m = -1$ as the omitted category. Each panel displays β_m estimates and their 95 percent confidence intervals. Standard errors are clustered at the firm level.

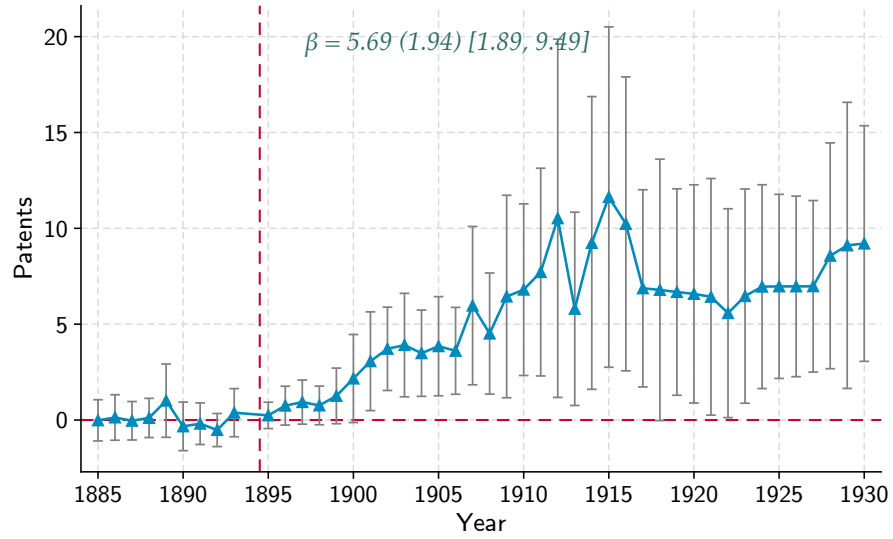
C.3 Robustness

Table C4: Product-Market Robustness—Alternative Rules and Fixed-Effects Variants

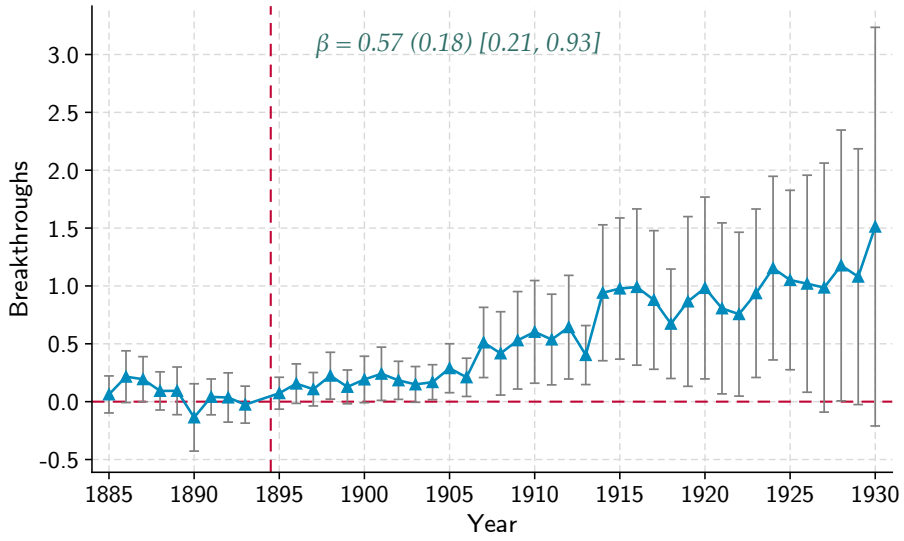
	(1) Top-1 subgroup	(2) Top-1 subclass	(3) Classified sample	(4) Industry × year	(5) Both FE sets
Panel A: Patents					
GMW Firm	6.240	6.303	5.895	5.520	5.707
	(1.983)	(1.984)	(1.928)	(1.905)	(1.935)
	[2.352, 10.127]	[2.413, 10.193]	[2.115, 9.675]	[1.786, 9.253]	[1.914, 9.500]
% of treated pre-mean	323%	327%	306%	286%	296%
Treated pre-mean	1.929	1.929	1.929	1.929	1.929
Panel B: Breakthroughs					
GMW Firm	0.656	0.629	0.573	0.531	0.568
	(0.199)	(0.193)	(0.184)	(0.175)	(0.183)
	[0.266, 1.046]	[0.250, 1.008]	[0.211, 0.934]	[0.188, 0.875]	[0.209, 0.927]
% of treated pre-mean	624%	599%	545%	506%	541%
Treated pre-mean	0.105	0.105	0.105	0.105	0.105
Controls dropped (%)	4.1%	33.6%	–	–	–
CPC section × year FE	Y	Y	Y	N	Y
Industry × year FE	N	N	N	Y	Y
# Treated	79	79	79	79	79
# Controls	11,244	7,788	5,795	5,795	5,795
N	520,858	361,882	270,204	270,204	270,204

Note: This table complements Table 2. Columns 1–2 exclude control firms whose 1885–94 patenting overlaps a treated niche under looser and stricter rules than the main-group rule of the text: any patent in a consolidation's largest CPC subgroup (column 1) and a modal subclass equal to a consolidation's largest subclass (column 2). Columns 3–5 restrict to firms with a name-based industry classification and vary the fixed effects as indicated. Table B5 in Appendix B reports the full menu of exclusion rules; rules based on the top two or three subgroups are infeasible, since the nine control firms with more than 75 pre-period patents all patent in some consolidation's top-two subgroups, emptying the largest reweighting stratum. Stratification weights are recomputed within each sample. Standard errors clustered at the firm level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Figure C5: Product-Market Robustness—Event Studies Under the Combined Specification



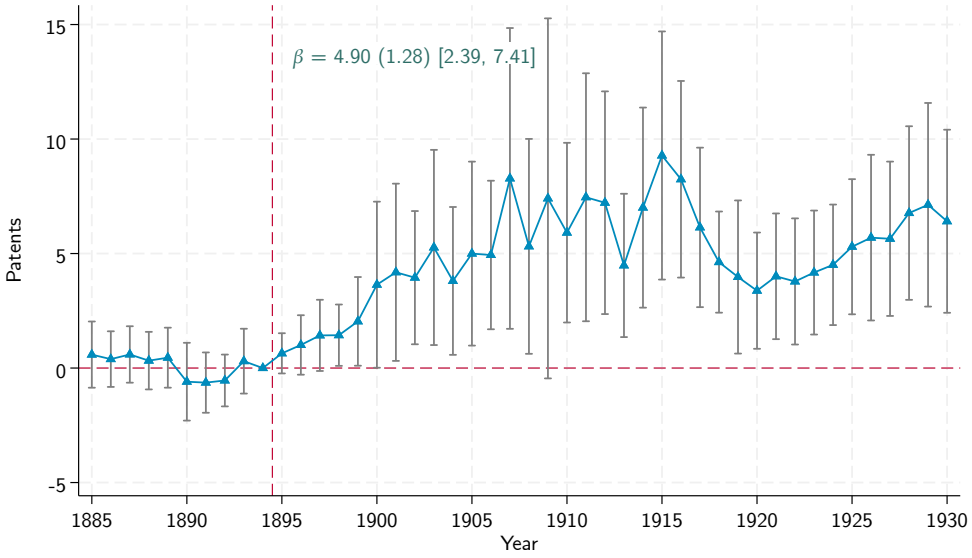
(a) Patents



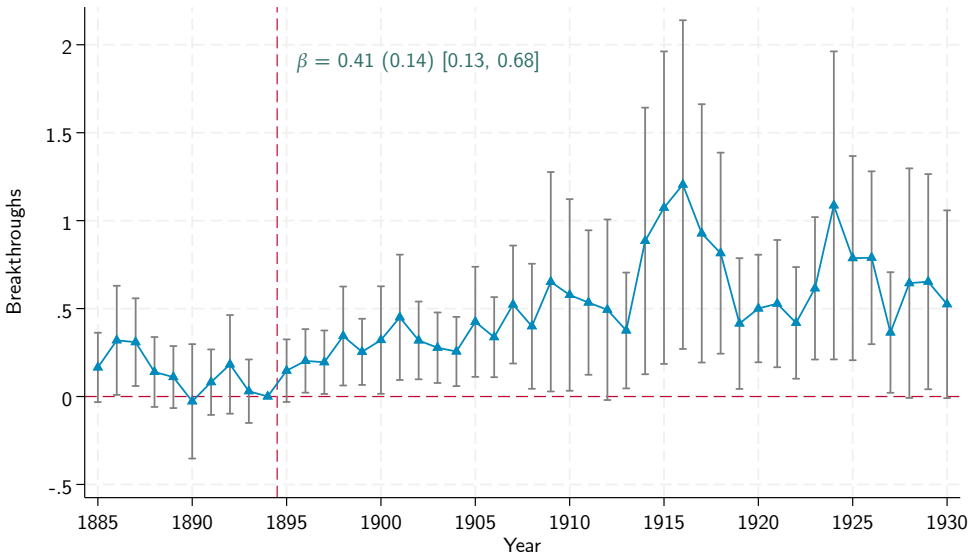
(b) Breakthroughs

Note: This figure presents event study estimates from equation (3) under the combined product-market specification of Table 2, column 3: control firms in a consolidation's largest pre-1895 CPC main group are excluded, the sample is restricted to firms with a name-based industry classification, and industry-by-year fixed effects are added to the technology-by-year effects. Each panel also reports β, the static estimate from equation (4), with its SE in parentheses and confidence interval in brackets. SEs are clustered at the firm level.

Figure C6: Horizontal consolidations



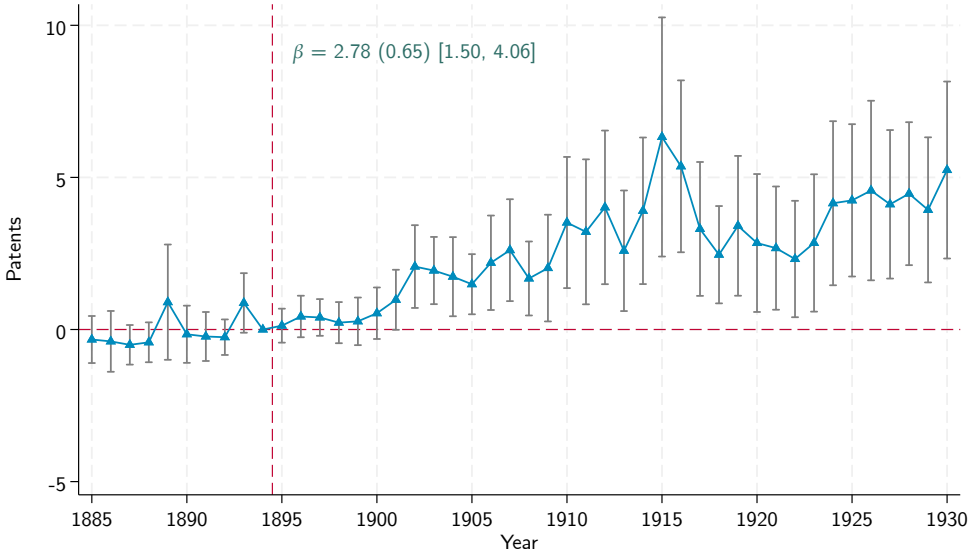
(a) Event study for all patents



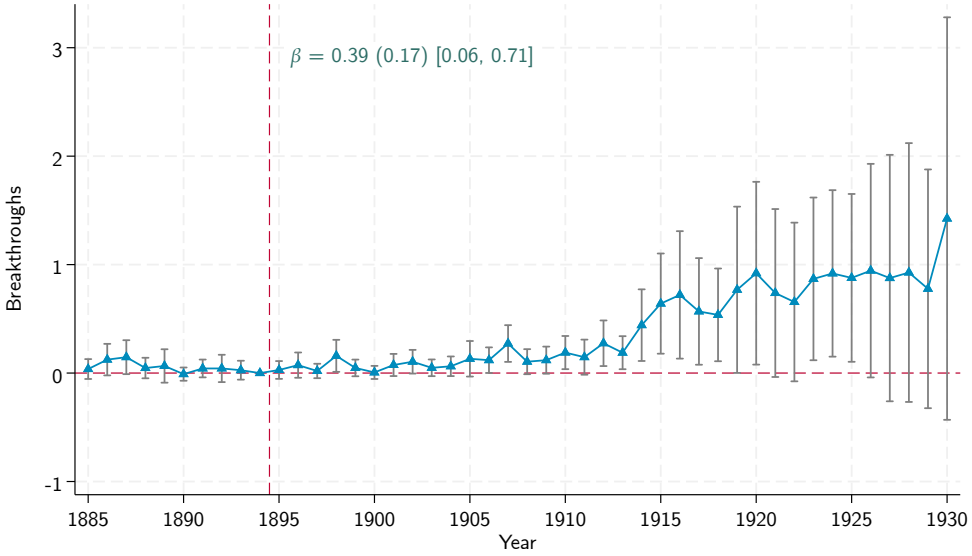
(b) Event study for breakthrough patents

Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on innovation, restricting the sample to consolidations classified as horizontal mergers. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. Each panel displays β_m estimates and their 95 percent confidence intervals. Standard errors are clustered at the firm level.

Figure C7: Trimmed sample



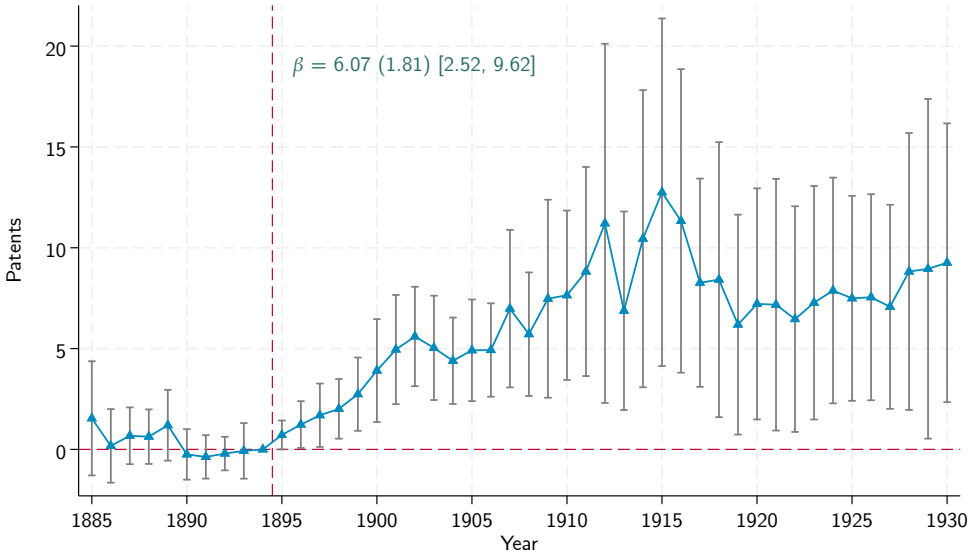
(a) Event study for all patents



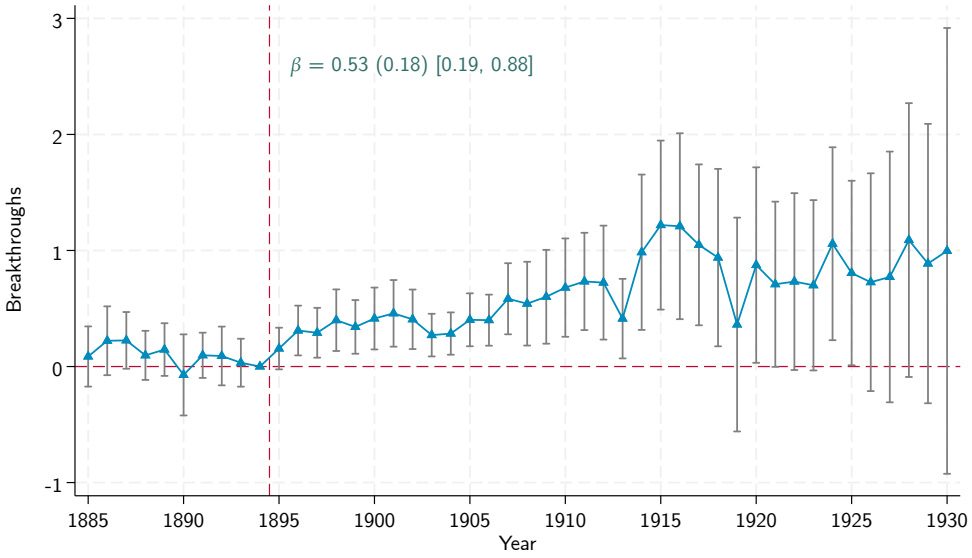
(b) Event study for breakthrough patents

Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on innovation using a trimmed sample that excludes top and bottom 5 percent of total patenting distributions by GMW status. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. Each panel displays β_m estimates and their 95 percent confidence intervals. Standard errors are clustered at the firm level.

Figure C8: Conditioning on Firm Survival—Intensive Margin



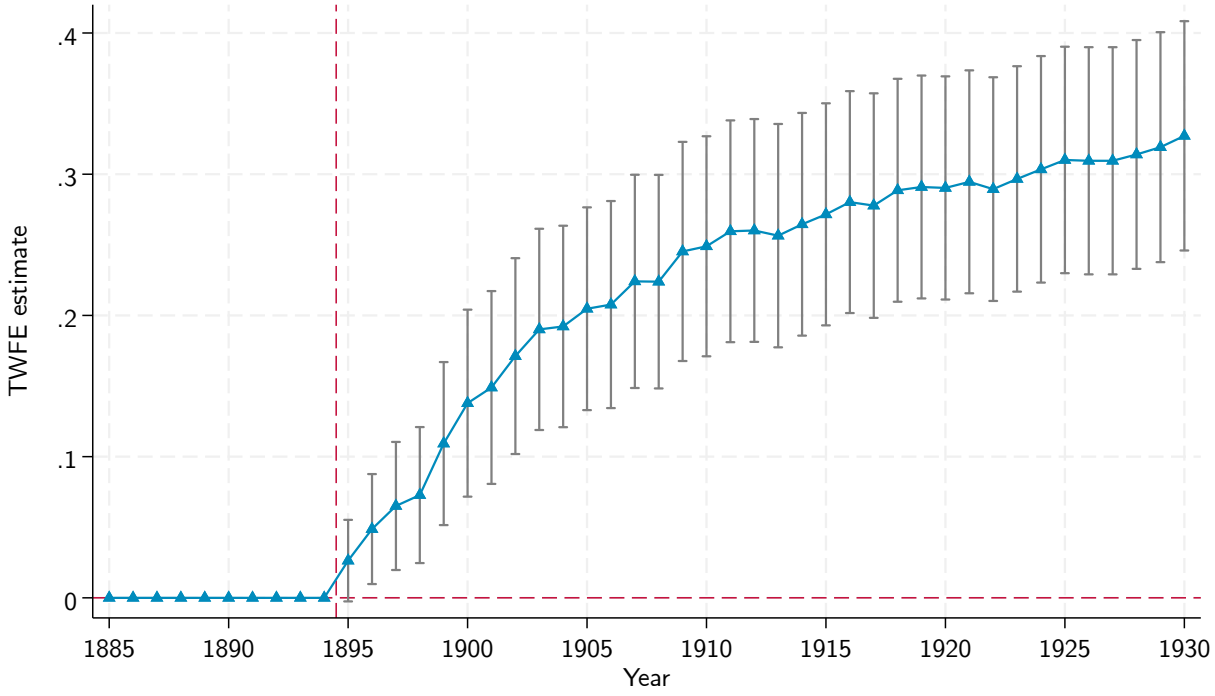
(a) Event study for all patents



(b) Event study for breakthrough patents

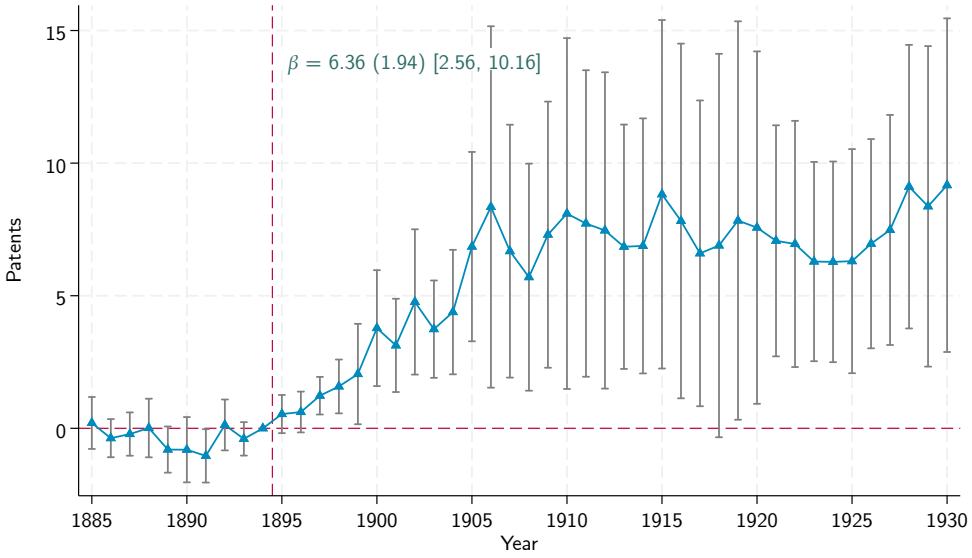
Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on innovation using an unbalanced panel where I only keep observations within three years of a firm's first and last patent. This also allows me to include consolidations that were later bought by other firms. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. Each panel displays β_m estimates and their 95 percent confidence intervals. Standard errors are clustered at the firm level.

Figure C9: Including later acquired consolidations (unbalanced sample)—Extensive Margin

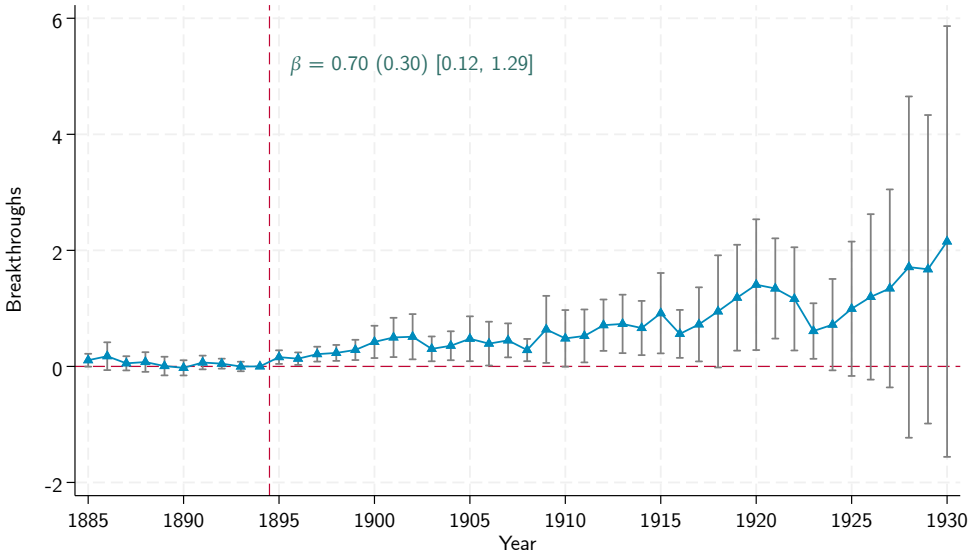


Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on the probability of beginning to patent, using an unbalanced panel that includes consolidations which were themselves acquired by other firms during the sample period. The figure displays β_m estimates and their 95 percent confidence intervals for firms that had not patented prior to 1895. Standard errors are clustered at the firm level.

Figure C10: Filing year—Intensive Margin



(a) Event study for all patents



(b) Event study for breakthrough patents

Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on innovation using patent filing dates rather than grant dates. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. Each panel displays β_m estimates and their 95 percent confidence intervals. Standard errors are clustered at the firm level.

Figure C11: Filing year—Extensive Margin

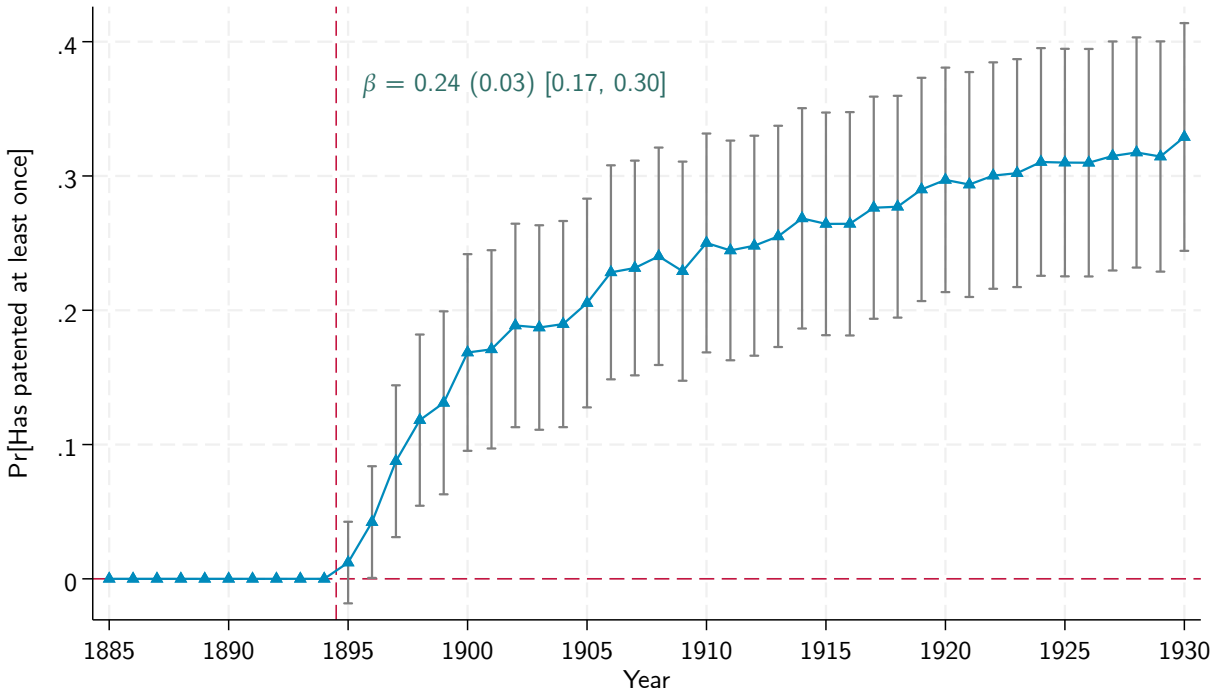


Figure C12: Event study estimates

Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on the probability of beginning to patent, using patent filing dates rather than grant dates. The figure displays β_m estimates and their 95 percent confidence intervals for firms that had not patented prior to 1895. Standard errors are clustered at the firm level.

Table C5: Effect of Consolidation on Firm-Level Patenting—Filing year

Margin:	Intensive			Extensive	
Outcome:	Patents	Breakthroughs	Has lab	Started patenting	Has lab
	(1)	(2)	(3)	(4)	(5)
GMW Firm	6.363	0.703	0.158	0.235	0.045
	(1.939)	(0.298)	(0.032)	(0.034)	(0.014)
	[2.562, 10.164]	[0.119, 1.286]	[0.096, 0.220]	[0.168, 0.303]	[0.016, 0.073]
Controls	Y	Y	Y	Y	Y
Treated pre-mean	1.884	0.093	0.017	0.000	0.001
Control post-mean	0.665	0.060	0.017	0.134	0.011
# Firms	12,156	12,156	12,156	834	834
N	559,176	559,176	559,176	38,310	38,310

Note: This table reports static difference-in-differences estimates of the effect of consolidation on firm-level innovation outcomes (Equation 4), using patent filing year rather than issue year. Notice that firms are resorted between intensive and extensive margin once pre-1895 patenting is measured with filing dates. patenting Columns 1-3 show results for the intensive margin (firms with pre-merger patents), while columns 4-5 show extensive margin effects (firms with no pre-merger patents). The dependent variables are: patent counts (col. 1), breakthrough patent counts (col. 2), probability of having at least one R&D lab (col. 3 and col. 5), probability of patenting at least once (col. 4). "GMW Firm" denotes the post-1895 interaction $\mathbb{1}[t \geq 1895] \times \text{GMW Firm}_i$ from equation (4); the firm fixed effects absorb the level consolidation indicator. Intensive margin regressions include time fixed effects by technological area, while extensive margin include time fixed effects by economic sector. Standard errors clustered at the firm level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

C.4 Paper Mergers: Robustness

Tables C8 and C9 probe the comparison between completed consolidations and paper mergers of Section 4.3. Table C8 reports the two-way fixed effects post-announcement estimate across thirteen specifications: alternative fixed effects, retaining absorbed constituents' post-absorption years, excluding influential deals one at a time (including the three largest patenters), excluding deals with weaker inclusion evidence or announcements before 1895, using Nelson's dating, and excluding constituent firms that publicly refused their deals. The estimate is positive throughout, ranging from about +1.1 to +2.4 patents per deal-year around a baseline of +2.0 (wild-cluster bootstrap p-value 0.023); these two-way fixed effects estimates complement the interaction-weighted specification reported in the body. Table C9 splits the comparison by the documented cause of abandonment. The estimates are stable across causes, though target-holdout deals hold nearly all of the pool's patents, so the split is descriptive rather than a separate design.

Table C6: Consolidated Firms' Innovation Around Announcement—Within-Firm Estimates

	(1) Patents	(2) Breakthroughs
Avg. post-announcement effect	2.022	0.225
	(0.694)	(0.067)
	[0.663, 3.382]	[0.094, 0.356]
Pre-announcement mean	0.826	0.020
# Consolidations	232	232
N	10,672	10,672

Note: This table reports within-firm event-study estimates of consolidated firms' innovation around the earliest public announcement of their consolidation (Equation 1), pooling all GMW consolidations with no control group. The reported effect is the average of the post-announcement event-time coefficients (years 0 through 25), with a standard error computed from the full covariance matrix and clustered at the consolidation level. The pre-announcement mean is the average outcome over the five years before announcement. Column 1 reports patents and column 2 breakthrough patents. 95 percent confidence intervals are reported in square brackets. Because the specification includes only consolidation fixed effects, the estimates are descriptive and net out no secular or technology-specific trends.

C.5 Estimator Robustness: Fixed-Effects Poisson

The specifications in Section 4 estimate treatment effects in levels. Table C10 re-estimates the intensive-margin patent effect with a fixed-effects Poisson model (Correia, Guimarães, and Zylkin 2020), in calendar time and in time relative to announcement. The two estimators answer different questions, and their magnitudes are not directly comparable: the linear model identifies an average effect in levels under parallel trends in levels—approximately 310 percent of the pre-wave mean—whereas the Poisson model identifies a proportional effect under parallel trends in growth rates, implying a larger multiple because treated firms start from a low base while control firms mean-revert. A model-free

Table C7: Completed Consolidations versus Paper Mergers—Deal-Level Estimates

	(1) Patents	(2) Breakthroughs
Avg. post-announcement effect	2.217	0.222
	(0.884)	(0.067)
	[0.484, 3.950]	[0.090, 0.354]
Pre-announcement mean	0.768	0.024
# Completed consolidations	265	265
# Abandoned deals	24	24
N	10,016	10,016

Note: This table compares completed consolidations to abandoned "paper" mergers at the deal level, in time relative to each deal's earliest public announcement (Equation 2). The reported effect is the average of the post-announcement event-time coefficients (years 0 through 25). Column 1 (patents) uses the interaction-weighted estimator of Sun and Abraham (2021), with a standard error computed from its covariance matrix; column 2 (breakthroughs) uses two-way fixed effects, since the interaction-weighted standard errors collapse for so sparse an outcome. The pre-announcement mean is the average outcome over the years before announcement. Standard errors are clustered at the deal level; 95 percent confidence intervals are reported in square brackets.

Table C8: Completed consolidations versus paper mergers: robustness

Specification	Post×Consolidated		Deals
Baseline (deal + event-time fixed effects)	+1.94	(0.88)	24
Sector × event-time fixed effects	+1.48	(0.71)	24
Excluding deals with weaker inclusion evidence	+1.91	(0.97)	20
Excluding the National Screw attempt	+1.49	(0.77)	23
Excluding the American Plow attempt	+2.06	(0.89)	23
Excluding the street-railway car attempt	+2.34	(0.80)	23
Excluding the three largest abandoned deals	+2.02	(0.65)	21
Excluding deals announced before 1895	+1.94	(0.93)	21
Excluding the National Saw consolidation	+1.95	(0.88)	24
Retaining post-acquisition consolidation years	+1.82	(0.85)	24
Nelson dating for completed consolidations	+1.94	(0.88)	24
Excluding constituents that publicly refused	+1.60	(0.77)	24

Note: Each row reports the post-announcement difference-in-differences coefficient (patents per deal-year) from a two-way fixed effects specification comparing completed consolidations to abandoned deals at the deal level, window -10 to $+25$ around announcement; SEs in parentheses are clustered at the deal level. "Deals" counts the abandoned deals entering each specification.

comparison of group-level growth ratios lies between the two. I do not report Poisson estimates for breakthroughs, where most firms have all-zero outcomes and the conditional likelihood selects the sample, nor for the extensive margin, where the outcome is an absorbing indicator. The announcement-anchored relative-time Poisson column warrants more caution than the others, since its pre-trends are jointly significant, unlike

Table C9: Completed consolidations versus paper mergers, by cause of abandonment

Cause of abandonment	Deals	Patents	DiD: patents		DiD: breakthroughs	
Target holdout or refusal	9	1197	+1.76	(1.70)	+0.17	(0.08)
Exogenous shock	3	6	+2.18	(0.66)	+0.25	(0.06)
Financing or promotion breakdown	9	5	+2.02	(0.65)	+0.25	(0.07)
Unknown / silent	3	10	+2.02	(0.66)	+0.25	(0.07)
All abandoned deals	24	1218	+1.94	(0.88)	+0.22	(0.07)

Note: Each row restricts the abandoned-deal group to deals with the stated cause of abandonment and re-estimates the two-way fixed effects comparison for patents and breakthroughs per deal-year; SEs in parentheses are clustered at the deal level. Patent counts are totals over 1885–1940 after the absorption sample restriction.

those of the calendar-time Poisson and both linear specifications.[52]

Table C10: Linear and fixed-effects Poisson estimates of the intensive-margin patent effect

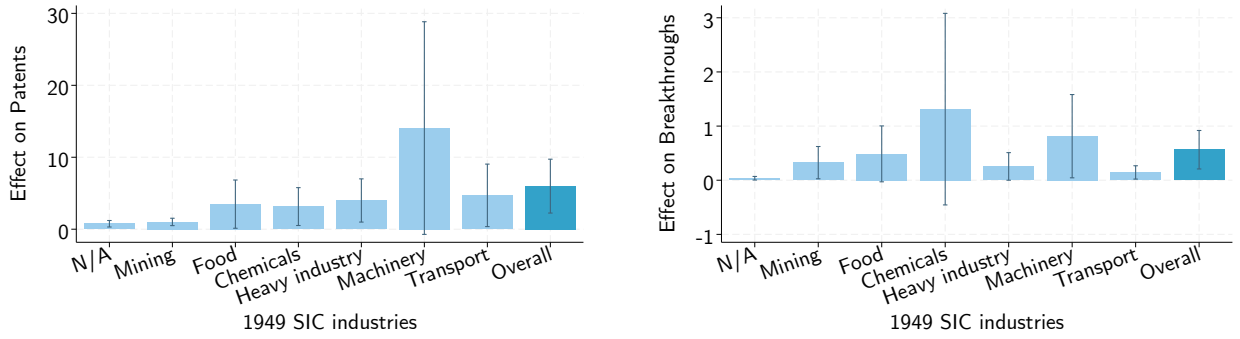
Clock:	Calendar time $\mathbf{1}[t \geq 1895]$ (1)	Relative to announcement (2)
Linear (levels)	5.989	6.485
	(1.910)	(2.103)
% of pre-treatment mean	+310%	+328%
FE-Poisson (log link)	2.626	2.759
	(0.367)	(0.433)
implied ratio e^{β}	×13.8	×15.8

Note: Difference-in-differences estimates for patents among intensive-margin firms, IPW weights and clustering as in the main specification. The linear coefficient is an average effect in patents per firm-year; the Poisson coefficient is on the log scale, with the implied proportional effect e^{β} reported below it. Columns differ in the event clock: calendar time ($t \geq 1895$) and time relative to each deal's earliest public announcement.

C.6 Heterogeneity

52. The joint test of the relative-time Poisson pre-trends rejects in the announcement-anchored specification (p-value 0.005); the calendar-time specification rejects for neither estimator.

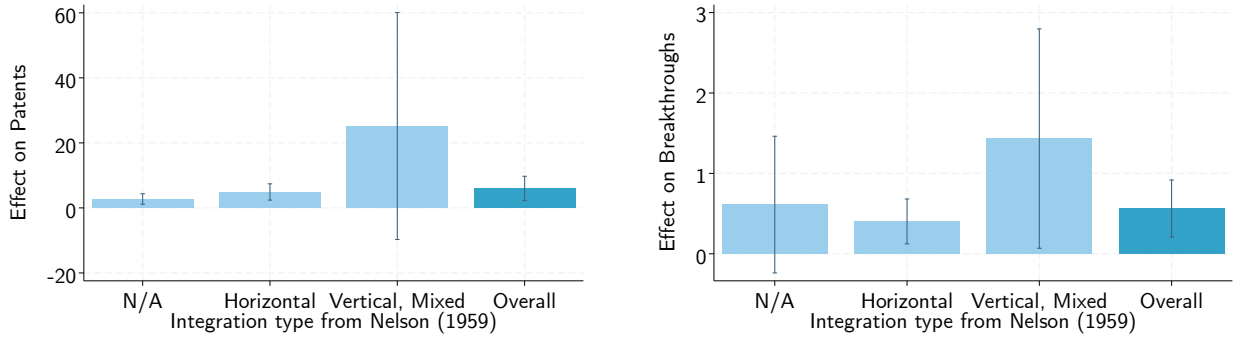
Figure C13: Heterogeneity in Consolidation Effects by Sector



(a) All patents (b) Breakthrough patents

Note: This figure shows how the effect of consolidation on innovation varies across manufacturing sectors defined by 1949 SIC industry codes, as reported in Nelson (1959). Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. Each bar shows the estimated effect of consolidation with error bars representing 95 percent confidence intervals computed from standard errors clustered at the firm level.

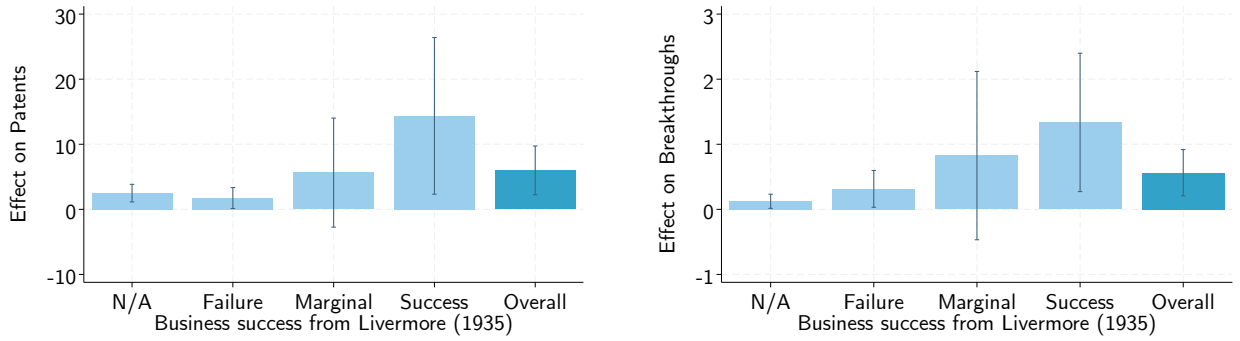
Figure C14: Heterogeneity in Consolidation Effects by Integration Type



(a) All patents (b) Breakthrough patents

Note: This figure shows how the effect of consolidation on innovation varies by the type of integration, as reported in Nelson (1959). Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. Each bar shows the estimated effect of consolidation with error bars representing 95 percent confidence intervals computed from standard errors clustered at the firm level.

Figure C15: Heterogeneity in Consolidation Effects by Success (Livermore 1935)



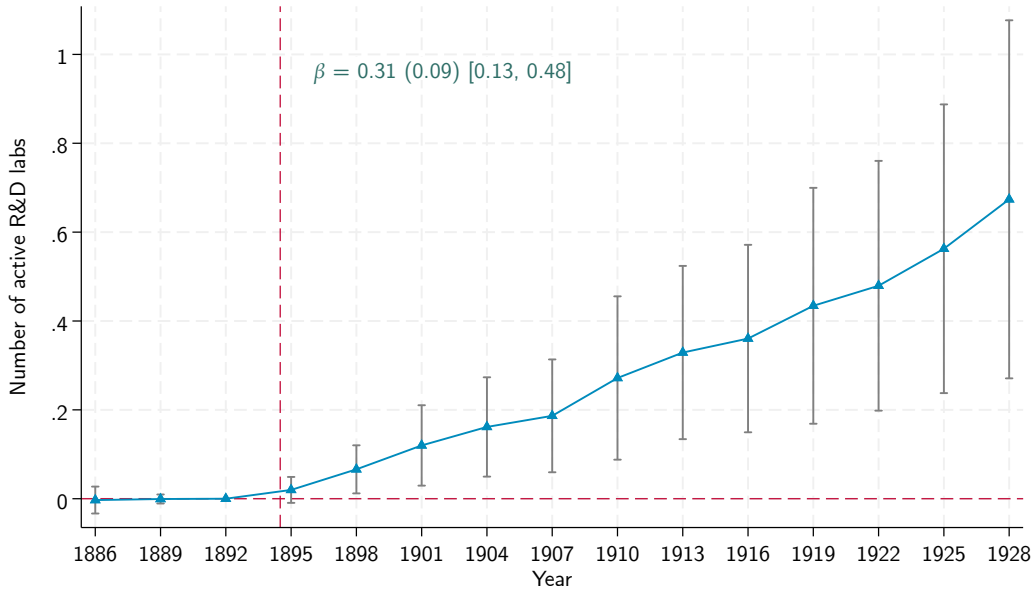
(a) All patents

(b) Breakthrough patents

Note: This figure shows how the effect of consolidation on innovation correlates with long-term business success, as classified by Livermore (1935). Firms are categorized as successful, moderately successful, or unsuccessful based on their subsequent business performance. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. Each bar shows the estimated effect of consolidation with error bars representing 95 percent confidence intervals computed from standard errors clustered at the firm level.

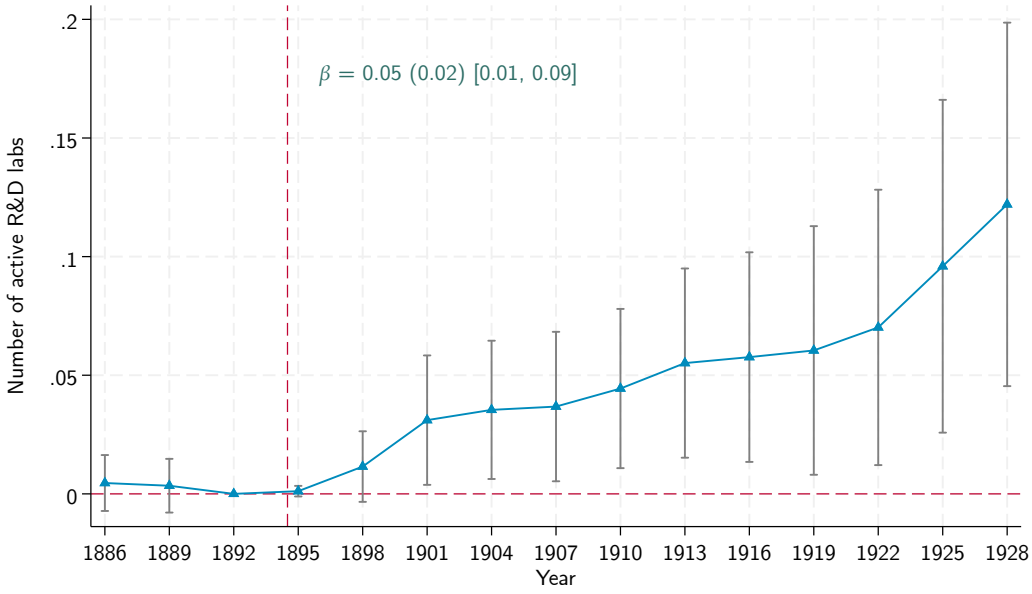
C.7 R&D Labs

Figure C16: Number of active labs — intensive margin firms



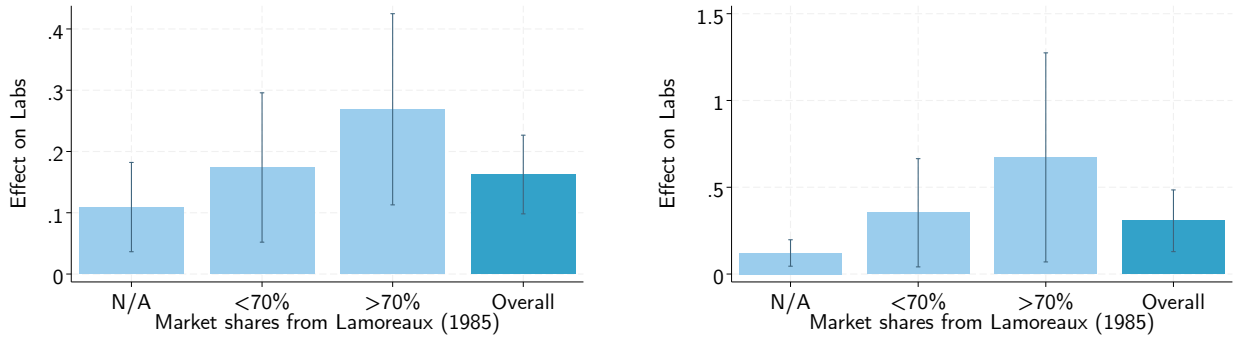
Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on the number of active R&D laboratories for firms that had patented prior to 1895. Observations are grouped in 3-year bins due to limited variation in laboratory establishment. The figure displays β_m estimates and their 95 percent confidence intervals. The x-axis reports the first year in each bin. Standard errors are clustered at the firm level.

Figure C17: Number of active labs — extensive margin firms



Note: This figure presents event study estimates from equation (3) showing the firm-level effect of consolidation on the number of active R&D laboratories for firms that had not patented prior to 1895. Observations are grouped in 3-year bins due to limited variation in laboratory establishment. The figure displays β_m estimates and their 95 percent confidence intervals. The x-axis reports the first year in each bin. Standard errors are clustered at the firm level.

Figure C18: Heterogeneity in Consolidation Effects by Market Concentration

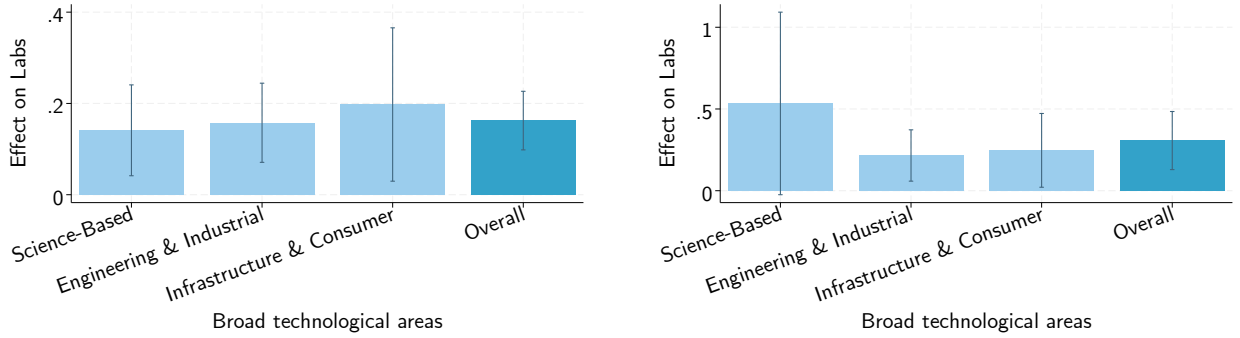


(a) Probability have having a lab

(b) Number of active labs

Note: This figure shows how the effect of consolidation on R&D laboratory establishment varies with market concentration. Firms are categorized into three groups: those without reported market share information (likely smaller and less successful mergers), those achieving substantial but less than 70% market share, and those capturing over 70% market share. Market share data comes from Lamoreaux (1985). Panel (a) shows the estimated effect on probability of having at least one lab, Panel (b) shows the estimated effect on number of active labs. Error bars represent 95 percent confidence intervals computed from standard errors clustered at the firm level.

Figure C19: Heterogeneity in Consolidation Effects by Broad Technological Area



(a) Probability have having a lab

(b) Number of active labs

Note: This figure shows how the effect of consolidation on R&D laboratory establishment varies across broad technological areas. Firms are categorized into three groups according to the CPC section where they patent the most: (i) sections C, G and H for science-based technologies, (ii) sections B, F and Y for engineering and industrial technologies, (iii) sections A, D and E for infrastructure and consumer-oriented technologies. Panel (a) shows the estimated effect on probability of having at least one lab, Panel (b) shows the estimated effect on number of active labs. Error bars represent 95 percent confidence intervals computed from standard errors clustered at the firm level.

C.8 Other descriptive statistics on GMW consolidations

Table C11: Additional Descriptive Statistics on GMW Firms

Year	Stopped Patenting (%)		Has Active & Dated Lab (%)		Has Any Lab (%)	
	GMW	Non-GMW	GMW	Non-GMW	GMW	Non-GMW
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Intensive Margin Firms						
1900	5.1	67.6	10.1	1.3		
1910	7.6	72.6	17.7	1.6		
1920	13.9	78.3	26.6	2.2		
1930	24.1	85.5	34.2	2.6		
By 1946					54.1	4.7
Panel B: Extensive Margin Firms						
1900	3.6	6.4	1.3	0.2		
1910	19.3	20.8	4.6	0.4		
1920	36.1	41.0	7.2	1.0		
1930	47.0	65.8	12.4	2.0		
By 1946					20.3	4.4

Note: This table presents additional descriptive statistics on R&D laboratory establishment and firm exit. In columns 1 and 2, the outcome 'Stopped Patenting' is computed for firms that ever patent in the sample period, and records whether the firm has been issued its last patent before a given year. Because later patenting is itself an outcome affected by consolidation, this measure reflects the continuation of patenting as much as true firm exit, and I do not read it as an independent measure of firm survival. Columns 3 and 4 report laboratories with non-missing establishment dates, while columns 5 and 6 show an indicator for whether the firm ever appears in the NRC surveys before 1946 (regardless of establishment date). Results are shown separately for intensive margin firms (those with pre-merger patents) and extensive margin firms (those without pre-merger patents), and GMW status.

D Alternative inference approaches

This appendix presents results from three alternative inference approaches: placebo mergers, synthetic control, and synthetic difference-in-differences. These methods create synthetic counterfactuals that aggregate multiple control units to more closely mirror actual consolidation structure.

Figures D1, D2, and D3 present event study estimates consistent with the main findings. Bayesian Model Averaging combines evidence across all four approaches to account for model uncertainty (Table D1).

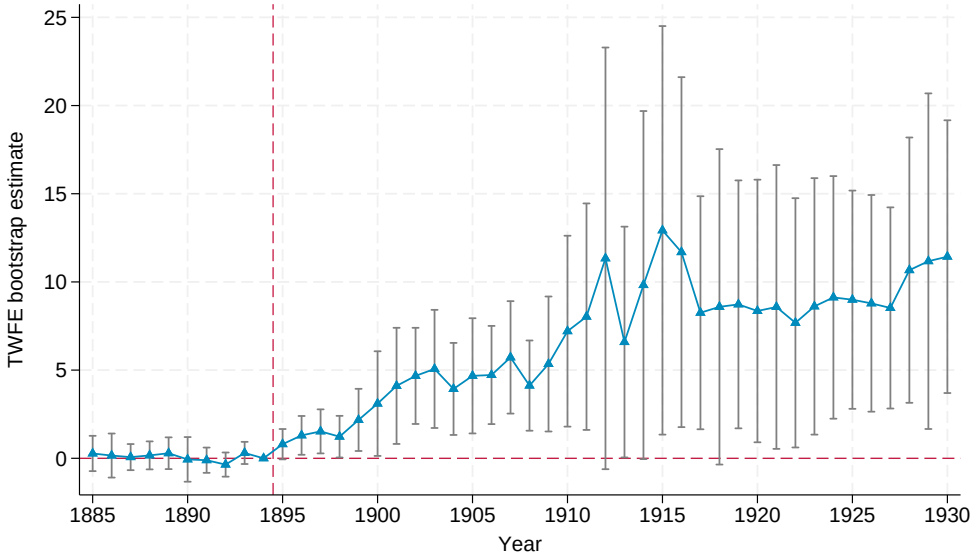
D.1 Results

The three alternative approaches each construct counterfactuals that aggregate multiple control units, more closely resembling the structure of treatment units. The placebo merger approach randomly groups control firms within technological areas to match treatment units' pre-merger innovation levels, then fits the same specification as Equation 3. Standard errors that account for variability in placebo construction are computed via bootstrap (1,000 iterations). The synthetic control method constructs optimal weighted combinations of control firms to match each treatment unit's pre-merger trajectory using constrained optimization. Synthetic difference-in-differences combines elements of both difference-in-differences and synthetic control by assigning weights to both units and time periods.

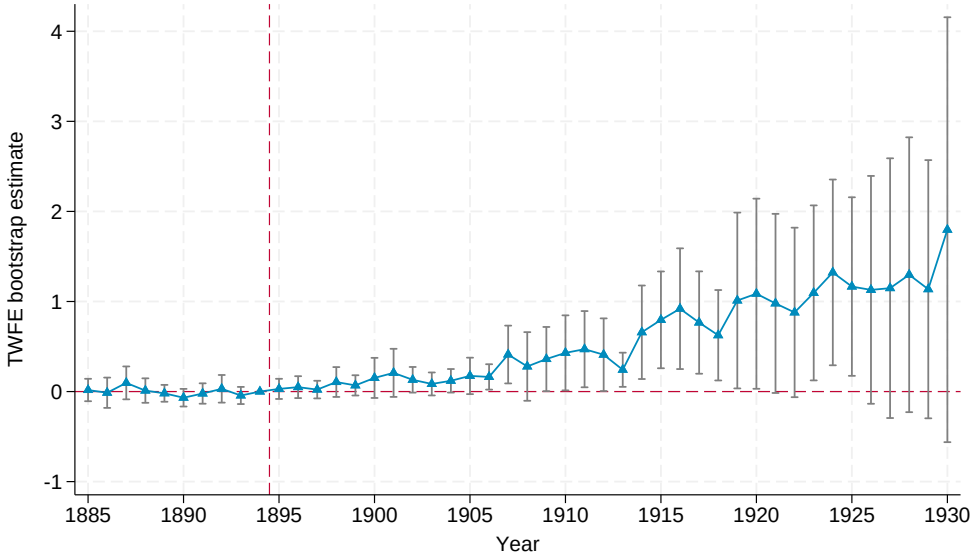
For the synthetic control and SDID approaches, outcomes are first residualized with respect to year-by-technological-area fixed effects to control for technology-specific trends outside their specific optimization routines (Arkhangelsky et al. 2021).

Figures D1, D2, and D3 present event study estimates from each approach. All three methods yield results that closely align with the main difference-in-differences findings.

Figure D1: GMW Firms vs. Placebo Mergers



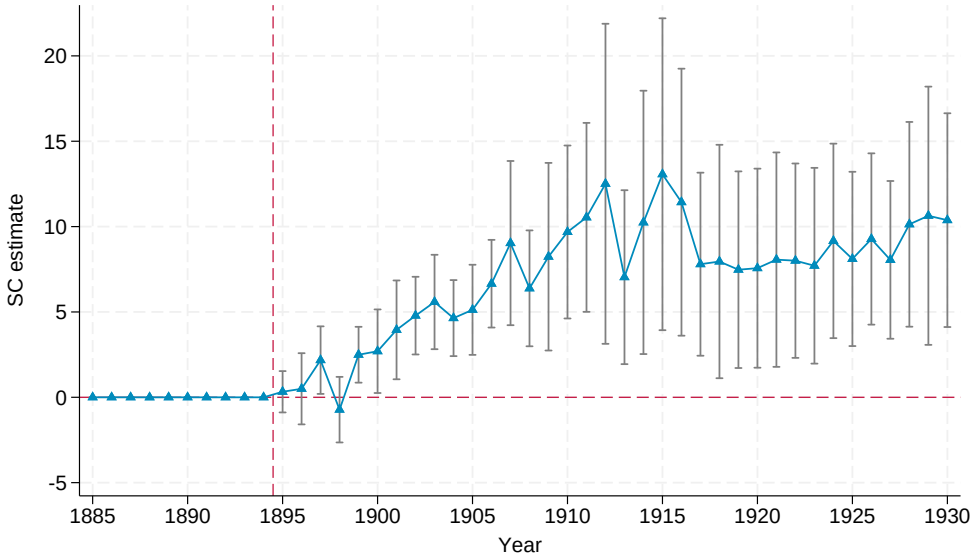
(a) Event study for all patents



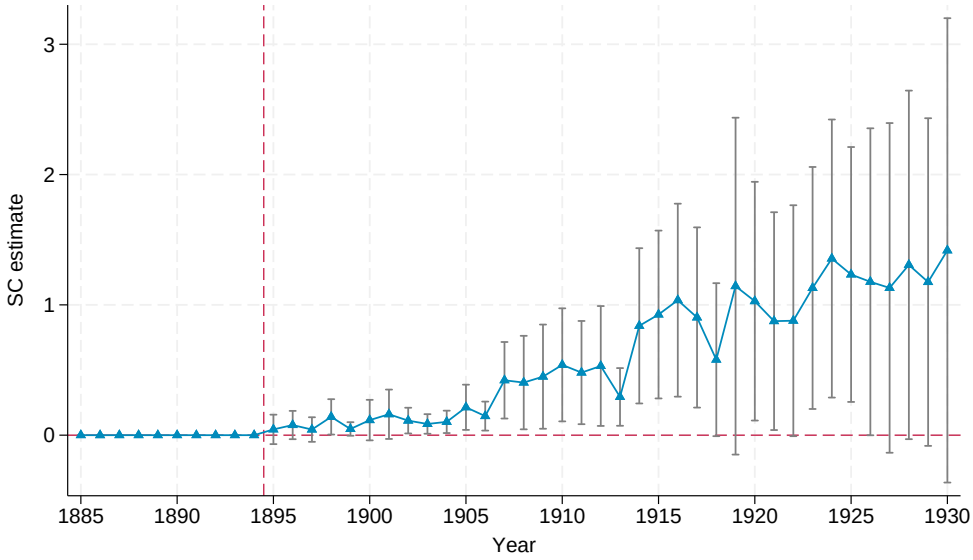
(b) Event study for breakthrough patents

Note: This figure presents event study estimates comparing actual GMW consolidations to placebo mergers constructed from control firms. Within each technological area, control firms are randomly selected to match treatment units' pre-merger innovation levels, creating synthetic consolidations. Results are based on 1,000 bootstrap iterations. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. The analysis shows that randomly constructed mergers fail to generate the sustained innovation increases observed for actual GMW consolidations, supporting the causal interpretation of the main results.

Figure D2: SC mergers



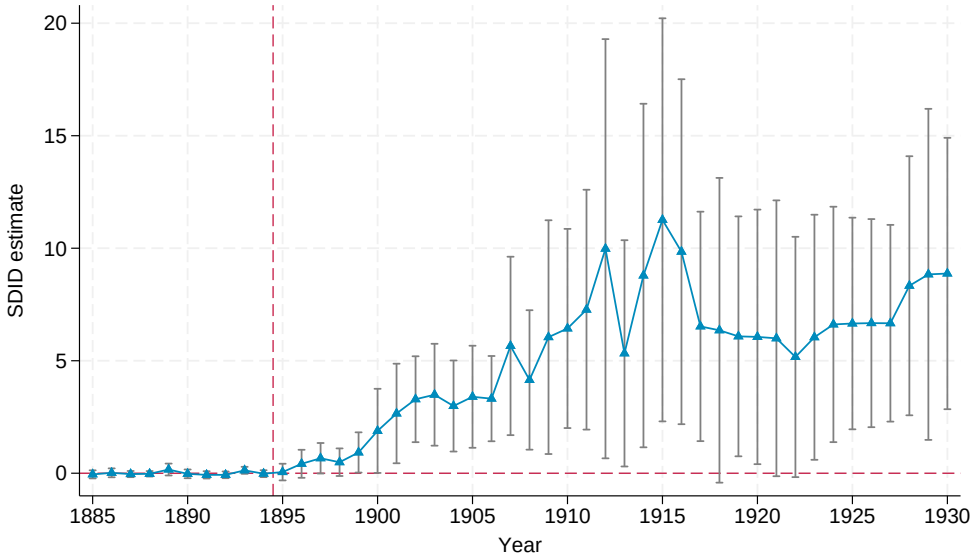
(a) Event study for all patents



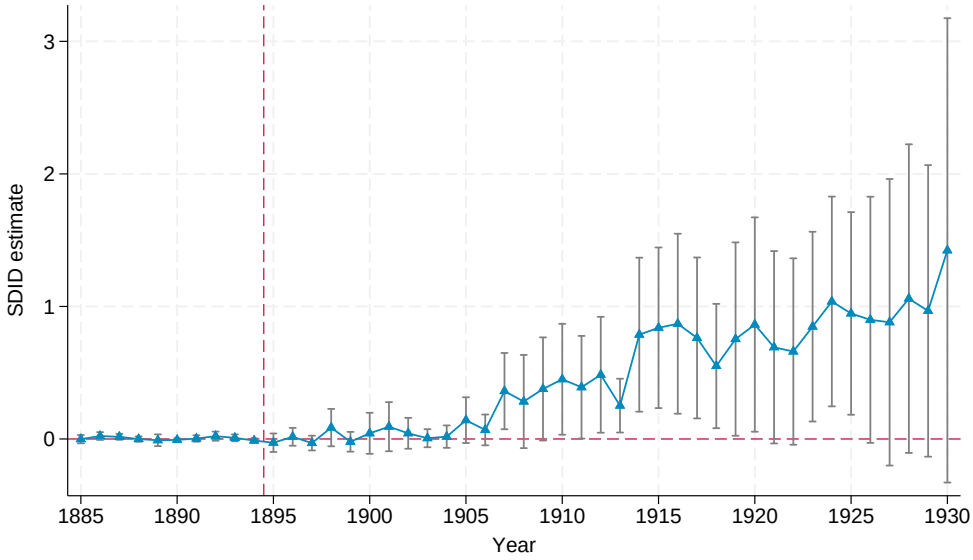
(b) Event study for breakthrough patents

Note: This figure presents event study estimates from synthetic control analysis following Abadie (2021). Optimal weighted combinations of control firms are constructed to match each treatment unit's pre-merger innovation trajectory. Outcomes are residualized with respect to year-by-technological-area fixed effects before applying the synthetic control procedure. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. The results closely align with the main difference-in-differences findings, confirming substantial post-merger innovation increases.

Figure D3: SDID mergers



(a) Event study for all patents



(b) Event study for breakthrough patents

Note: This figure presents event study estimates from synthetic difference-in-differences analysis following Arkhangelsky et al. (2021). This method assigns weights to both units and time periods, combining elements of difference-in-differences and synthetic control. Outcomes are residualized with respect to year-by-technological-area fixed effects before implementing SDID. Panel (a) shows results for total patents, Panel (b) shows results for breakthrough patents. The approach confirms substantial post-merger innovation increases consistent with alternative specifications.

D.2 Bayesian Model Averaging

Bayesian Model Averaging (BMA) is a framework for combining results from multiple competing models while explicitly acknowledging model uncertainty. Rather than selecting a single 'best' specification, BMA treats each empirical approach as a plausible representation of the underlying data-generating process and combines their estimates using posterior model probabilities as weights.

General BMA Framework. In the standard BMA framework, the posterior distribution of a parameter of interest θ given data D is computed as a weighted average across M competing models:

$$P(\theta|D) = \sum_{m=1}^{M} P(\theta|M_m, D) \cdot P(M_m|D) \tag{D1}$$

where $P(\theta|M_m, D)$ represents the posterior distribution of θ under model M_m, and $P(M_m|D)$ denotes the posterior probability of model M_m given the observed data. The posterior model probabilities are derived using Bayes' rule:

$$P(M_m|D) = \frac{P(D|M_m) \cdot P(M_m)}{\sum_{j=1}^{M} P(D|M_j) \cdot P(M_j)} \tag{D2}$$

where $P(D|M_m)$ is the marginal likelihood (or model evidence) and $P(M_m)$ represents the prior probability assigned to model M_m. The marginal likelihood measures how well model M_m predicts the observed data, integrating over the uncertainty in model parameters.

RMSE-Based Approximation. Computing exact marginal likelihoods across the heterogeneous estimation approaches employed here—difference-in-differences, placebo analysis, synthetic control, and synthetic difference-in-differences—presents significant computational challenges. Traditional approaches using Bayesian Information Criteria are not readily applicable given the distinct methodological frameworks and non-standard procedures.

I therefore employ a simplified, *ad hoc* weighting based on overall in-sample root mean squared error (RMSE) as a predictive-fit summary. This approach treats RMSE as an inverse measure of model fit, with the weight on model m proportional to:

$$P(M_m|D) \propto \frac{1}{\text{RMSE}_m} \tag{D3}$$

After normalization across all models, the posterior probability becomes:

$$P(M_m|D) = \frac{1/\text{RMSE}_m}{\sum_{j=1}^{4} 1/\text{RMSE}_j} \tag{D4}$$

This approximation is reasonable for several reasons. First, RMSE provides a natural measure of predictive accuracy that is comparable across the different estimation procedures. Second, models with lower prediction errors should intuitively receive greater weight in the averaging process, which this approach achieves. Third, RMSE is loosely related to model fit under Gaussian errors, where the log marginal likelihood declines

in the residual sum of squares. These weights are best read as an *ad hoc* predictive-fit summary rather than exact posterior model probabilities, which would also depend on the likelihood scale, sample size, priors, and complexity penalties; the exercise is a robustness check on the averaged estimates, not a formal Bayesian model comparison.

BMA Estimation. The BMA point estimate combines individual model estimates using posterior probabilities as weights:

$$\hat{\beta}_{BMA} = \sum_{m=1}^{4} \hat{\beta}_m \cdot P(M_m|D) \tag{D5}$$

The BMA variance accounts for both within-model parameter uncertainty and between-model uncertainty arising from the spread of estimates across specifications:

$$\text{Var}(\hat{\beta}_{BMA}) = \sum_{m=1}^{4} P(M_m|D) \left[\text{Var}(\hat{\beta}_m|M_m, D) + (\hat{\beta}_m - \hat{\beta}_{BMA})^2\right] \tag{D6}$$

The first term within brackets captures the parameter uncertainty conditional on each model, while the second term reflects the uncertainty about which model best represents the true data-generating process.

Table D1 presents the BMA results alongside individual model estimates. The posterior weights reveal that all four approaches receive substantial support from the data, though the IPW and Placebo difference-in-differences receive the most. The BMA estimates of 6.35 additional patents and 0.58 additional breakthroughs per firm per year align closely with the main inverse probability weighted results.

Table D1: BMA

Model:	IPW DiD	Placebo DiD	Synth. Control	Synth. DiD	BMA
Panel A: Patents					
GMW Firm	5.99	6.80	7.13	5.37	6.35
	(1.91)	(2.49)	(2.00)	(1.91)	(2.23)
	[2.24, 9.73]	[1.92, 11.69]	[3.20, 11.06]	[1.63, 9.11]	[1.97, 10.72]
BMA weight	0.380	0.380	0.112	0.128	—
Panel B: Breakthroughs					
GMW Firm	0.56	0.60	0.63	0.50	0.58
	(0.18)	(0.23)	(0.21)	(0.18)	(0.21)
	[0.21, 0.92]	[0.16, 1.05]	[0.21, 1.04]	[0.14, 0.85]	[0.17, 0.98]
BMA weight	0.348	0.376	0.128	0.147	—

Note: This table presents Bayesian Model Averaging results combining evidence from four identification strategies: inverse probability weighted difference-in-differences (IPW DiD), placebo mergers, synthetic control, and synthetic difference-in-differences. Model weights are computed from RMSE-based predictive fit (an *ad hoc* approximation, not exact posterior model probabilities). Panel A shows results for total patents, Panel B shows results for breakthrough patents. The BMA estimates incorporate both parameter uncertainty within specifications and model uncertainty across specifications. Standard errors in parentheses, 95 percent confidence intervals in brackets. BMA weights show the RMSE-based weight assigned to each approach.

E Additional results — Corporate R&D Labs and Innovative Productivity

This appendix presents additional results and validation tests for the R&D laboratory analysis.

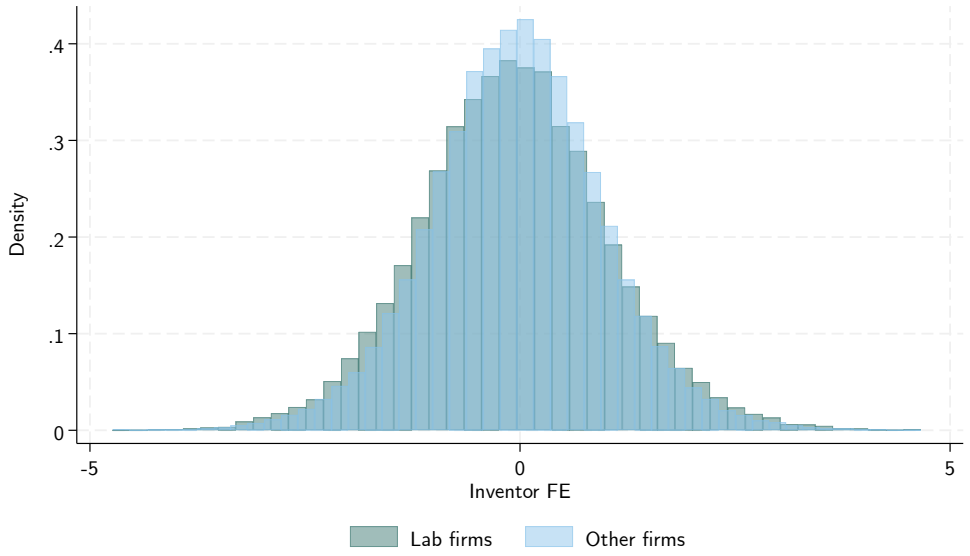
Table E1 reports breakthrough rates by firm assignment status and lab proximity. Table E2 shows AKM variance decomposition results, comparing plug-in and bias-corrected estimates across multiple outcomes. Figures E3, E4, and E5 validate core AKM framework assumptions through symmetry tests, exogenous mobility assessments, and additive separability checks. Section E.1 implements the Bonhomme, Lamadon, and Manresa (2019) framework as an alternative to AKM, using discrete worker and firm types as a check on the main findings.

Table E1: Breakthrough rates by firm assignment and lab proximity

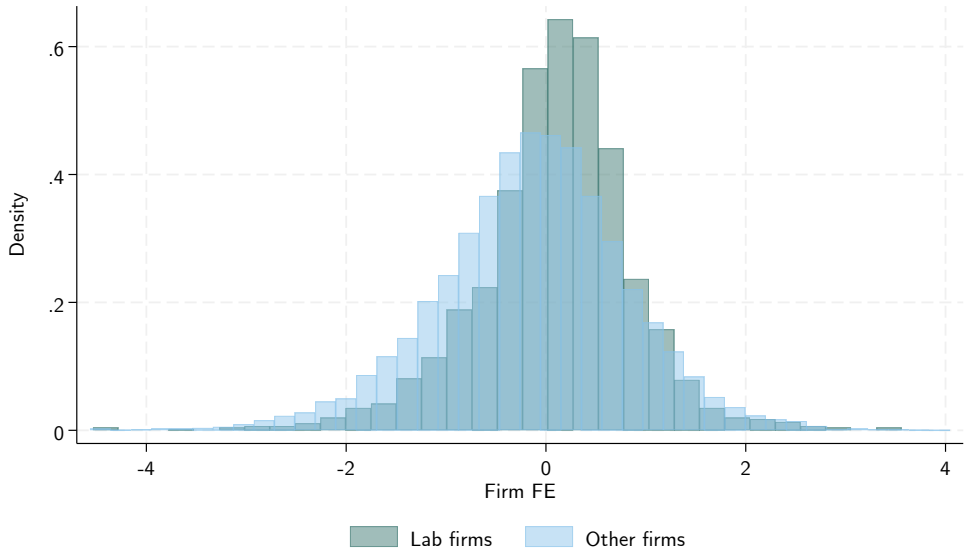
	Patents N (% of total) (1)	Breakthroughs N (% of total) (2)	Breakthrough Rate (%) (3)
Within 50km of a dated & open lab	108,720 (7.2)	22,519 (16.2)	20.7
Within 50km of an undated lab	24,925 (1.7)	5,829 (4.2)	23.4
Outside 50km of any lab (but firm has one)	87,757 (5.8)	14,929 (10.7)	17.0
Firms without labs	357,062 (23.8)	39,278 (28.2)	11.0
Not assigned to firms	922,224 (61.5)	56,699 (40.7)	6.1
Total	1,500,688 (100.0)	139,254 (100.0)	9.3

Note: This table shows breakthrough rates by patent assignment status and proximity to R&D laboratories, 1905-1940. Patents are categorized by inventor location relative to firm labs and assignment type. Breakthrough rates calculated using the Kelly et al. (2021) measure. Results show substantial quality variation across organizational configurations.

Figure E1: Distribution of Firm and Inventor Fixed Effects by Lab Status



(a) Inventor Fixed Effects by Lab Status



(b) Firm Fixed Effects by Lab Status

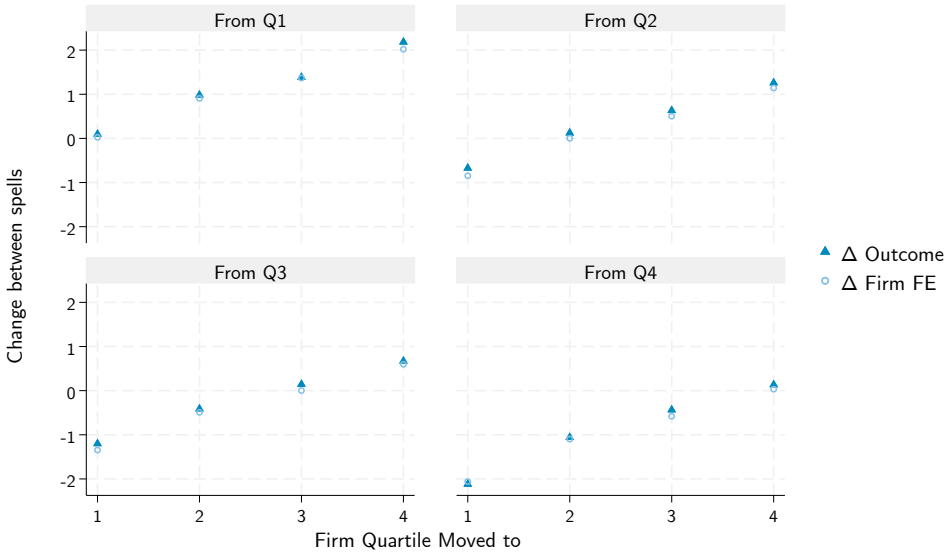
Note: Panel (a) shows the distribution of estimated inventor fixed effects from the AKM model, separated by whether inventors work at firms with or without R&D laboratories. Panel (b) shows the corresponding distribution of firm fixed effects. The virtually identical inventor distributions are evidence against simple inventor sorting, while the shifted firm distributions show higher firm effects among lab-owning firms; because firm effects also absorb other persistent firm characteristics, this comparison is consistent with, but does not on its own establish, a productivity effect of laboratories.

Table E2: AKM Variance decomposition: additional outcomes and naive estimates

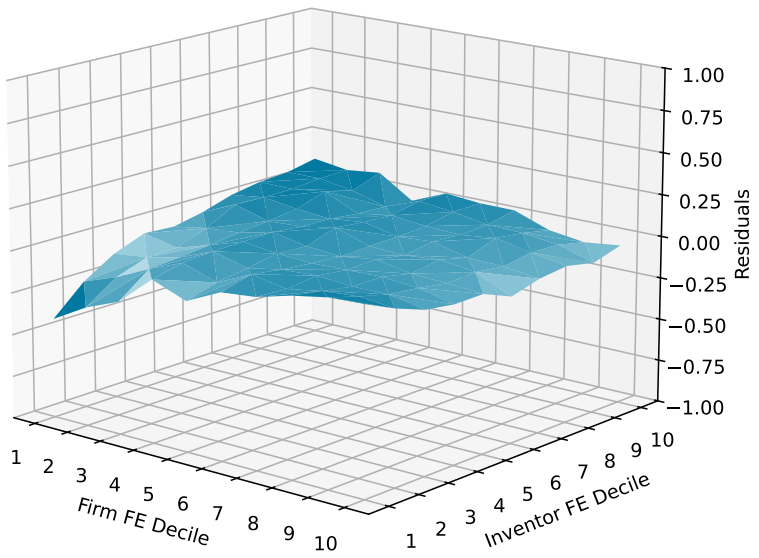
Outcome:	QW Patents	Breakthroughs	UW Patents
Var(y)	1.784	0.197	1.158
Var(ε)	0.664	0.086	0.496
Panel A. Plug-in Variance Decomposition (% of Var($\psi + \alpha$))			
Var(ψ)	30.77	26.31	36.97
Var(α)	90.91	83.64	97.84
$2 \cdot$ Cov(ψ, α)	-21.68	-9.95	-34.80
Panel B. Bias-corrected Variance Decomposition (% of Var($\psi + \alpha$))			
Var(ψ)	32.84	39.09	66.71
Var(α)	75.50	37.38	88.19
$2 \cdot$ Cov(ψ, α)	-8.34	23.54	-54.90
# Spells	227,284	227,284	227,284
# Firms	18,286	18,286	18,286
# Inventors	94,040	94,040	94,040
% Movers	61.38	61.38	61.38

Note: This table reports variance decomposition results from the AKM model (Equation 8) for three outcome measures: quality-weighted (QW) patents, breakthrough patent counts, and unweighted (UW) patent counts. Panel A shows plug-in estimates that do not account for limited mobility bias; Panel B shows bias-corrected estimates using the Kline, Saggio, and Sølvsten (2020) method. The substantial differences between plug-in and bias-corrected estimates highlight the importance of addressing limited mobility bias in AKM estimation.

Figure E2: Validation of AKM Model Assumptions



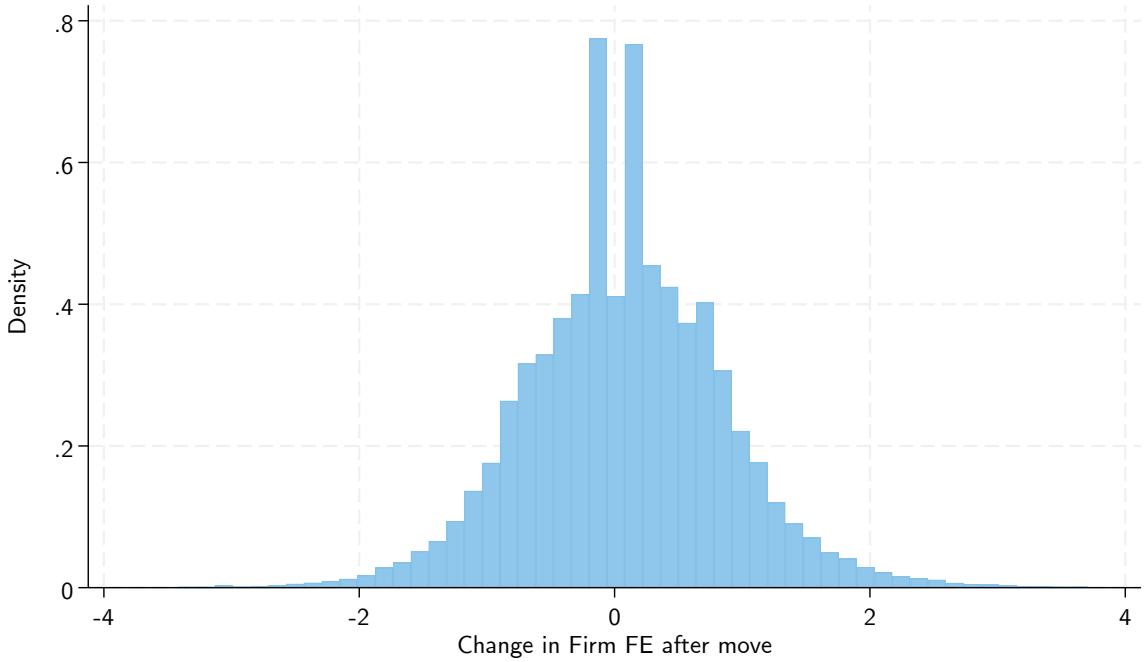
(a) Output changes vs Firm effect changes



(b) Residual by Inventor and Firm Effects

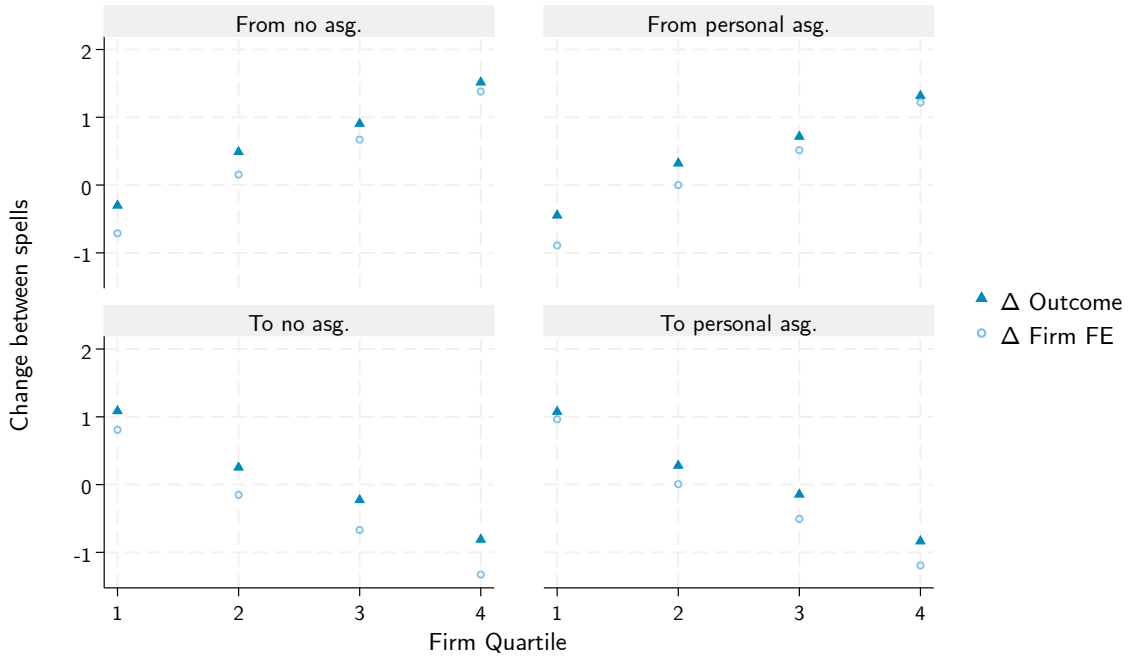
Note: This figure tests the assumptions of the AKM framework using quality-weighted patent output. Panel (a) plots changes in inventor output between spells at different firms (triangles) against changes in firm fixed effects (circles). Moves are organized by quartile of firm effects, e.g., the top left quadrant shows changes when moving from first quartile firms to firms in quartiles one through four. Panel (b) shows average model residuals by deciles of inventor and firm effects.

Figure E3: AKM Validation: Symmetry Test for Inventor Mobility



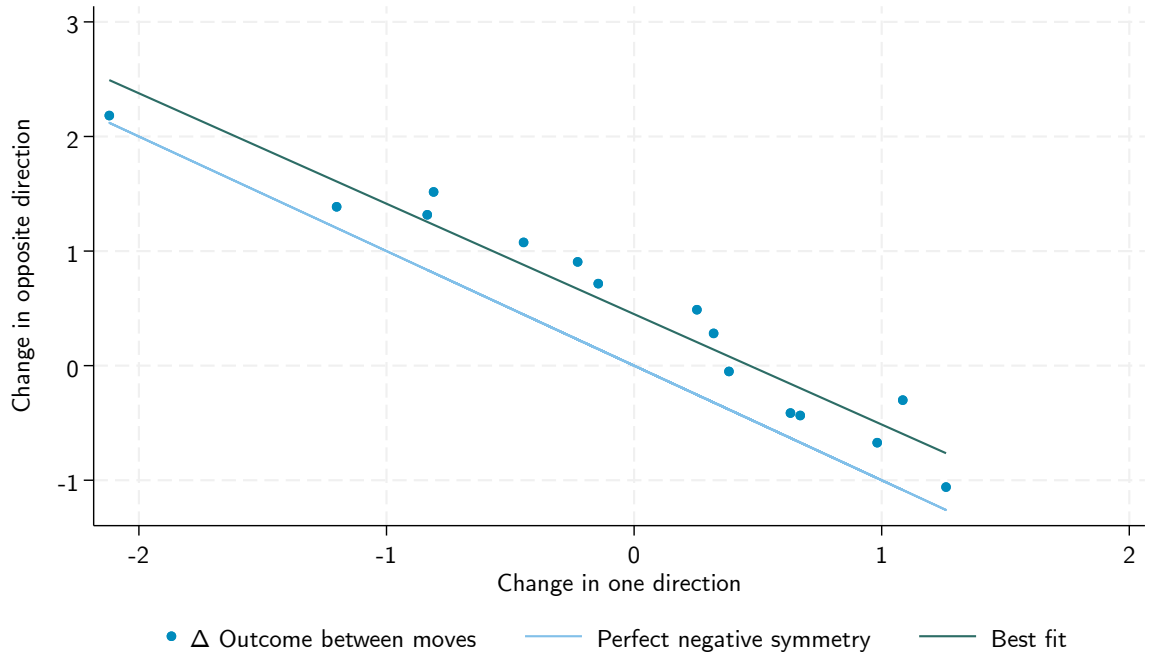
Note: This figure examines the distribution of changes in firm fixed effects experienced by mobile inventors. The histogram shows changes in firm fixed effects when inventors move between firms. Exogenous mobility concerns the relation between moves and transitory productivity shocks; a roughly symmetric distribution of these changes around zero is a reassuring descriptive check, though symmetry is neither necessary nor sufficient for the assumption to hold.

Figure E4: AKM Validation: Output changes vs Firm effect changes Including Pseudo-Firms



Note: This figure extends the main validation test shown in Figure E2 by including transitions to and from pseudo-firms (unassigned patents and personal assignments). The plot shows changes in inventor output (triangles) and corresponding changes in firm fixed effects (circles) for moves involving different firm quartiles and pseudo-firm states.

Figure E5: AKM Validation: Additive Mobility and Perfect Negative Symmetry Test



Note: This figure tests the additive separability assumption of the AKM framework by examining bidirectional moves between firms. The x-axis shows changes in inventor output when moving in one direction (e.g., from firm A to firm B), while the y-axis shows changes when moving in the opposite direction (from firm B to firm A). Under perfect negative symmetry, these changes should be equal in magnitude but opposite in sign, yielding a slope of -1. The observed best-fit line lies close to the perfect-negative-symmetry line, consistent with the additive-separability assumption underlying the AKM model, though the figure does not by itself quantify how close is close enough.

E.1 BLM

As an additional robustness check, I implement the semi-structural approach of Bonhomme, Lamadon, and Manresa (2019) (BLM) as an alternative to the AKM framework. The BLM method addresses several limitations of the standard AKM approach by modeling both worker and firm heterogeneity as discrete types rather than continuous fixed effects. This parametric approach requires weaker identifying assumptions and avoids the limited mobility bias that can affect AKM estimation in samples with sparse firm-to-firm transitions. Additionally, the BLM framework naturally accommodates complementarities between worker and firm types, allowing for richer interaction patterns than the purely additive AKM specification. However, relying on few firm clusters determined from the productivity distribution greatly limits our investigation of the relationship between R&D labs and firm-level productivity.

The BLM estimation procedure uses k-means clustering to classify firms into a small number of discrete types based on their earnings distributions (in this application, inventor productivity distributions), then estimates a finite mixture model with discrete worker types. Because no variance correction is necessary, I apply this method to all the data I have. The specification includes four firm types and three worker types. This dimensional reduction necessarily limits the share of variance that can be attributed to firm heterogeneity, since the model constrains all firm-level variation to operate through just four discrete categories. Figure E6 shows the estimated work-type distribution across firm classes.

Table E3 presents the BLM variance decomposition results alongside the same sample breakdown used in the main AKM analysis. Consistent with the AKM findings, firm effects account for a substantial share of productivity variance, with this share increasing markedly between the early period (1875-1904) and the later period (1905-1950) as industrial R&D became more prevalent. The firm contribution is somewhat lower than in the AKM specification, but this mechanical attenuation reflects the constraint of fitting all firm heterogeneity into just four discrete types.

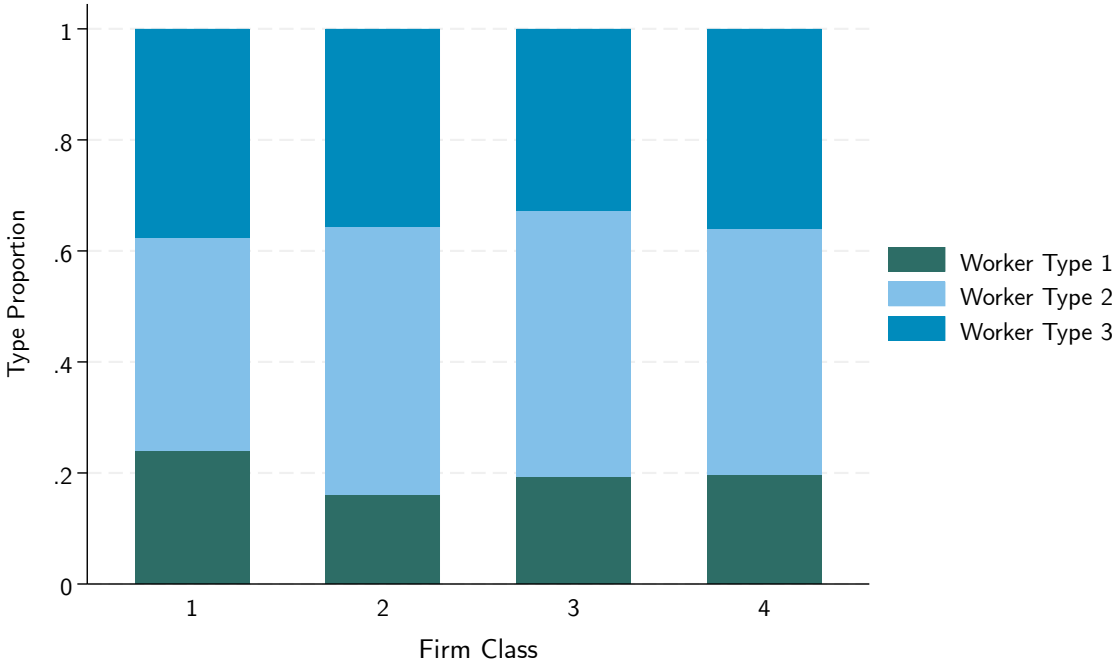
Figure E7 shows average productivity by worker type and firm type, revealing nearly parallel productivity profiles across firm types for each worker category. This parallel pattern indicates negligible complementarities between worker and firm types—highly productive workers gain approximately the same benefit from working at high-productivity firms as do workers of other types. Figure E8 shows that the higher-productivity firm types disproportionately include large firms and lab-owning organizations, consistent with R&D infrastructure being associated with enhanced innovative productivity even within this more structured parametric framework. The share of large firms rises across productivity classes and is highest in the top class, while the share of lab-owning firms is comparably high in the top two classes.

Table E3: BLM Variance decomposition

Sample:	Full sample (1)	1875–1904 (2)	1905–1950 (3)
$Var(y - X'\beta)$	1.701	1.435	1.966
R^2	0.678	0.665	0.667
$Var(\psi)/Var(\psi+\alpha)$	0.174	0.083	0.195
$Var(\alpha)/Var(\psi+\alpha)$	0.816	0.879	0.822
$2\,Cov(\psi,\alpha)/Var(\psi+\alpha)$	0.010	0.038	-0.017
$Corr(\psi,\alpha)$	0.014	0.070	-0.022
Spells	1,164,122	361,017	842,147
Firms	82,007	21,940	66,224
Inventors	944,795	310,204	677,079
Movers (%)	11.25	8.89	12.65

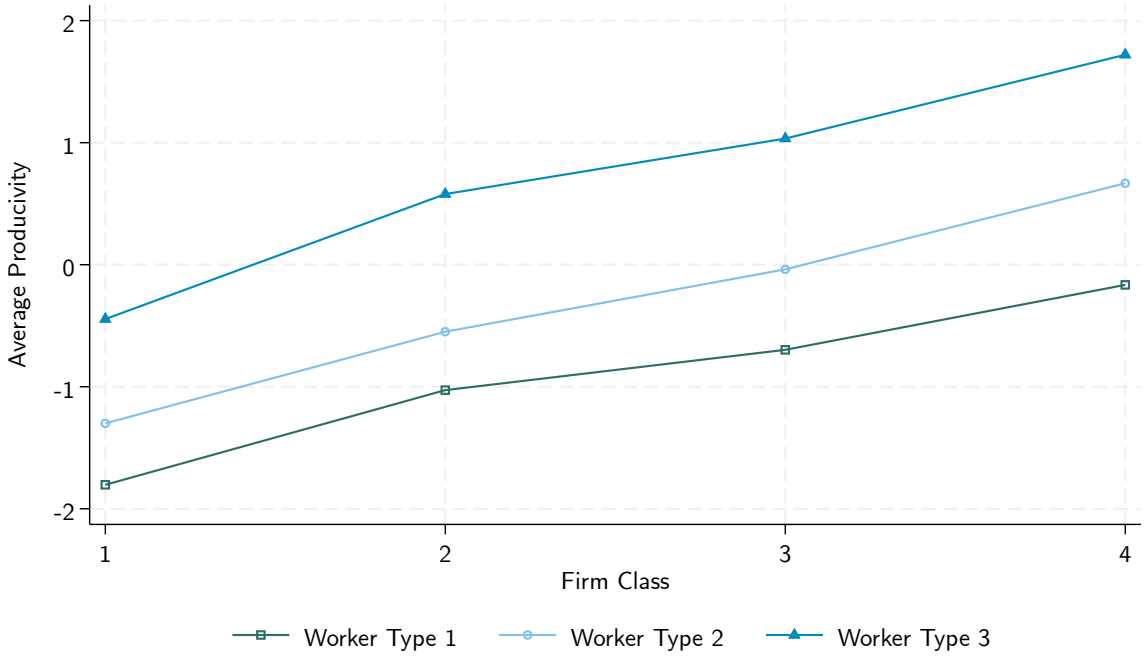
Note: This table reports BLM variance decomposition using discrete worker and firm types (4 firm classes, 3 worker types). Column 1 shows results for the full sample (1875-1950); columns 2-3 show results for early (1875-1904) and later (1905-1950) sub-samples. The sample is larger than in the AKM analysis because BLM does not require the same bias correction restrictions. Results are consistent with the main AKM findings, with firm share increasing from early to later period.

Figure E6: Worker Type Composition Across Firm Classes



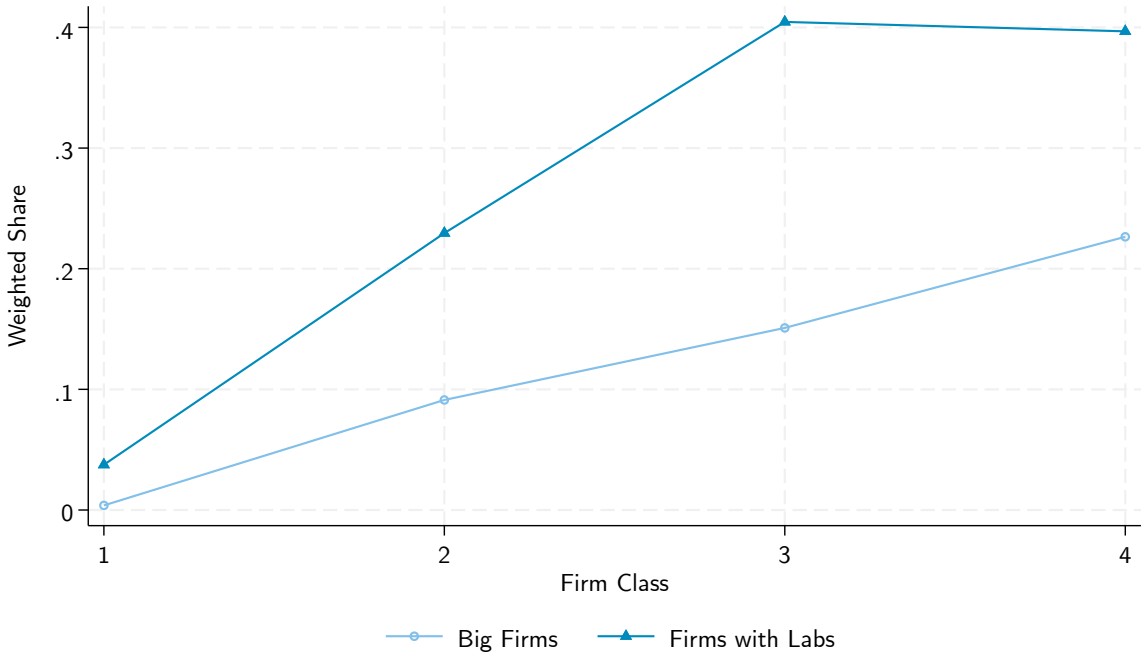
Note: This figure shows the distribution of worker types across firm classes in the BLM model. The stacked bars represent the proportion of each worker type (Types 1, 2, and 3) within each of the four firm classes.

Figure E7: Average Productivity by Worker Type and Firm Class



Note: This figure shows average productivity by worker type across firm classes. Nearly parallel lines indicate negligible complementarities: all worker types benefit similarly from higher-productivity firms, consistent with additive structure.

Figure E8: Firm Characteristics by Productivity Class



Note: This figure shows the spell-weighted share of large firms ("Big Firms") and laboratory-owning firms ("Firms with Labs") within each of the four firm classes estimated by the BLM model. Firm classes are ordered by average productivity, with Class 4 representing the highest-productivity category. The figure shows that the share of large corporations rises across productivity classes, reaching its maximum in the highest class (Class 4), while the share of firms with dedicated R&D laboratories is comparably high in the top two classes, consistent with the association between firm size, R&D infrastructure, and innovative productivity.

F Additional results — Spillover Effects of the Great Merger Wave

This appendix presents additional results and robustness checks for the spillover analysis of Section 6.

Table F1 examines effects on number of active firms, inventors, and firm patent share. Table F2 shows robustness to alternative numbers of size-balanced technology clusters. Table F3 tests sensitivity to vintage fixed effects inclusion. Table F4 explores heterogeneity by GMW exposure intensity using terciles rather than binary treatment. Tables F5 and F6 provide additional emerging technology results, examining exposure intensity effects and robustness to vintage controls. Section F.1 calculates the overall magnitude of GMW impact on American innovation, translating empirical estimates into aggregate breakthrough counts (Table F7).

Table F1: Effect of Consolidation Exposure on Technology-Level Innovation—Additional Outcomes

Outcome:	Active Firms	Active Inventors	Firm Patent Share
Panel A: Overall Effect			
Post × GMW	2.182	3.175	2.602
	(0.460)	(1.571)	(0.728)
	[1.279, 3.084]	[0.093, 6.257]	[1.173, 4.030]
Panel B: Heterogeneity			
Post × GMW × Science-Based	3.662	6.016	3.234
	(1.582)	(4.087)	(1.617)
	[0.558, 6.766]	[-2.003, 14.035]	[0.061, 6.407]
Post × GMW × Other Tech	1.950	2.731	2.503
	(0.468)	(1.680)	(0.806)
	[1.032, 2.868]	[-0.566, 6.028]	[0.920, 4.085]
Year x CPC Section FE	Y	Y	Y
Year x Vintage Tercile FE	Y	Y	Y
Technologies	977	977	977
N	59,597	59,597	59,597

Note: This table reports difference-in-differences estimates of the effect of consolidation exposure on additional technology-level outcomes (Equation 12). The dependent variables are: number of active firms (col. 1), number of active inventors (col. 2), and share of patents assigned to firms versus other assignees (col. 3). Panel A shows overall effects; Panel B shows heterogeneity by science-based (CPC sections C, G, H) versus other technologies. Standard errors clustered at the technology level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Table F2: Effect of Consolidation Exposure on Technology-Level Innovation—By Number of Clusters

	Patents			Breakthroughs		
Target Clusters (k):	750	1000	1500	750	1000	1500
Panel A: Overall Effect						
Post × GMW	2.304	1.139	0.838	0.719	0.374	0.374
	(2.508)	(1.664)	(0.895)	(0.356)	(0.267)	(0.158)
	[-2.621, 7.228]	[-2.126, 4.403]	[-0.917, 2.593]	[0.020, 1.418]	[-0.151, 0.899]	[0.064, 0.683]
Panel B: Heterogeneity						
Post × GMW × Science-Based	8.919	7.388	1.962	4.677	3.472	2.690
	(7.025)	(4.740)	(2.781)	(1.986)	(1.563)	(0.955)
	[-4.873, 22.711]	[-1.914, 16.690]	[-3.493, 7.416]	[0.779, 8.575]	[0.405, 6.539]	[0.817, 4.562]
Post × GMW × Other Tech	1.247	0.162	0.658	0.087	-0.111	0.004
	(2.662)	(1.746)	(0.937)	(0.239)	(0.175)	(0.096)
	[-3.979, 6.472]	[-3.264, 3.587]	[-1.180, 2.496]	[-0.382, 0.556]	[-0.455, 0.234]	[-0.184, 0.192]
Year x CPC Section FE	Y	Y	Y	Y	Y	Y
Year x Vintage Tercile FE	Y	Y	Y	Y	Y	Y
Technologies	742	977	1,435	742	977	1,435
N	45,262	59,597	87,535	45,262	59,597	87,535

Note: This table tests robustness to alternative numbers of target clusters used to construct size-balanced technology domains. The baseline analysis uses k=1000 clusters; this table shows results using k=750 and k=1,500 clusters as well. The dependent variables are total patents (columns 1-3) and breakthrough patents (columns 4-6) at the technology-year level. GMW exposure is defined as having any pre-1895 patents held by firms that subsequently participated in the Great Merger Wave. Standard errors clustered at the technology level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Table F3: Effect of Consolidation Exposure on Technology-Level Innovation—No Vintage FEs

Outcome:	Patents	Breakthroughs	Patents	Breakthroughs
Panel A: Overall Effect				
Post × GMW	2.447	0.510	2.161	0.490
	(1.730)	(0.272)	(1.713)	(0.264)
	[-0.949, 5.842]	[-0.025, 1.045]	[-1.200, 5.522]	[-0.028, 1.008]
Panel B: Heterogeneity				
Post × GMW × Science-Based	11.198	3.907	10.861	3.673
	(4.863)	(1.577)	(4.811)	(1.528)
	[1.655, 20.740]	[0.813, 7.002]	[1.421, 20.302]	[0.674, 6.672]
Post × GMW × Other Tech	1.045	-0.034	0.768	-0.020
	(1.841)	(0.173)	(1.823)	(0.168)
	[-2.568, 4.659]	[-0.373, 0.305]	[-2.809, 4.345]	[-0.350, 0.310]
Excluding GMW firms	N	N	Y	Y
Year x CPC Section FE	Y	Y	Y	Y
Year x Vintage Tercile FE	N	N	N	N
Technologies	977	977	977	977
N	59,597	59,597	59,597	59,597

Note: This table reports difference-in-differences estimates excluding vintage fixed effects from the main specification (Equation 12). The baseline specification includes year fixed effects by three vintage groups based on terciles of technology-level median earliest patent year. This robustness check removes these vintage controls (ν_{tm}) while maintaining year fixed effects by CPC section and technology fixed effects. The dependent variables are total patents and breakthrough patents, with columns 3 and 4 excluding direct contributions from GMW firms. Standard errors clustered at the technology level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Table F4: Effect of Consolidation on Technology-Level Innovation—Exposure Intensity

Outcome:	Patents	Breakthroughs
Panel A: Overall Effect by Intensity		
GMW Tercile 1 × Post	2.517	0.431
	(2.014)	(0.399)
	[-1.435, 6.470]	[-0.352, 1.213]
GMW Tercile 2 × Post	1.852	0.416
	(1.973)	(0.335)
	[-2.020, 5.725]	[-0.241, 1.073]
GMW Tercile 3 × Post	-1.062	0.271
	(1.963)	(0.316)
	[-4.915, 2.791]	[-0.350, 0.892]
Panel B: Heterogeneity by Intensity		
GMW Tercile 1 × Post × Science-Based	14.408	4.181
	(6.730)	(2.495)
	[1.202, 27.614]	[-0.714, 9.076]
GMW Tercile 2 × Post × Science-Based	5.410	2.820
	(5.462)	(2.052)
	[-5.309, 16.128]	[-1.207, 6.848]
GMW Tercile 3 × Post × Science-Based	1.085	3.261
	(5.682)	(1.936)
	[-10.065, 12.236]	[-0.539, 7.061]
GMW Tercile 1 × Post × Other Tech	0.583	-0.172
	(2.032)	(0.211)
	[-3.405, 4.570]	[-0.587, 0.242]
GMW Tercile 2 × Post × Other Tech	1.317	0.040
	(2.080)	(0.222)
	[-2.764, 5.398]	[-0.395, 0.475]
GMW Tercile 3 × Post × Other Tech	-1.440	-0.194
	(2.075)	(0.215)
	[-5.513, 2.632]	[-0.615, 0.228]
Year x CPC Section FE	Y	Y
Year x Vintage Tercile FE	Y	Y
Technologies	977	977
N	59,597	59,597

Note: This table examines heterogeneity by exposure intensity among technologies that had any GMW exposure before 1895. Treated technologies are divided into exposure terciles based on GMW patent share. The specification follows Equation 12 but replaces the binary exposure indicator with tercile indicators. The dependent variables are patents (column 1) and breakthrough patents (column 2). Standard errors clustered at the technology level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Table F5: Effect of Consolidation on Emerging technologies—Intensity of Exposure

Outcome:	First Patent	First Breakthrough
Panel A: Overall Effect		
ln (GMW exposure)	1.003	1.063
	(0.050)	(0.098)
	[0.911, 1.105]	[0.888, 1.272]
Panel B: Heterogeneity		
ln (GMW exposure) × Science-Based	1.098	1.110
	(0.102)	(0.166)
	[0.916, 1.317]	[0.827, 1.489]
ln (GMW exposure) × Other Tech	0.951	1.008
	(0.051)	(0.093)
	[0.856, 1.056]	[0.842, 1.208]
Controls	Y	Y
Groups	1,696	1,696
Subclasses	237	237
N	48,145	65,701

Note: This table reports hazard ratios from the Cox proportional hazard model in Equation 13, but it restricts to exposed technologies and replaces the binary indicator for exposure with an intensive exposure measure: the logarithm of the share of pre-1895 GMW patents in the given CPC subclass. Panel A shows overall effects; Panel B shows heterogeneity by science-based (CPC sections C, G, H) versus other technologies. Hazard ratios above one indicate accelerated emergence; ratios below one indicate delayed emergence. Standard errors clustered at the subclass level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

Table F6: Effect of Consolidation on Emerging Technologies—No Vintage Control

Outcome:	First Patent	First Breakthrough
Panel A: Overall Effect		
GMW	1.086	0.888
	(0.073)	(0.104)
	[0.952, 1.240]	[0.706, 1.116]
Panel B: Heterogeneity		
GMW × Science-Based	1.117	0.982
	(0.124)	(0.142)
	[0.898, 1.390]	[0.739, 1.303]
GMW × Other Tech	1.057	0.700
	(0.079)	(0.093)
	[0.914, 1.223]	[0.539, 0.909]
CPC Class FE	Y	Y
Groups	2,898	2,898
Subclasses	474	474
Vintage	N	N
N	84,109	109,995

Note: This table reports hazard ratios from Cox proportional hazards models excluding vintage controls from the baseline emerging technology specification (Equation 13). The baseline specification includes controls for the earliest year any pre-1895 patent appeared in each group's subclass (EarliestYear_g). This robustness check removes this vintage control while maintaining CPC class fixed effects and section-specific baseline hazards. The sample includes CPC groups with no patents before 1895. Panel A shows overall effects; Panel B shows heterogeneity by science-based versus other technologies. Hazard ratios above one indicate accelerated emergence; ratios below one indicate delayed emergence. Standard errors clustered at the subclass level are shown in parentheses. 95 percent confidence intervals are reported in square brackets.

F.1 Implied causal effect of the GMW

To assess the overall magnitude of the Great Merger Wave's impact on American technological development, I perform back-of-the-envelope calculations to estimate the total effect on breakthrough innovations attributable to consolidation, across both established and emerging technology domains.

The approach differs between established and emerging technologies due to data structure constraints. For established technologies, I directly apply the difference-in-differences coefficients from Table 7 to the number of technology-year observations exposed to consolidation during 1905-1940. Specifically, I multiply the estimated treatment effects by the corresponding number of exposed observations in each category. This yields the total number of breakthrough innovations added or lost due to GMW exposure over the sample period.

For emerging technologies, the survival analysis framework precludes direct extrapolation of breakthrough counts from hazard ratios. The Cox proportional hazards model estimates the relative likelihood of first breakthrough occurrence but does not readily translate into absolute numbers of missed or accelerated innovations. To address this limitation, I estimate a complementary regression specification that mirrors the identification strategy of the survival analysis while providing interpretable coefficients in levels. The specification includes CPC class fixed effects, vintage-by-year fixed effects, and CPC section-by-year fixed effects, capturing the same sources of variation as the Cox model. Formally, for CPC group g in calendar year t, I estimate:

$$B_{gksmt} = \beta_S \, \text{SExp}_g + \beta_O \, \text{OExp}_g + \gamma_k + \nu_{tm} + \sigma_{ts} + \varepsilon_{gksmt}, \tag{F1}$$

where B_{gksmt} is the number of breakthrough patents in group g (in class k in section s) and year t, γ_k are CPC class fixed effects, ν_{tm} are vintage by year fixed effects, and σ_{ts} are CPC section by year fixed effects. SExp_g and OExp_g are exposure indicators equal to one if group g belongs to a subclass exposed to consolidation and falls, respectively, in science-based sections (C, G, H) or in other sections (A, B, D, E, F, Y). Standard errors are clustered at the subclass level. The sample is restricted to $t \in [1905, 1940]$ and to CPC classes with within-class variation in exposure (as in the Cox model). This regression essentially captures the average difference in the number of breakthroughs between exposed and unexposed groups, conditional on the controls outlined before.

The aggregate results, presented in Table F7, reveal that consolidation exposure accounts for approximately 13,047 breakthrough innovations during 1905-1940, representing a 13.2 percent increase relative to the counterfactual without the Great Merger Wave. This aggregate effect masks substantial heterogeneity across technological domains. Among science-based technologies, GMW exposure generated 16,080 additional breakthroughs, constituting a 30.3 percent increase above counterfactual levels. In contrast, non-science-based domains experienced a net reduction of 3,033 breakthroughs, representing a 6.7 percent decline.

Table F7: Net Breakthrough Innovations Attributable to the Great Merger Wave

	Science-Based (1)	Other (2)	Total (3)
Established	10,874	-2,321	8,553
	(29.4%)	(-5.4%)	(10.7%)
Emerging	5,206	-713	4,493
	(32.3%)	(-28.5%)	(24.1%)
Total	16,080	-3,033	13,047
	(30.3%)	(-6.7%)	(13.2%)

Note: This table reports the estimated net number of breakthrough innovations attributable to Great Merger Wave exposure during 1905-1940. For established technologies, effects are calculated by multiplying difference-in-differences coefficients from Table 7 by the number of exposed technology-year observations. For emerging technologies, effects are estimated using a complementary linear regression specification that mirrors the identification strategy of the survival analysis in Table 8. Science-based technologies encompass CPC sections C, G, and H (chemistry, metallurgy; scientific instruments, computing; electronics, telecommunications). Numbers in parentheses show the percentage change relative to the counterfactual scenario without GMW exposure, calculated as the net effect divided by the implied baseline (observed total minus net effect).