Servitization and Firm Scope*

Salomé Baslandze
Federal Reserve Bank of Atlanta

Elvira Sojli
UNSW Sydney

Wing Wah Tham
UNSW Sydney

August 4, 2026

Preliminary and incomplete. Please do not circulate.

Abstract

The shift toward services is usually understood as a reallocation of economic activity from goods-producing sectors toward service-producing sectors. We show that servitization is also a change in firm scope: goods producers increasingly add service offerings around the products they sell. Standard data cannot observe this margin: industry codes assign firms a primary activity, segment reporting is managerial, and services are rarely patented. We measure it using the goods and services statements filed with U.S. trademark applications. Parsing 10.8 million marks and 51.7 million dated items, we classify individual offerings as goods or services and link firm portfolios to Compustat and the National Establishment Time Series. We document that the service share of offerings rose sharply after 1990, including within firms and sectors classified as goods producing. Entry is highly structured: machinery firms add installation and repair, pharmaceutical firms add medical and monitoring services, and electronics firms add software and technical support. Service entry is followed by substantial firm growth. We explain these facts with a model in which services combine diagnosis with execution. Advances in IT, software, sensors, and intangible capital turn products in use into sources of data and remote coordination, lowering the cost of related services and making them more tradable and scalable. Goods become inputs into service production, giving producers pair-specific advantages, shifting activity from independent providers, raising the role of installed-base concentration, and limiting Baumol's cost disease to the execution component of services.

*Acknowledgements and disclosures: This research did not receive external financial support. The authors have no financial relationships or other potential conflicts of interest to disclose. The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any errors are our own.

1 Introduction

The shift from goods toward services is one of the defining structural transformations of advanced economies. In the United States, services accounted for roughly half of employment in 1950 and more than four-fifths in recent years, with similar changes across high-income countries (Buera and Kaboski, 2012; Herrendorf, Rogerson, and Valentinyi, 2014). The standard account describes this transformation as a reallocation of economic activity from goods-producing sectors toward service-producing sectors. It also carries a pessimistic implication: because many services resist mechanization, a rising service share may slow aggregate productivity growth (Baumol, 1967; Nordhaus, 2008). Yet services are also expanding within sectors conventionally classified as goods producing. Caterpillar's VisionLink platform uses data generated by connected machinery to monitor equipment, diagnose faults, and coordinate maintenance. Tesla uses vehicle data to identify component failures, schedule repairs, deliver software updates, and price insurance. Both remain classified as goods producers even as they supply activities traditionally associated with specialized service firms.

We study this process, which we call servitization, and ask two questions. First, how much of the shift toward services appears in the goods and services firms bring to market, rather than in the industry classifications on which conventional measures of structural transformation rely? Second, why are goods producers increasingly entering service markets, and how does their entry change firm scope and the organization of service provision? Our argument is that the modern service shift is not only a reallocation of activity across sectors. It also occurs within goods-producing sectors, as firms add services to their product portfolios and their products become inputs into service production.

Neither question can be answered with standard data, because the goods-service boundary is not observed within the firm. Industry classification assigns each firm a primary activity, so a manufacturer of equipment and a manufacturer of equipment sold with monitoring, maintenance, and analytics receive the same code. Financial statements report revenue by operating segment rather than by the nature of the product offering, and segment definitions are set by management and revised at its discretion. Text-based measures constructed from product descriptions in 10-K filings recover the breadth of a firm's markets (Hoberg and Phillips, 2025), but they do not separate goods from services, and the underlying language is drafted by each firm's counsel without a common vocabulary. Patents record invention in codified, patentable technologies (Kogan et al., 2017), and services are rarely patentable. The result is that firms are observed to move toward services in the aggregate while the underlying decision, which service a firm adds and when, is unobserved.

We measure that decision directly, from the goods and services statements filed with U.S.

trademark applications. A trademark protects a brand used in commerce, and registration requires the applicant to state what is sold under it. The statements are itemized: each lists discrete offerings, and every item falls in one of forty-five Nice classes, thirty-four for goods and eleven for services. An examining attorney reviews each item against a published controlled vocabulary and requires amendment before registration, so a given offering is described in common terms across filings. Registration requires a specimen showing the mark in actual use in commerce. Thereafter the registrant must declare under oath, in the sixth year and at each renewal, that every item remains in use and must delete those that do not, and the USPTO audits these declarations and strikes items whose use cannot be proven. What we observe is therefore an examiner-validated list of offerings, item by item, and pays to keep. We use the universe of USPTO trademark filings to construct a granular record of commercial offerings using 10.8 million trademarks including 51.7 million distinct goods and service items. We classify each mark as pure goods, pure services, or a bundle listing both under a single brand, date entry by first use in commerce, and observe exit when items lapse. We link the resulting portfolios to two firm panels: CRSP-Compustat, which supplies financial outcomes and independent measures of service activity against which the trademark measures can be validated, and the National Establishment Time Series, which extends the sample from listed firms to employers of every size. Together they yield what is, to our knowledge, the first record of firm scope that observes goods and services symmetrically, at the level of the individual offering and across the full size distribution of U.S. firms, and therefore the first in which the relation between a good and the services that accompany it can be measured.

We validate the measure in four ways. Firms that add service items subsequently reclassify from goods-producing to service industry codes, and the trademark portfolio registers the shift three to five years before the classification does. Among firms reporting business segments, the service indicator agrees with the presence of service-oriented segment revenue in 90 percent of firm-years. Service items appear before service language enters the business description of the annual report, and 97 percent of the firms identified as service providers by their portfolios are confirmed in that text. Firm histories accord with the record: the measure tracks IBM's transition into services, Caterpillar's into equipment monitoring, and Amazon's from retail services into proprietary goods, and it declines for General Electric over its divestiture of service businesses, while consumer goods manufacturers exhibit no movement in either direction. The portfolio therefore tracks the composition of what a firm sells, and not the accumulation of filings.

We document four facts. First, the shift toward services appears in the goods and services firms bring to market. The service share of offerings rose from 21% in 1990 to 46% in 2020,

with the bulk of the move between 1995 and 2000 and a plateau thereafter. Second, the shift extends well beyond firms that industry classifications identify as service producers. Among goods-producing firms with trademark activity, the share holding a service-related mark rose from 22% in 1995 to 45% in 2020. The increase holds within cohorts defined by a firm's first filing year: goods firms first filing in the 1940s entered with essentially no service marks and reach 46% by 2019, level with firms first filing in the 2000s and above those first filing in the 2010s. The rise is therefore not the replacement of goods producers by differently constituted entrants, and it occurs without any accompanying change in industry classification. Third, the services these firms add are far from a random draw. Relative to a benchmark in which goods firms add services in proportion to the frequency with which each service is added economy-wide, entry is concentrated xxx times above chance in pairs such as machinery with installation and repair services, pharmaceuticals with medical and monitoring services, and electronics with technical and software services. Fourth, service entry is followed by expansion. Employment rises 23% over five years relative to matched non-servitizing firms in the same industry, and 5.6% relative to goods firms expanding their trademark portfolios in the same year, a comparison with flat pre-trends that isolates the component specific to service entry.

Diversification toward a growing sector accounts for neither the sorting of entry across good-service pairs nor its timing, and it leaves the fourth fact unexplained: firms are expanding by moving into the sector where productivity growth is supposed to stall. Reconciling this requires distinguishing the components of a service that can be mechanized from those that cannot, and asking what the good contributes to the first. Doing so also raises a question the facts alone cannot settle. If the information a service requires is generated by the product, the producer holds an input its rivals cannot purchase, and service provision migrates toward firms with the largest installed bases. Whether that is a source of efficiency or of market power depends on the production structure, which is what we model next.

We develop a model of good and service supply. Each service line can be served by an independent provider or by the producer of an associated good, and the provider with the lower cost supplies it. Production requires *diagnosis* and *execution*, with weights that differ across lines. Execution is labor. Diagnosis is labor for the independent provider, but a goods producer can build a stock of intangible capital, in sensors, software, and the systems that read them, and use it to obtain diagnosis from products already in use. How much it obtains depends on whether the good generates information relevant to that service, on the size of the installed base, and on the productivity of the available data technology. The stock is costly and convex and each additional line carries a setup cost, so a producer chooses its intangible capital and its set of service lines jointly, spreading the stock across the lines it

enters. Entry requires both a high diagnostic weight and a good capable of supplying the diagnosis, which is why the third fact appears as a matrix of pairs rather than as a ranking of services. Improvements in data technology raise diagnostic quality for the pairs that have the link and lower the cost of adding a line, moving pairs across the entry threshold and delivering the timing in the first fact.

Two implications follow that the facts do not contain. First, the service the goods producer adds was already supplied. The producer's advantage is the independent provider's disadvantage, since what falls is the cost of the diagnostic labor the product replaces, so part of the expansion recorded in the fourth fact is a transfer, and the increase in service provision at the entering firm overstates the increase in the economy. Second, the producer's gain is a lower marginal cost and not a higher markup. In the model, prices fall, quantities rise, and share concentrates on the producer with the largest installed base, so concentration in these service lines is a consequence of cost advantage rather than of pricing power. We test the first implication directly, asking whether employment at independent providers of a service line declines where goods producers in the associated product category have entered. (tbd) Execution, the component of a service to which Baumol's argument applies, is unaffected throughout: what the intangible layer displaces is diagnosis, and the technician still performs the repair.

The mechanism has three consequences beyond the firms that adopt it. First, official statistics miss part of the shift to services. When a machinery producer begins selling monitoring contracts, the revenue is recorded in manufacturing, because the firm is still classified as a manufacturer. The service is produced and sold, but the accounts assign it to goods, so the sectoral series understate the reallocation by whatever share of it occurs inside goods-producing firms. Second, Baumol's argument binds on some components of a service and not others. The cost of diagnosis has fallen; the cost of an hour of a technician's time has not. A service is exposed to the cost disease in proportion to its execution share, which differs across services and can be measured, so the pace at which a service sector's relative price rises is a function of its composition rather than of its being a service. Third, these lines are increasingly supplied by the firms that sold the underlying good, and the advantage is a lower marginal cost rather than a higher markup. Concentration arising from a cost advantage has different welfare consequences from concentration arising from pricing power, and the two are separated by where the advantage originates rather than by the level of concentration itself.

The paper makes three contributions. The first is measurement. A trademark application states, item by item, the goods and services sold under the mark, and each item is assigned to a class that separates the two. Classifying items therefore yields a firm's offerings, dated

and decomposed into goods and services, for every U.S. employer that files. Prior trademark research counts marks and relates holding to firm size and patenting (Graham et al., 2013; Dinlersoz et al., 2018); work on servitization measures revenue earned from services, in surveys and in French administrative data (Crozet and Milet, 2017). Neither identifies which services a firm supplies, and that identification is what the correspondence between goods and services requires.

The second contribution is that correspondence. Services attach to goods systematically, and we estimate the resulting matrix: the frequency with which producers in each goods category enter each service category, relative to the frequency at which that service is entered economy-wide. The matrix is a primitive that has not been measured, it separates servitization from diversification, and any theory of firm scope that spans goods and services must reproduce it.

The third contribution is the mechanism that generates it. Firms enter the services their goods can inform, because the good supplies the diagnostic input those services require and the independent provider must purchase it with labor. The result situates the growth of services inside goods producers within work on data and intangible capital (Jones and Tonetti, 2020; Farboodi and Veldkamp, 2023), supplies structural transformation with a margin interior to the firm (Ngai and Pissarides, 2007; Herrendorf, Rogerson, and Valentinyi, 2014), and extends the analysis of firm scope from goods to the boundary between goods and services (Bernard, Redding, and Schott, 2010; Mayer, Melitz, and Ottaviano, 2014).

2 Data and Measurement

This section describes our approach to measuring firm-level servitization using trademark filings. We first explain why trademarks are well suited for tracking goods and services produced by firms, then describe the structure of trademark data and our key innovation of item-level measurement, followed by our matching procedure to link trademarks to firm-level data, construction of time-varying portfolios, variable definitions, and validation exercises.

The central measurement challenge is that the paper's question is about products, not sectors. We need to observe the goods and services available in the economy and, crucially, whether they are introduced by goods producers, service producers, or firms that integrate the two. Standard datasets are poorly suited to this task. Scanner data cover only selected consumer goods; trade data observe internationally traded products; industry codes assign firms or establishments to a primary activity; financial statements report broad segments; and patents capture patentable invention rather than commercial offerings. Trademarks provide a different object: a legally reviewed record of branded goods and services offered in

commerce. They therefore let us observe product-scope choices directly, including services that would be largely invisible in patent data (Graham et al., 2013; Dinlersoz et al., 2018).

The scale of the data is important for the paper. The full USPTO file contains 10.8 million trademarks and 51.7 million goods and service items over 1925–2020. In linked firm data, trademarking is selected but economically broad: a small share of firms file trademarks, yet trademarking firms account for a much larger share of employment; patenting firms frequently also file trademarks, while most trademarking firms do not patent; and in our public-firm panel, live trademark portfolios are observed for over 80% of Compustat firm-years. These coverage facts are not peripheral. They are what make trademarks useful for measuring the modern service margin: the data are broad enough to capture commercialized products, rich enough to distinguish goods from services, and timed to track firm scope changes.

2.1 Why Trademarks?

We obtain trademark data from the USPTO Trademark Case Files Dataset, which contains the complete history of trademark applications filed with the USPTO starting from 1870 (Graham et al., 2013). A trademark is a word, name, phrase, symbol, logo, design, sound—or a combination of these—that identifies the source of particular goods or services and distinguishes them from competitors. A trademark is a legal protection for a brand used in commerce. To obtain and maintain that protection, the applicant must specify the goods and services sold under the mark, and the application is reviewed by the USPTO. This makes trademarks a unique large-scale source of data on products (goods and services) available in the economy.

Trademarks are uniquely suited to study service provision and goods-service bundling for several reasons. First, the Nice International Classification system provides an internationally standardized taxonomy that explicitly distinguishes goods from services. Goods are covered by Classes 1–34, while services are covered by Classes 35–45. The Nice Classification came into effect at the USPTO on September 1, 1973. Trademark applications filed before this date were categorized under a U.S. national classification system comprising 52 goods classes, Classes 1–52, and 8 service classes, Classes 100–107, plus classes for collective membership marks and certification marks. The USPTO has harmonized the U.S. classes to Nice classes for the full dataset. Tables A1 and A2 provide further details on the Nice and prior U.S. classifications.

Second, trademarks are registered at the level of specific goods and services. Unlike financial statements that aggregate across business segments, or patents that cover tech-

nologies rather than commercial products, trademarks directly reflect what firms offer in the marketplace. Crucially, trademark rights require continuous use in commerce, which means firms must demonstrate active use of their marks in connection with registered goods or services, or risk cancellation (see discussions in Graham et al., 2013; Leo Liu and Tham, 2026, for additional details). This legal requirement ensures that trademarks capture what firms are selling, rather than what they plan to sell, used to sell, or could potentially commercialize. While patents can be held without commercialization and financial disclosures may include discontinued segments or forward-looking projections, our trademark-based measure provides a real-time, verified record of firms' active market presence.

Third, trademarks have precise timing. Each trademark includes exact filing dates, first-use-in-commerce dates, registration dates, and cancellation or expiration dates. This temporal precision enables us to construct time-varying portfolios that capture firms' entry into and exit from specific product markets, which is essential for studying the dynamics of servitization and bundling.

Fourth, trademarks provide broad coverage of commercial activity. Unlike patents, which cover only patentable innovations, trademarks span the full universe of commercial products and services offered by a firm. This scope is critical for studying services: while goods-producing firms frequently patent their technologies, service activities are rarely patentable but routinely trademarked. Our data thus capture service provision that would be invisible in patent-based or R&D-based measures of firm activity.

Overall, filing a trademark requires payment of government fees, precise legal description of offerings, and subsequent demonstration of use in commerce. Unlike survey responses or voluntary disclosures, trademarks are legally binding commitments backed by enforcement mechanisms. This combination of cost, specificity, and legal consequence makes trademarks a useful indicator of firms' product-market strategies.

Coverage is also central. Figure 16 shows that trademark-based indices contain information distinct from patent-based measures of innovation. Figures 17 and 18 show that trademark use is widespread over time and rises sharply with firm size. This selection is useful for our question: the firms most likely to hold trademarks are the firms for which product-scope adjustment is economically consequential. At the same time, the trademark/patent overlap facts documented in Dinlersoz et al. (2018) show why trademarks expand the innovation measurement frontier rather than merely duplicating patent data. Many patenting firms trademark, but most trademarking firms do not patent, which is especially important for services and other non-patentable commercial offerings.

2.2 Structure of Trademark Filings

Each trademark filing contains several components critical to our analysis. A trademark is assigned a unique serial number that enables tracking over time. The owner information field records the legal name and address of the trademark holder, which we use to match trademarks to firms. The owner field is updated when trademarks are transferred through mergers, acquisitions, or sales. Each trademark records the filing date, registration date, status, and end date, where the end date captures whether a trademark has been canceled, expired, abandoned, or remains active. Trademarks must be renewed every 10 years and require periodic filings demonstrating continued use in commerce. These dates enable us to construct time-varying trademark portfolios.

Each trademark is registered in one or more of 45 Nice classes, and each mark can be identified as a goods mark, a service mark, or both. The classification comprises:

- Goods classes, Classes 1–34, covering tangible products, from chemicals in Class 1 to pharmaceuticals in Class 5, electronics in Class 9, and clothing in Class 25.
- Service classes, Classes 35–45, covering intangible services, from business services in Class 35 to financial services in Class 36, telecommunications in Class 38, and technology services in Class 42.

2.2.1 Goods and Services Descriptions: Item-Level Measurement

For our purpose, the most granular and useful component of each trademark filing is the goods and services description. For each Nice class in which a trademark is registered, the applicant must provide a detailed statement specifying the particular goods or services covered. Crucially, each statement contains one or more distinct items, usually separated by semicolons. An item is a single, discrete good or service description extracted from the semicolon-delimited goods/services statement within each Nice class. This grammatical structure reflects USPTO requirements: each item must describe goods or services that would naturally be understood as a cohesive unit.[1] Consider Apple Inc.'s trademark for "iBeacon" (Serial No. 86111001), filed in November 2013:

> **Serial No. 86111001 – IBEACON**
> **Filing Date:** November 5, 2013
> **Class 35 (Advertising and Business Services):** "Business management; data processing services; advertising agency services; advertising, marketing, and

[1] See TMEP §1402.01(a). Examiners require applicants to separate items that improperly combine unrelated offerings.

promotion services; advertising and marketing consultation; sales promotion services; promoting the goods and services of others; conducting market research; analysis of advertising response and market research; design, creation, preparation, production, and dissemination of advertisements and advertising material for others; media planning services; administration of consumer loyalty programs; arranging and conducting incentive rewards programs to promote the sale of goods and services; providing business, consumer, and commercial information over computer networks and global communication networks; business services, namely, providing computer databases regarding the purchase and sale of a wide variety of products and services of others; retail store services; online retail store services provided via the Internet and other electronic and communications networks; arranging and conducting of commercial, trade, and business conferences, shows, and exhibitions; information, advisory and consultancy services relating to all the aforesaid."

This single trademark contains 19 distinct service items, all separated by semicolons:

1. Business management
2. Data processing services
3. Advertising agency services
4. Advertising, marketing, and promotion services
5. Advertising and marketing consultation
6. Sales promotion services
7. Promoting the goods and services of others
8. Conducting market research
9. Analysis of advertising response and market research
10. Design, creation, preparation, production, and dissemination of advertisements and advertising material for others
11. Media planning services
12. Administration of consumer loyalty programs
13. Arranging and conducting incentive rewards programs to promote the sale of goods and services
14. Providing business, consumer, and commercial information over computer networks
15. Business services, namely, providing computer databases regarding the purchase and sale of products
16. Retail store services
17. Online retail store services

18. Arranging and conducting of commercial, trade, and business conferences, shows, and exhibitions
19. Information, advisory and consultancy services relating to all the aforesaid

This example illustrates several key features of our measurement approach. First, despite iBeacon being widely known as a Bluetooth Low Energy hardware technology, Apple's trademark registration classifies it as a service offering focused on location-based marketing and retail services. This demonstrates how trademark data capture firms' commercial positioning of their offerings, not just the underlying technology.

Second, the granularity of item-level data reveals the scope of Apple's intended service bundle. Rather than simply registering "iBeacon services," Apple specified 19 distinct service activities, from advertising consultation to loyalty program administration to online retail services. This specificity is legally required: the use-in-commerce requirement applies at the item level, meaning Apple must demonstrate actual commercial use of each listed service or risk cancellation.

A critical institutional feature of the trademark system is that the goods and services descriptions we analyze are not self-reported text that passes through the USPTO unchecked. Each trademark application is assigned to an examining attorney who reviews the goods and services statement line by line against the USPTO Acceptable Identification of Goods and Services Manual, the "ID Manual," a controlled vocabulary of pre-approved descriptions maintained by the Office.[2] If the applicant's language is vague, overbroad, or inconsistent with ID Manual conventions, the examining attorney issues an office action requiring amendment. The application cannot proceed to registration until every item in the goods and services statement satisfies the examiner's requirements. This adversarial, iterative review process imposes a standardization discipline on the underlying data that is absent from self-reported sources such as 10-K business descriptions, survey responses, or financial segment disclosures.

The iBeacon application illustrates this process concretely. Apple's original filing described 19 service items in Class 35. The examining attorney required amendments to several of these items before allowing registration. Table 1 displays representative corrections. Each correction moves the language in the same direction: from vague to precise, from potentially overbroad to bounded, and from self-referential to standardized. The examining attorney required Apple to specify what kind of service is being offered, for whom it is being performed, and what subject matter it covers. The correction adding the "of others" qualifier

[2]The ID Manual is publicly available at `https://idm-tmng.uspto.gov/id-master-list-public.html`. It contains pre-approved entries organized by Nice class. Examining attorneys use it as the primary reference for evaluating the acceptability of goods and services descriptions; see TMEP §1402.01.

is particularly informative for economic measurement. Class 35 of the Nice Classification covers services performed on behalf of third parties. Without the qualifier, the description could be read as the applicant administering loyalty programs for its own benefit, which is not a registrable service under the classification.[3] The examiner enforces this distinction, ensuring that the final description accurately reflects the commercial relationship the applicant maintains with its customers.

Table 1
Examiner Amendments to iBeacon Goods and Services Description

The table presents selected amendments required by the USPTO examining attorney for Apple Inc.'s iBeacon trademark application (Serial No. 86111001, Class 35, filed November 5, 2013). *Original filing language* is the language of the initial application; *examiner-corrected language* is the final registered description. Corrections are identified from the Tracked Text Description field in the USPTO Trademark Case Files Dataset, which records amendments between the original and final goods and services statements.

Original filing language	Examiner-corrected language	Nature of correction
Media planning services	Planning advertising media for others	Specifies the type of media and identifies the beneficiary of the service
Retail store services	Retail store services featuring a wide variety of consumer goods	Identifies the subject matter of the retail activity, eliminating an overbroad entry
Administration of consumer loyalty programs	Administration of consumer loyalty programs to promote the sale of goods and services of others	Establishes that the service is performed on behalf of third parties, a legal requirement for Class 35 registration
Arranging and conducting incentive rewards programs to promote the sale of goods and services	Arranging and conducting incentive rewards programs to promote the sale of goods and services of others	Adds the "of others" qualifier required under Class 35
Providing business, consumer, and commercial information over computer networks and global communication networks	Providing business, consumer product, and commercial information over computer networks and global communication networks	Replaces the ambiguous "consumer" with "consumer product," specifying the type of information provided

[3]The distinction between performing services for oneself and for others is a fundamental requirement of trademark registration in service classes. An applicant that provides services only to itself is not engaged in "commerce" within the meaning of the Trademark Act; see TMEP §1401.03(b).

This institutional mechanism has three consequences for the quality of our measurement. First, it ensures cross-firm comparability. Because examining attorneys enforce the same ID Manual standards on every applicant, the item "retail store services featuring a wide variety of consumer goods" carries an identical meaning whether it appears in an Apple filing, a Walmart filing, or the filing of a small regional retailer. This standardization is what makes item-level aggregation across firms economically meaningful. By contrast, 10-K business descriptions are drafted by each firm's legal counsel with no external vocabulary constraint: two firms performing identical activities may describe them in entirely different language, making systematic comparison infeasible at scale.

Second, the institutional setting provides economic precision at the item level. The examiner corrections systematically eliminate the ambiguity that would otherwise compromise our ability to classify items as goods versus services and to identify bundling. When the final description states "planning advertising media for others," we can classify this as a service with high confidence. When goods and service items survive jointly under the same trademark after examiner review, we can interpret the bundle as a deliberate commercial commitment rather than a drafting artifact, because the examining attorney has verified that each item conforms to its respective class and that the applicant has demonstrated, or intends to demonstrate, use in commerce for each listed item.

Third, the process limits strategic scope inflation. Trademark applicants have an incentive to draft broad descriptions to maximize the scope of legal protection. The examining attorney's role is precisely to resist this tendency and to narrow descriptions to recognized, bounded entries that correspond to identifiable commercial activities. This adversarial dynamic means that the items surviving into the final registration are more likely to reflect genuine offerings than aspirational claims. The requirement of demonstrated use in commerce, verified through specimens submitted to the examining attorney, reinforces this constraint: items for which the applicant cannot show actual commercial activity are subject to cancellation.[4]

The examining attorney process means that our item-level data inherit a quality-control mechanism built into the legal infrastructure of the trademark system. Each item in our dataset has been reviewed against a standardized taxonomy by an independent examiner, amended where necessary to conform to established conventions, and verified through evidence of commercial use. This institutional foundation distinguishes trademark-based mea-

[4]Under 15 U.S.C. §1058, registrants must file a declaration of continued use between the fifth and sixth years after registration and at each subsequent ten-year renewal. Failure to demonstrate use for specific goods or services results in cancellation of those items from the registration. This ongoing verification requirement means that our stock measure of active trademarks reflects current commercial offerings rather than historical filings.

surement from approaches that rely on unaudited self-reporting and provides the precision that enables the granular scope classification central to our analysis.

A potential concern is that large firms like Apple file multiple trademark applications for the same mark covering different goods and services. Does this artificially inflate our item counts or distort our bundling measures? This is not a measurement problem, but part of the variation we exploit. Under U.S. trademark law, firms cannot expand the goods or services covered by an existing registration. When a firm enters a new product market, it must file a new trademark application, even for an existing brand name. Each registration thus represents a distinct commercial commitment: the first-use-in-commerce date records when the firm entered that specific market, and the legal requirement to demonstrate continued use means each registration is independently verified.

Consider Apple Inc., which holds numerous registrations for the mark APPLE. Serial No. 73120444 covers APPLE for computers and computer software in Class 9, with first use in April 1976. Serial No. 75857151 covers APPLE STORE for online retail store services provided via a global computer network featuring computers, computer hardware, computer systems, computer software, computer peripherals, and accessories in Class 35, with first use on November 1, 1997. Serial No. 75826806 covers IMUSIC for online chat rooms, bulletin boards, and broadcasting of audio and audio-visual material in Class 38, with first use in August 1995 and cancellation in August 2008. These registrations reflect economically distinct events: Apple's original personal computer business, its entry into digital entertainment services, and its expansion into online retail services. Aggregating these to a single brand-level measure would erase the dynamics we seek to study.

We verify this interpretation through a manual audit of a random sample of firms with multiple registrations for the same mark and confirm that the large majority represent genuine market expansions, such as new product categories, geographic markets, or business models, rather than administrative duplicates. Furthermore, in portfolio measures we use unique items in the firm portfolio, not the number of times an item is introduced.

From the full USPTO database, we extract item-level observations from unique trademarks. Figure 15 shows the breadth of service items appearing in the data, and Table 2 summarizes the full USPTO hierarchy and the CRSP-Compustat matched sample. This item-level granularity is the key innovation in our measurement approach. Prior studies using trademark data typically count the number of Nice classes in which a firm files. Our item-level approach offers three advantages. First, items capture product-specific information lost when aggregating to the 45 Nice classes. A firm filing in Class 9 might be offering database software, smartphones, or scientific instruments, distinctions that matter for understanding product strategy but are invisible at the class level. Second, each item is classified

into a specific Nice class, allowing us to precisely identify whether each offering represents a good or a service. Third, items bundled within the same trademark reveal information about product relationships. When a firm bundles software products with technical support services under one trademark, it signals that these offerings are marketed as an integrated solution rather than separate products.

2.2.2 Trademark Types: Pure Goods, Pure Services, and Bundles

Based on item-level classification, we categorize each trademark and the items within a trademark into one of three types:

- **Pure goods trademark:** contains only goods items, all in Classes 1–34. Example: athletic footwear; athletic apparel; sports bags.
- **Pure service trademark:** contains only service items, all in Classes 35–45. Example: retail store services featuring athletic apparel; online retail ordering services.
- **Bundle trademark:** contains both goods and service items. Example: computer hardware; computer software; technical support services; software installation and maintenance services.

Bundle trademarks are particularly informative about servitization strategy. When a firm files a bundle trademark, it makes a legal commitment to offer goods and services together under a unified brand, a stronger form of integration than filing separate trademarks for goods and services. Table 3 gives additional examples of how item-level parsing distinguishes pure goods, pure services, and bundles. This measure connects the paper to work on goods-service complementarity and product-service integration (Wise and Baumgartner, 1999; Cusumano, Kahl, and Suarez, 2015; Crozet and Milet, 2017).

2.3 Matching Trademarks to Firms

Linking trademark owners to firms requires reconciling inconsistent name conventions across databases. For example, Compustat may record a trademark owner as "International Business Machines Corp," while the USPTO lists "IBM Corporation" or "International Business Machines." We develop a multi-stage procedure that combines exact matching, fuzzy matching, and subsidiary linkages.

The procedure operates in several stages. First, we standardize company names in both databases. We then carry out exact matching in two tiers: on standardized names with suffixes retained, and on suffix-stripped names. Second, unmatched firms enter a fuzzy matching stage that combines Levenshtein distance, token-sort ratio, and partial string scores. To limit

false positives, we block on the first four characters of the company name, impose a maximum length ratio between candidate names, and require a minimum token overlap. Only matches above a high composite similarity threshold are retained. Lower-confidence fuzzy matches are excluded from the analysis.

Third, because publicly traded firms often hold trademarks through subsidiaries, we supplement direct name matching with subsidiary information from Exhibit 21 disclosures in firms' 10-K filings and M&A data. When a trademark owner matches a disclosed subsidiary or acquired target, the trademark is assigned to the parent company in the relevant years. Fourth, we account for mergers and acquisitions. When a target firm is acquired, its trademarks are reassigned to the acquiring entity from the completion date onward. This step ensures that the trademark portfolio of the surviving firm reflects its full post-merger scope, and that trademarks are not lost from the panel when target firms exit.

2.4 Linked Firm Panels: Compustat and NETS

In the empirical analysis, the trademark data are linked to two firm panels. The CRSP-Compustat match provides public-firm accounting outcomes, segment reporting, and validation exercises over 1980–2019, covering roughly 700,000 matched trademarks, 2 million items, and 19,000 firms. Table 2 reports the corresponding matched trademark and item counts in the current paper sample. We also construct a NETS firm panel that links trademark owners to firm-family-year employment histories. The baseline NETS sample keeps firms that ever reach at least three employees, giving us a broader employment panel for studying firm entry, growth, and scope changes outside the public-firm universe.

2.5 Constructing Time-Varying Trademark Portfolios

Our empirical strategy requires measuring each firm's product portfolio at specific points in time. We construct time-varying portfolios using the temporal information in trademark filings.

2.5.1 Active Trademark Definition

For each trademark, we extract the year of item introduction using the earliest available date among the first-use-in-commerce date, the filing date, and the registration date. The first-use-in-commerce date is the date on which the firm introduced a particular item to the market, verified by an examining attorney through a specimen. Because this date is sometimes missing, we use the trademark filing year when necessary. We refer to this variable as the filing year.

For each firm i and year t, we define the set of active trademarks as those that have been filed but not yet canceled, abandoned, or expired:

$$\mathcal{M}_{i,t} = \{m : \text{filing year}_m \leq t \text{ and } (\text{end year}_m > t \text{ or end year}_m = \varnothing)\}. \tag{1}$$

This definition captures the stock of trademark protection at each point in time, accounting for both inflows, new trademark filings that enter the portfolio, and outflows, cancellations, expirations, and abandonments that exit the portfolio.

2.5.2 Portfolio Aggregation

For each firm-year, we aggregate across all active trademarks to calculate:

$$\begin{aligned}
\text{Goods Items}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \#\{\text{goods items in trademark } m\}, \\
\text{Service Items}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \#\{\text{service items in trademark } m\}, \\
\text{Total Items}_{i,t} &= \text{Goods Items}_{i,t} + \text{Service Items}_{i,t}, \\
\text{Pure Goods Items}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \#\{\text{goods items in pure-goods trademark } m\}, \\
\text{Pure Service Items}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \#\{\text{service items in pure-service trademark } m\}, \\
\text{Bundle Items}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \#\{\text{items in bundle trademark } m\}.
\end{aligned}$$

Similarly, we count trademarks by type:

$$\begin{aligned}
\text{Pure Goods TMs}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \mathbf{1}\{\text{trademark } m \text{ is pure goods}\}, \\
\text{Pure Service TMs}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \mathbf{1}\{\text{trademark } m \text{ is pure service}\}, \\
\text{Bundle TMs}_{i,t} &= \sum_{m \in \mathcal{M}_{i,t}} \mathbf{1}\{\text{trademark } m \text{ is a bundle}\}.
\end{aligned}$$

We construct servitization measures using both stock, or cumulative portfolio, variables and flow variables based on new filings. Our first measure of servitization is the service share of a firm's active trademark portfolio:

$$\text{Service Share (Stock)}_{i,t} = \frac{\text{Service Items}_{i,t}}{\text{Total Items}_{i,t}}.$$

This measure captures the fraction of a firm's product-market presence, as revealed by trademarks, devoted to services. We also construct extensive-margin indicators:

$$\text{Has Services}_{i,t} = \mathbf{1}\{\text{Service Items}_{i,t} > 0\},$$
$$\text{Has Bundle}_{i,t} = \mathbf{1}\{\text{Bundle Items}_{i,t} > 0\}.$$

For analyses examining changes in strategy, we construct flow measures based on firm i's new trademarks in year t. Pure Goods Share (Flow) is the share of pure-goods trademarks in total new trademarks, Pure Service Share (Flow) is the share of pure-service trademarks in total new trademarks, and Bundle Share (Flow) is the share of bundle trademarks in total new trademarks. These three shares sum to one and capture how firms allocate new trademark filings across goods, services, and bundled offerings. Flow measures are defined only for firm-years with at least one trademark filing. Table 4 reports summary statistics for the core servitization variables and the Compustat controls used in the validation analysis.

The portfolio measures also vary sharply across industries, which is useful for interpreting both the level of servitization and the scope for goods-producing firms to add services. Table 7 reports service and bundle shares by one-digit SIC industry over the full sample and at the 1980 and 2019 endpoints. The table shows that service trademarks are common outside narrowly defined service industries and that manufacturing industries move toward service and bundle activity over time.

2.6 Validation

Our analysis relies on trademark filings to measure firms' service offerings. A natural concern is whether trademark activity reflects actual service provision or merely legal precaution. We present several pieces of evidence that trademark filings capture real economic activity.

We validate our measures first by focusing on notable case studies of firms that move from pure services or pure goods to providing both. We then show that service trademark filings predict changes in firms' reported business activities through industry reclassification, correlate with service revenue in segment data, and precede textual mentions of services in SEC filings. Together, these tests suggest that our trademark-based measure captures meaningful variation in firms' service strategies.

2.6.1 Case Study Evidence

As a qualitative validation, we examine whether our trademark-based service share tracks known strategic changes and developments in prominent firms spanning technology, indus-

trial, and consumer goods sectors: IBM, Microsoft, Apple, Amazon, Google, Salesforce, GE, Caterpillar, P&G, J&J, and Walmart.

The technology firms IBM, Microsoft, and Apple exemplify classic servitization, the strategic shift from goods to services. Figure 19 shows before/after changes around major strategic events, Figure 20 shows the corresponding composition of trademark filings, and Table 5 reports the decade-level service and bundle shares together with key firm events. IBM's trajectory corresponds to its transformation from a hardware manufacturer to a services company, including the expansion of IBM Global Services and the acquisition of PwC Consulting (Gerstner, 2002). Microsoft's trajectory captures the transition from packaged software products, such as Windows and Office, to cloud services, including Azure, Office 365, and Xbox Live. Apple's shift reflects the expansion from Macintosh hardware to iTunes, the App Store, iCloud, Apple Music, and Apple Pay.

Amazon, Google, and Salesforce present reverse cases that further validate our measure. Amazon began as a service-native online retailer, then moved toward a mix that included proprietary goods and acquired physical retail assets. Google expanded from search and advertising into mobile phones, tablets, and computers. Salesforce, the pioneer of software-as-a-service, remains strongly service-oriented but also changes portfolio composition through acquisitions and software products. These declining service shares show that our measure captures strategic portfolio rebalancing in both directions.

General Electric and Caterpillar demonstrate servitization in traditional heavy manufacturing. GE's service share rose alongside the expansion of GE Capital, aviation services, engine maintenance contracts, and healthcare services, then declined with restructuring and divestiture. Caterpillar's increase corresponds to its explicit services strategy, including equipment financing, telematics, remote monitoring, aftermarket parts, and maintenance services. Procter & Gamble and Johnson & Johnson serve as control cases in consumer goods industries. Both maintain low service shares throughout the sample period, consistent with their core business of manufacturing consumer products.

The case evidence validates our trademark-based measure in three ways. First, the trajectories match known strategic transformations. Second, the measure captures portfolio shifts in both directions. Third, control firms maintain expected low service shares. The case-study summary in Table 5 summarizes the patterns. These patterns provide confidence that our measure reflects genuine shifts in firms' product-market strategies.

2.6.2 Trademark Filings and Industry Reclassification

If service trademark filings reflect genuine business transformation, they should predict changes in how firms describe their activities to regulators. We test whether service trade-

mark filings predict transitions from goods-producing to service-oriented industry codes and vice versa.

We define an industry switch as a firm changing its Compustat historical industry code from a goods-producing sector to a service sector in a given year, excluding temporary reclassifications that reverse within two years. Following the structure of the first digit of SIC codes, we classify industries 6, 7, and 8, Finance, Business Services, and Professional Services, as the service sector in our baseline definition, and exclude switches originating from Trade, SIC1 = 5, to avoid ambiguous cases where retail and wholesale activity overlaps with service provision.

Firms that subsequently reclassify into service industries exhibit systematically different trademark portfolios in the years before the switch. The service share of trademark filings rises before formal industry reclassification and continues rising afterward, consistent with a gradual shift in business orientation rather than an abrupt administrative change. Stock measures, which accumulate past filings, are higher throughout and trend upward around the switch. Figure 21 shows that service orientation embedded in trademark portfolios anticipates and accompanies formal industry reclassification.

To test this relationship more formally, we estimate whether current trademark activity predicts future industry switching. Specifically, we regress an indicator for switching to a service industry in year $t+k$ on an indicator for filing at least one service trademark in year $t-1$, controlling for firm size, employment, and profitability, with year and industry fixed effects:

$$\mathbf{1}\{\text{Goods} \to \text{Service SIC}\}_{j,t+k} = \alpha + \beta \cdot \text{ServiceTM}_{j,t-1} + X'_{j,t}\gamma + \delta_i + \delta_t + \varepsilon_{j,t}. \qquad (2)$$

The dependent variable equals one if firm j transitions from a goods SIC to a service SIC between years t and $t+k$.

The results, summarized in Table 6, show that filing a service trademark predicts future industry switching. Given a baseline switching rate well below one percent per firm-year, the estimated effects are economically meaningful. Firms with lower profitability are also more likely to switch, consistent with the interpretation that some reclassifications reflect firms in declining goods sectors reorienting toward service activities. Importantly, goods trademark filings do not predict industry transitions in the same way, suggesting that the result is specific to service activity.

Together, the event-study evidence and forward-looking regressions support the view that trademark portfolios contain information about firms' evolving industry orientation that precedes and predicts formal reclassification.

2.6.3 Trademark Filings, Compustat Segments, and 10-K Text

Compustat Segments data provide another window into firms' business activities. For a subset of firms, Compustat Segment data report revenue by business segment, which we classify as goods- or service-related based on segment descriptions. We examine whether service trademark filings correlate with the appearance of service-oriented business segments. We classify segments as service-oriented based on segment SIC codes and text descriptions, including keywords such as "service," "consulting," "support," "maintenance," and "subscription."

An alternative source of information is the business description in firms' 10-K filings. We analyze 10-K text for its correlation with increased discussion of services in annual reports. We count service-related keywords in the Business Description section. The positive correlations across these independent datasets indicate substantial convergent validity between our trademark-based measures and alternative service measures.

3 Empirical Facts

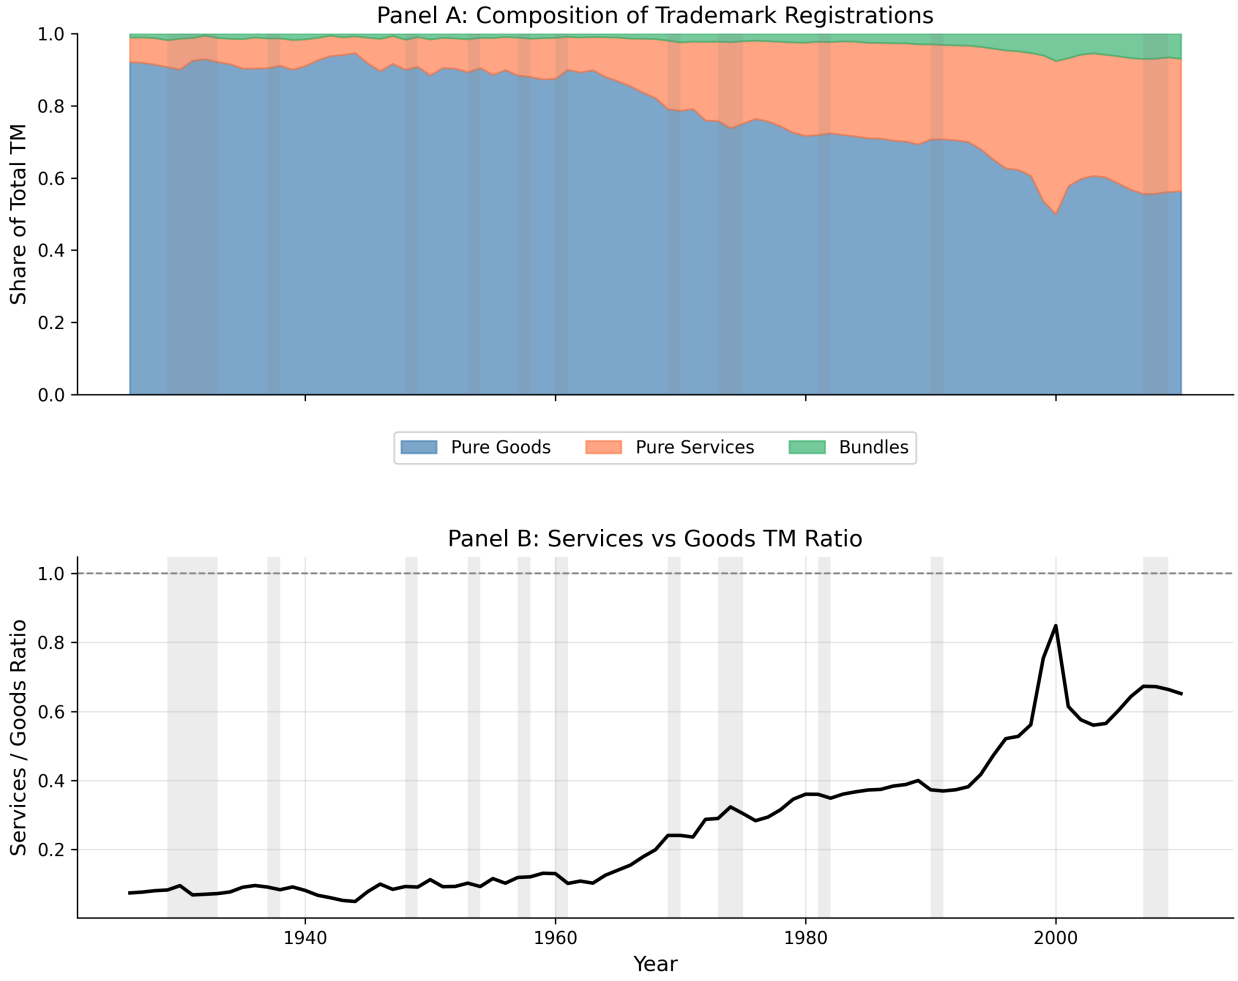
This section organizes the main descriptive and reduced-form evidence around four empirical facts. The first two establish the object to be explained: the shift toward services appears in product space, and it is not simply the expansion of firms already classified as service-sector firms. The third fact shows that the services added by goods firms are often related to the goods they already produce. The fourth fact links service entry to subsequent firm growth.

3.1 Fact 1: Structural Transformation in Product Space

The standard evidence on structural transformation comes from employment, value added, or industry classifications. Our data allow us to ask a different question: whether the product space itself is shifting toward services. Figure 1 shows the composition of trademark filings over time. Trademark filings increasingly contain service items and goods-service bundles, and the service-to-goods ratio rises sharply after the 1990s. This is a product-level analogue of structural transformation: firms increasingly bring service offerings, and service-complemented offerings, to market. This fact is useful because it moves the paper away from a sector-relabeling story. Industry codes assign each firm to a primary activity, but trademarks record the goods and services the firm offers. The rise of service-related trademarks therefore captures a change in the composition of commercial offerings, not just a reallocation of firms or workers across industries. This motivates the central measurement idea of the paper: servitization is a change in firm scope, observed at the product level.

Figure 1
Trademark Composition over Time

The figure presents the composition of trademark filings over time. Trademarks are categorized in three mutually exclusive groups: *Pure goods* includes all trademarks where all items are goods, *Pure services* includes all trademarks where all items are services, and *Bundle* includes all trademarks where there are both goods and services items in the trademark. All ratios are with respect to the total number of trademarks filed in year t. Shaded areas are NBER recession periods.



3.2 Fact 2: Goods Producers Move to Services

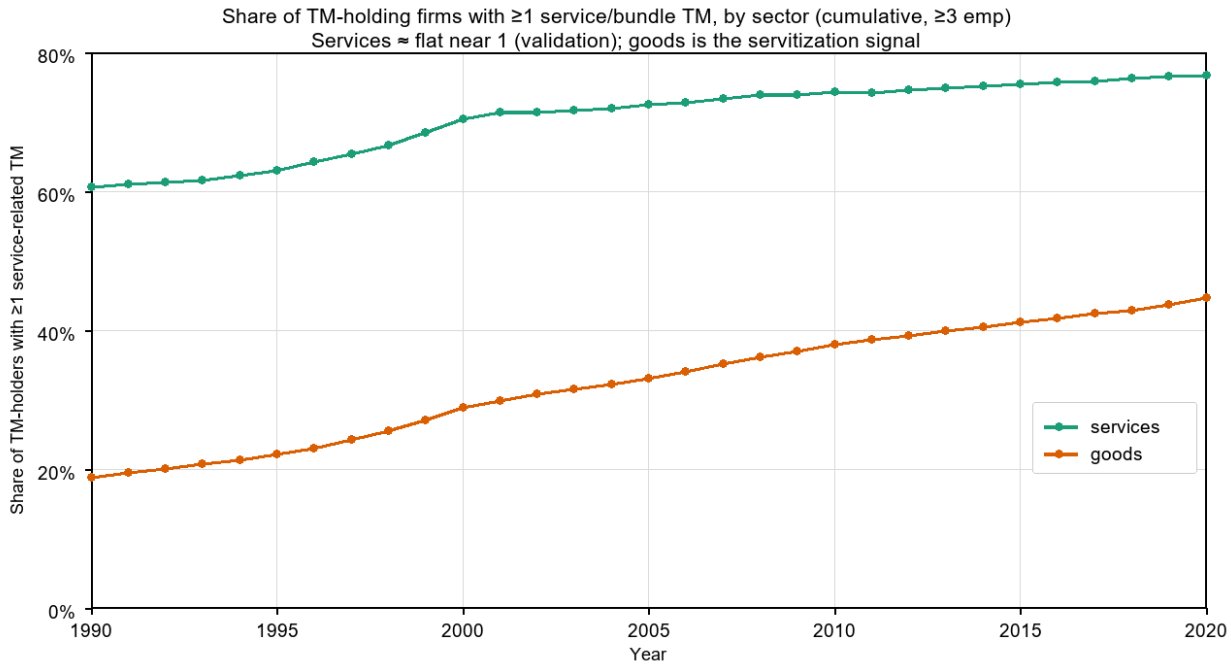
The second fact asks whether the product-space shift is driven only by service-sector firms, or whether goods-producing firms themselves add services. Figure 2 splits firms into goods, distribution, and service groups. Service firms are service-heavy throughout, as expected. The striking movement is the goods-firm line: among goods firms with trademark activity, the share with service-related trademarks rises from roughly 19 percent in 1990 to 44 percent in 2019. Distribution firms also become more service-oriented, while service firms remain high throughout the period.

The long-panel cohort evidence in Figure 3 provides a complementary view. Each cohort fixes a set of goods firms by first filing period and tracks the service-related share of its trademark portfolio over time. The service share rises within old and new goods-firm cohorts alike: for example, the 1940 cohort rises from essentially zero service-related trademark activity to about one-half by 2019, while the 1980 cohort rises from about 17 percent to 39 percent. The pooled line therefore combines two forces: within-firm scope change and cohort composition. Both point in the same direction.

This fact is the key empirical object for the theory. The paper is not only about a growing service sector. It is about goods producers entering service lines, adding service products to goods portfolios, and altering the boundary between goods and services inside firms.

Figure 2
Service Trademark Activity by Firm Sector

The figure presents the share of trademark-holding firms whose portfolio contains at least one service-related trademark, by year and firm sector. Firms are split into goods, distribution, and services groups on the basis of the NAICS sector of the firm's headquarters establishment in NETS. A firm counts from the first year its portfolio contains a service or bundle trademark and remains counted thereafter. The sample period is 1990 to 2020.

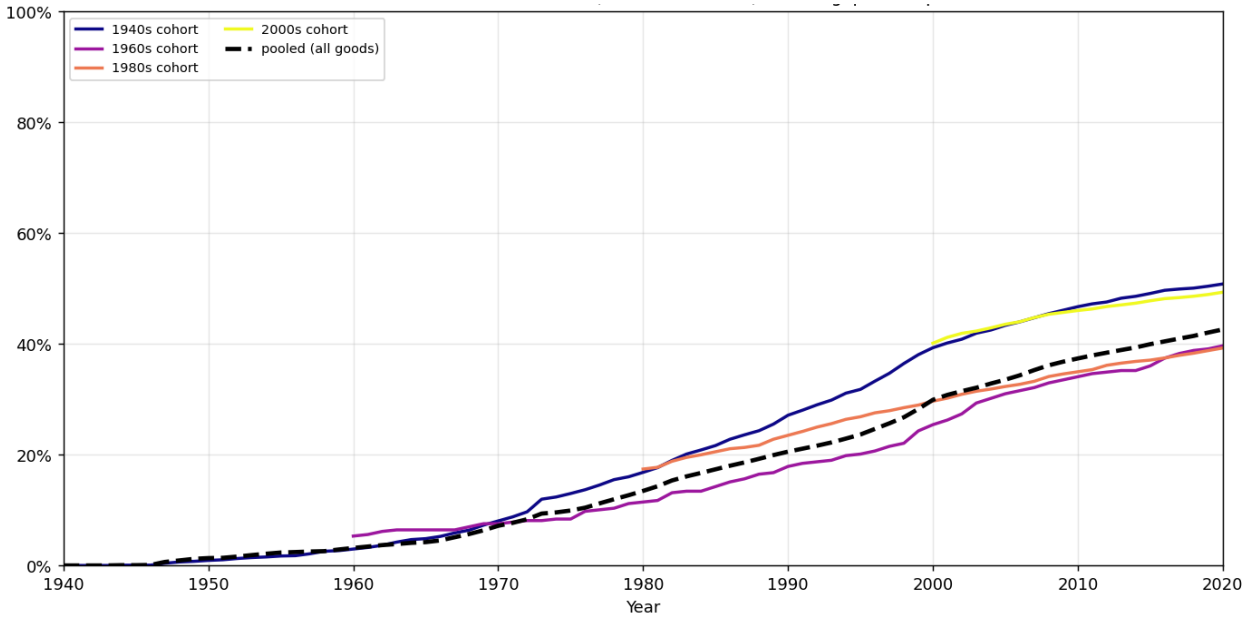


3.3 Fact 3: Added Services Are Related to Firms' Goods

The third fact asks what kinds of services goods firms add. Figure 4 shows that goods firms increasingly add service classes over time, with advertising and business services, Class 35,

Figure 3
Goods-Firm Service Shares by Filing Cohort

The figure presents the service-related share of the trademark portfolio for fixed cohorts of goods firms, together with the pooled line. A cohort is the set of goods firms whose first trademark filing falls in a given period, and each cohort is tracked in calendar time thereafter. The sample is restricted to firms that appear in NETS from 1990 onwards with at least three employees; the trademark spine is floored at 1940. Industry is the firm's NAICS sector in 1990 carried back.



and software and IT services, Class 42, accounting for a large share of the volume. Class 42 is especially prominent around the dot-com period, while Class 35 grows steadily and becomes the largest service category by the end of the sample. In the underlying service-class series, Class 35 rises from 325 trademark-class additions in 1990 to 4,209 in 2020, and its share of service additions rises from 11.9 to 22.1%. Class 42 remains large in levels, rising from 1,079 to 3,914 additions, but its share falls from 39.6 to 20.6% as service additions broaden.

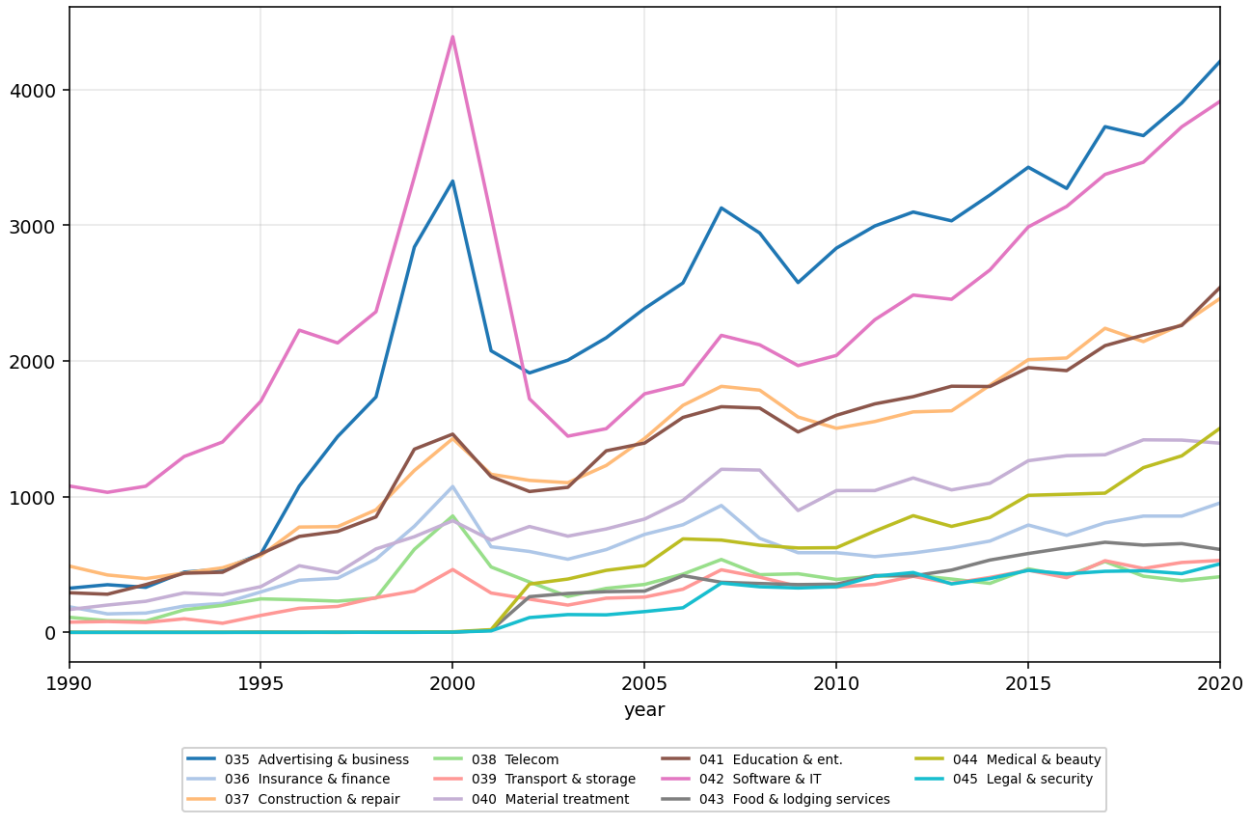
Figure 5 then links service additions to goods categories. The heatmap assigns each goods firm to a primary goods class at inception and asks which service classes those firms add. Three service columns dominate across many goods rows: advertising and business services, education and entertainment services, and software and IT services. At the same time, goods-specific pairings appear in economically sensible places. Pharmaceutical goods are paired with medical and beauty services; toys and sporting goods are paired with education and entertainment services; and many equipment and electronics categories are paired with software, technical support, and business services.

This pattern matters for interpretation. If goods firms were simply diversifying randomly into unrelated services, the rise of services would be hard to distinguish from generic scope

expansion. Instead, the services they add look related to the information, distribution, maintenance, training, software, or customer-support margins around their goods. This is the empirical bridge to a model in which goods create data, diagnostics, installed-base relationships, or customer access that make some services cheaper or more valuable when supplied by the goods producer.

Figure 4
Service Classes Added by Goods Firms

The figure presents service-class additions by goods firms over time, one line per NICE service class 35 to 45. A goods firm is a firm whose NETS headquarters sector is goods-producing. Each trademark contributes once to each distinct NICE class it covers. The sample period is 1990 to 2020.

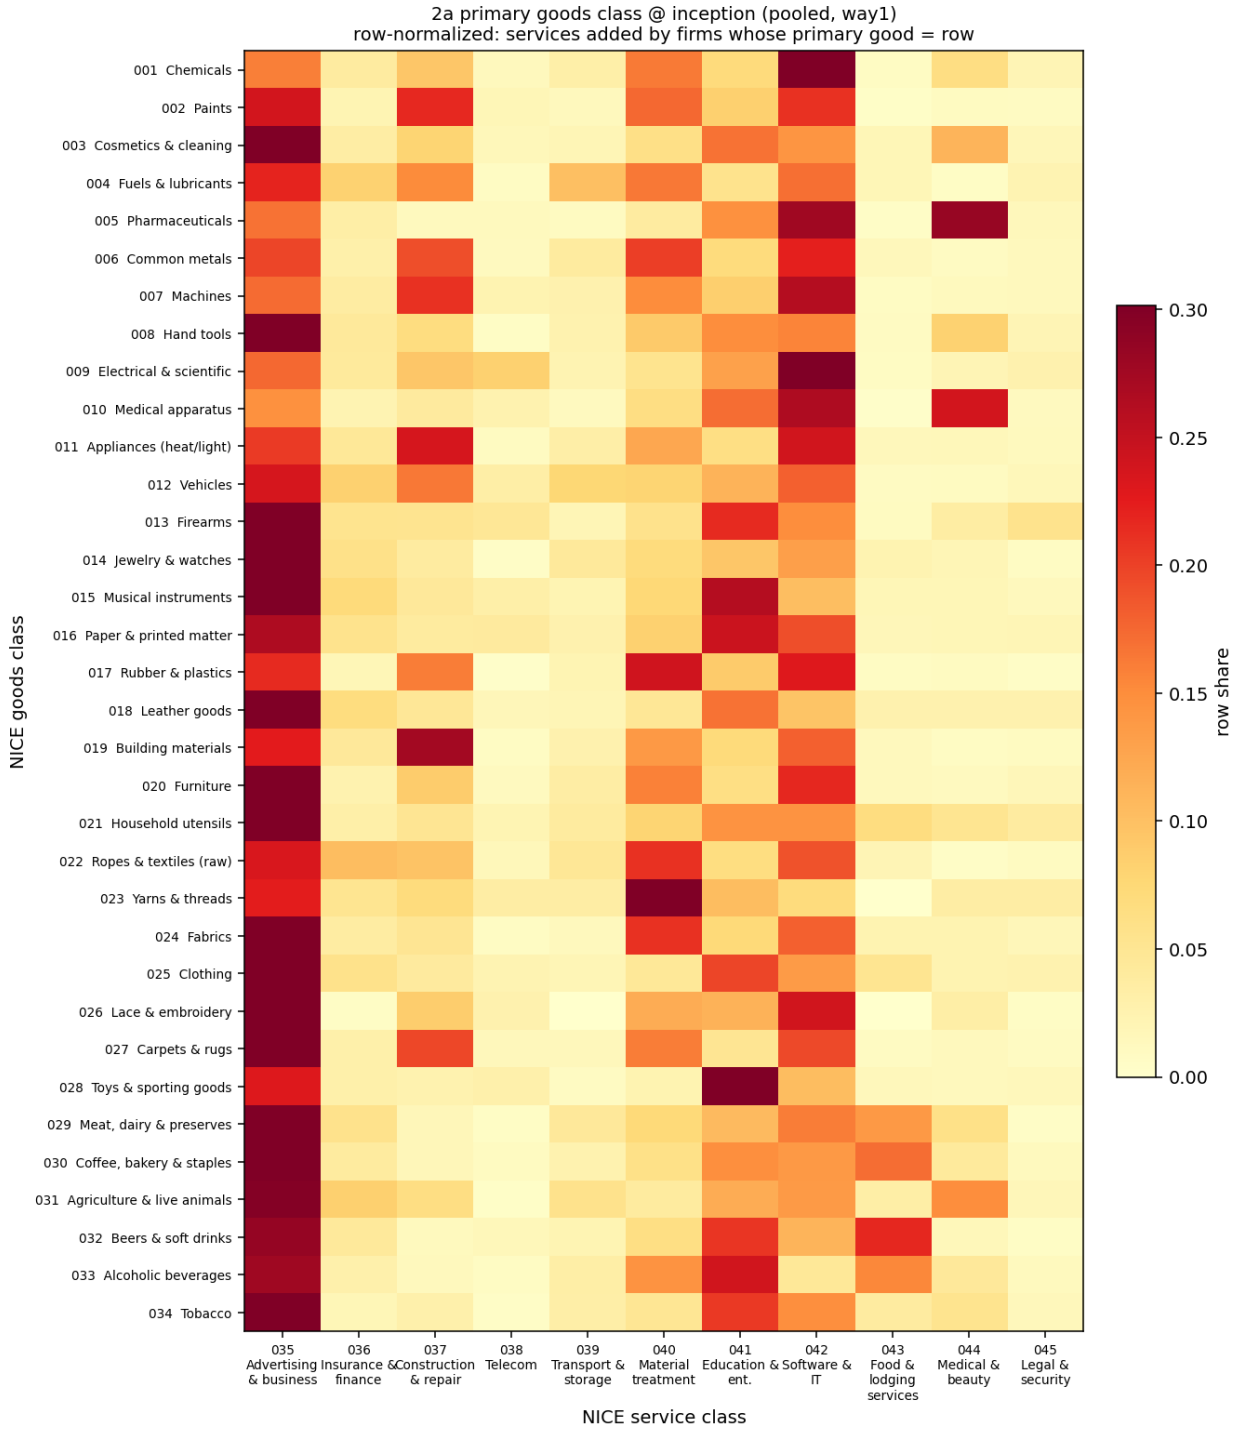


3.4 Fact 4: Firms Grow after Adding Services

The final motivating fact links first service entry to subsequent firm growth. The event is a goods firm's first introduction of a service or bundle trademark in year g, measured using both filing year and first use in commerce. Following the local-projection approach of Jordà

Figure 5
Goods-Service Pairings

The figure presents the mapping from goods classes to added service classes. Rows are the 34 NICE goods classes and columns the 11 NICE service classes. Each goods firm is assigned a single goods row, its modal NICE goods class among the goods trademarks held at or before the year of its first service addition, and every service that firm subsequently adds is attributed to that row. Cells are row-normalized, so each row sums to one. The sample period is 1990 to 2020.



(2005), we estimate stacked cohort-specific responses:

$$y_{i,g+h} - y_{i,g-1} = \beta_h \, \text{Servitize}_{ig} + \Gamma_h' X_{i,g-1} + \delta_g + \delta_{s(g-1)} + \varepsilon_{igh}, \qquad h \in \{-3, -2, 0, 1, 3, 5\}. \quad (3)$$

The outcomes are log employment and log sales. The baseline control group, $C0$, consists of goods firms that do not servitize through $g + 5$. A tighter control group, CG, consists of goods firms that add goods trademarks at g but do not add services through $g + 5$, which nets out generic product-scope expansion. Controls include lagged goods-trademark stock, lagged real activity, and leave-one-out NAICS2 service-trademark stock; specifications include cohort, sector, and year fixed effects, with standard errors clustered by firm.

Figure 6 shows the local projections using both clean non-servitizing controls and goods-trademark adders as a tighter comparison group. Relative to non-servitizing goods firms, employment rises steadily after the first service introduction. On the use-year clock, the headline effect is about 5 percent at event time, 11 percent after one year, 21 percent after three years, and roughly 26 percent by year five. The comparison to goods-trademark adders is smaller, as expected, but remains positive; this suggests that service entry is not merely a marker of generic trademark expansion. The central pattern is that service entry is followed by a persistent expansion in firm scale, rather than a short-lived filing-year spike.

Selection into servitization is still an important concern. Figures 7–8 therefore report a matched local-projection design. Each treated firm is matched to five non-servitizing controls within the same NAICS3 industry and baseline year. The matching ladder first matches on employment level, then on the pre-event employment growth path, then adds trademark-stock growth, and finally adds firm age. The key movement is from level matching to pre-trend matching: the five-year employment response falls from about 27 percent to about 24 percent, indicating selection on prior growth. Adding trademark trajectory and age changes little after that. The matched evidence therefore supports the interpretation that servitizing firms grow more after service entry, while making clear that the clean baseline partially reflects selection on pre-existing growth.

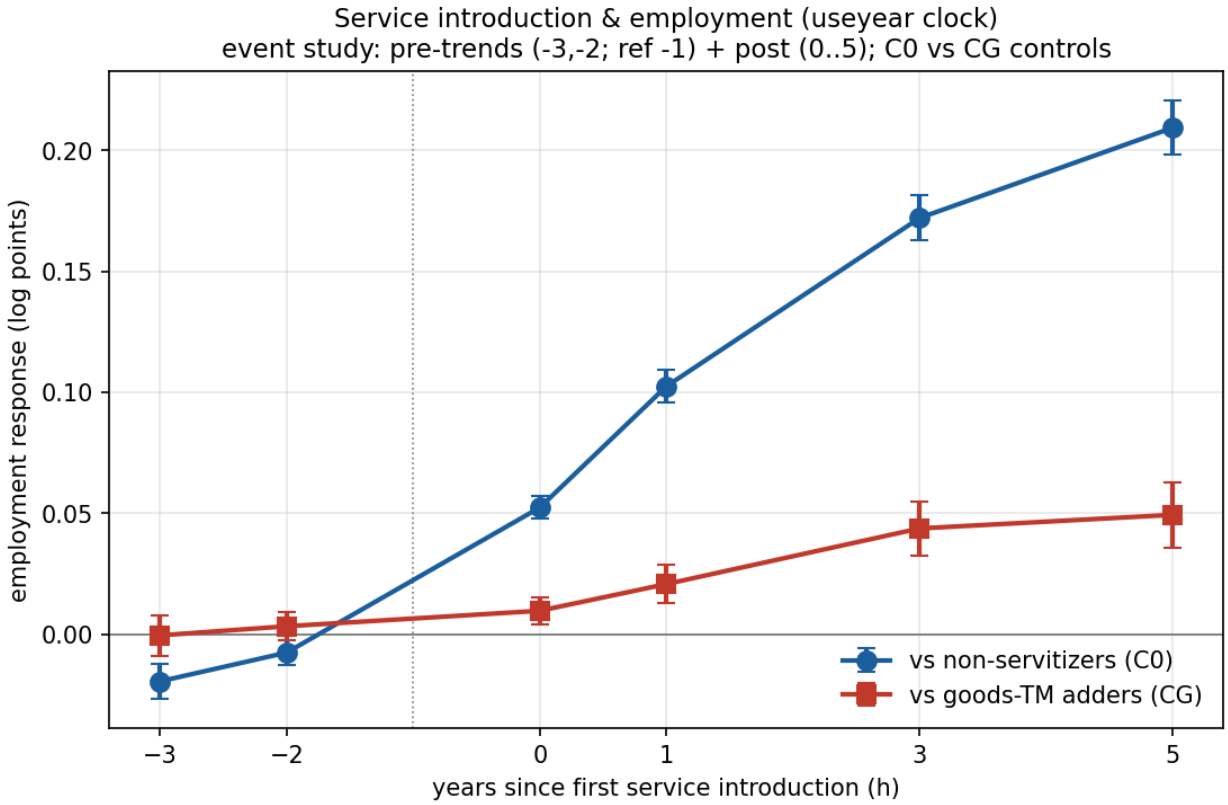
As a continuous-dose complement, we also estimate intensive-margin local projections among all goods firms:

$$y_{i,t+h} - y_{i,t-1} = \beta_h \, Z(\text{Service Flow}_{it}) + \Gamma_h' X_{i,t-1} + \text{FE} + \varepsilon_{ith}, \quad (4)$$

where service flow is the number of new service plus bundle trademarks in year t, standardized so that β_h is the response to a one-standard-deviation increase. Figures 9–10 report the visual evidence for employment and sales, contrasting a cross-firm sector-by-time specification with a within-firm specification. The responses are smaller than the first-entry event-study effects,

Figure 6
Employment after First Service Introduction

The figure presents stacked local-projection estimates of log employment around a goods firm's first service introduction, on the first-use-in-commerce clock. The blue series compares servitizing goods firms to goods firms that do not servitize through $g + 5$; the red series compares them to goods firms that add goods trademarks at g but no services through $g + 5$. Controls are the lagged goods-trademark stock, lagged real activity, and the leave-one-out NAICS2 service-trademark stock; specifications include cohort, sector, and baseline-year fixed effects. Standard errors are clustered at the firm level and bands are 95% confidence intervals. Horizon $h = -1$ is the reference.



as expected for a one-year filing flow measured among all goods firms, but they remain positive over the medium run. This exercise supports the event-study evidence by showing that growth is not only associated with crossing the first-service threshold; it also moves with the intensity of service additions.

Figure 7
Matched Employment Responses

The figure presents matched local-projection estimates of log employment. Each treated goods firm is matched to five non-servitizing controls, with replacement, requiring an exact match on 3-digit NAICS and baseline year. The matching ladder is nearest on the employment level, then on the pre-event employment growth path, then adding trademark-stock growth, then adding firm age. Standard errors are clustered at the firm level and bands are 95% confidence intervals.

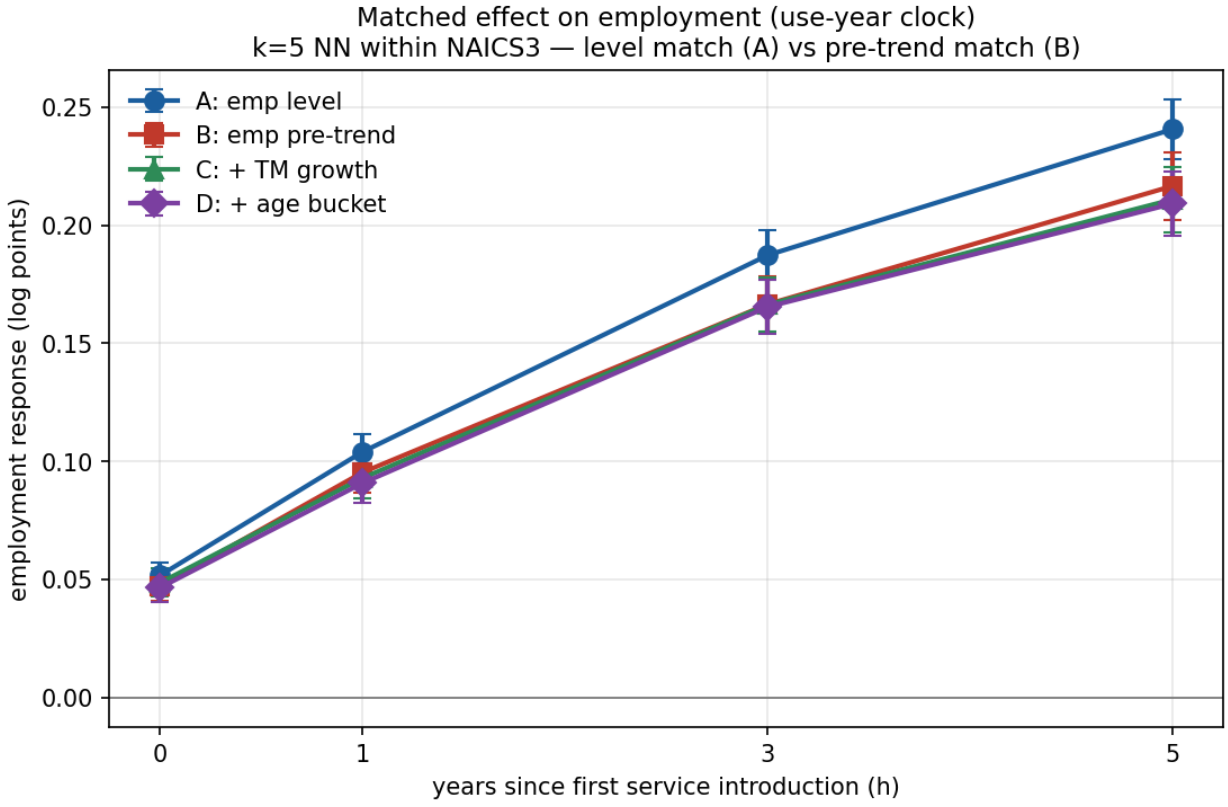


Figure 8
Matched Sales Responses

The figure presents matched local-projection estimates of log sales, using the same matching ladder as Figure 7. In NETS, sales are imputed from employment for a large share of firms, so the sales responses are corroborating rather than independent evidence. Standard errors are clustered at the firm level and bands are 95% confidence intervals.

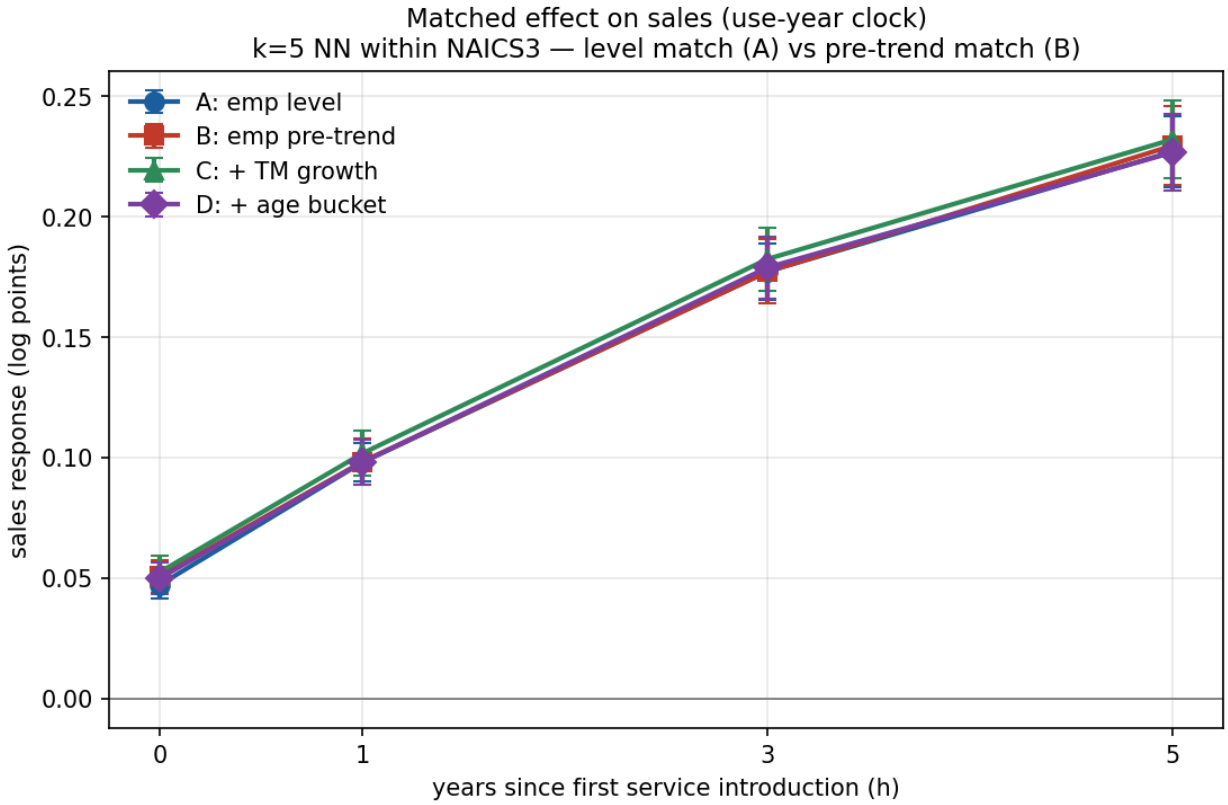


Figure 9
Intensive-Margin Employment Responses

The figure presents local projections of log employment among all goods firms, on the first-use-in-commerce clock. The treatment variable is standardized service flow, the number of new pure-service and bundle trademarks in year t, so the coefficient is the response to a one-standard-deviation increase. The figure contrasts a sector-by-time fixed-effects specification with a within-firm specification. Standard errors are clustered at the firm level and bands are 95% confidence intervals.

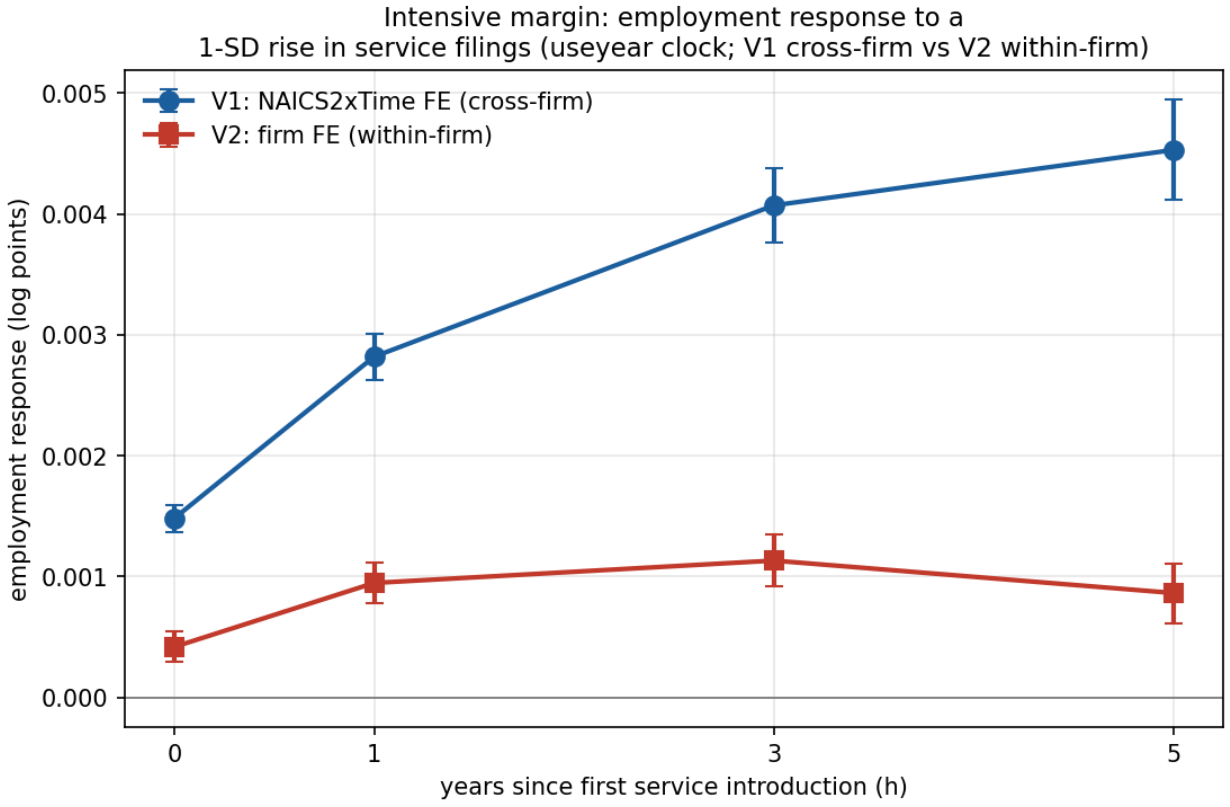
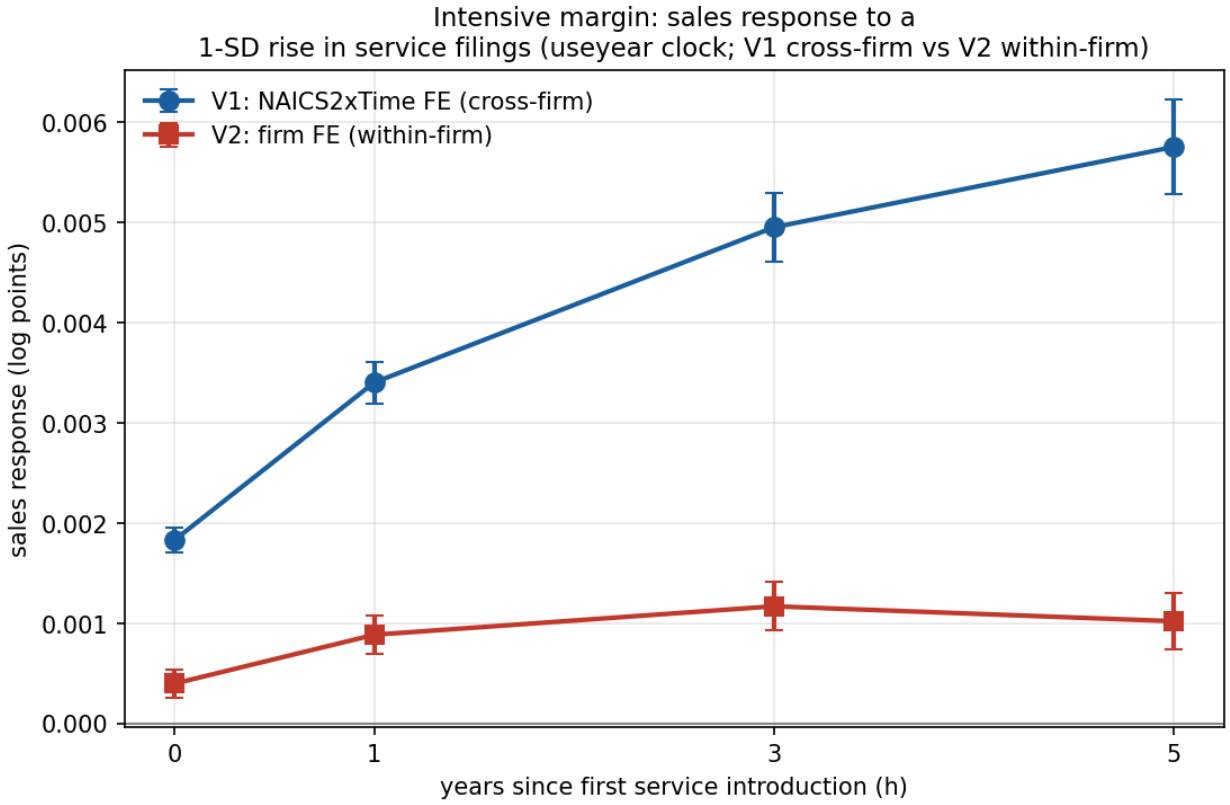


Figure 10
Intensive-Margin Sales Responses

The figure presents local projections of log sales among all goods firms, with standardized annual service flow as the treatment variable, contrasting a sector-by-time fixed-effects specification with a within-firm specification. In NETS, sales are imputed from employment for a large share of firms, so the sales responses are corroborating rather than independent evidence. Standard errors are clustered at the firm level and bands are 95% confidence intervals.

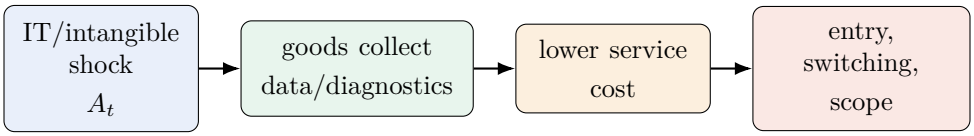


4 Model

The empirical facts show that servitization is not only the growth of services in aggregate. The object to explain is that goods producers increasingly add service lines and bundles. The model provides a production-side mechanism for this change in firm scope: goods become useful service-production technologies when an intangible layer turns product use into data, information, or diagnostics. The core idea is that many services combine two inputs. One input is information or diagnostics: measuring a machine, monitoring a patient, predicting a fault, routing a delivery, matching a customer, or identifying what needs to be done. The other input is execution time: repair, treatment, delivery, intervention, or other customer-facing service work. The model asks when a producer of good g can use goods-generated diagnostics to become an efficient provider of service j. The mechanism follows the logic summarized in the slide deck: This formulation deliberately places the key good-service

Figure 11
Main Model Mechanism

The figure presents the model mechanism. An IT or intangible shock A_t raises the productivity of product-generated data and diagnostics. For product-service pairs in which those diagnostics matter, goods producers become lower-cost service providers, which delivers service entry, switching away from standalone providers, and expanded firm scope.



complementarity in production rather than preferences. Consumers may value both goods and services, but they value a service flow regardless of whether it is supplied by a standalone service firm or by a goods producer. Servitization matters because it changes the cost and price of supplying that service flow. The exposition proceeds in the same order as the model slides: consumers and service demand; service production and standalone providers; goods production and product-service amenability; goods-generated data; bundled service costs and switching; entry, scope, and market-power implications; and finally quantification.

4.1 Goods Lines and Service Lines

There is a set of goods lines and a set of service lines:

$$\mathcal{G} = \{1, \ldots, G\}, \qquad \mathcal{J} = \{1, \ldots, J\}. \tag{5}$$

Goods lines are indexed by $g \in \mathcal{G}$. Service lines are indexed by $j \in \mathcal{J}$. The primitive question is when a producer of a good line can also become an efficient provider of a related service line.

The model separates three objects:

1. a service line's diagnostic intensity, α_j;

2. a service line's amenability to product-based servitization, m_j;

3. a goods producer's product scale, x_g.

The first object describes the service technology. The second describes whether product-generated information is useful for that service. The third describes the scale of product use generated by demand for the associated good.

4.2 Consumers

Consumers value goods and services. To keep notation light, suppress the consumer index. Let x_g denote consumption or use of good g, and let s_j denote consumption of service j. A convenient representation is a nested CES system:

$$U = \left[\omega_X X^{\frac{\rho-1}{\rho}} + \omega_S S^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}}, \tag{6}$$

where the goods aggregate is

$$X = \left(\sum_{g\in\mathcal{G}} \varphi_g x_g^{\frac{\varepsilon_X-1}{\varepsilon_X}}\right)^{\frac{\varepsilon_X}{\varepsilon_X-1}}, \tag{7}$$

and the services aggregate is

$$S = \left(\sum_{j\in\mathcal{J}} \psi_j s_j^{\frac{\varepsilon_S-1}{\varepsilon_S}}\right)^{\frac{\varepsilon_S}{\varepsilon_S-1}}. \tag{8}$$

The important point is that goods and services both enter consumer demand, but the servitization complementarity is not a direct preference complementarity. Consumers value the service flow s_j, regardless of whether it is provided by a standalone service firm or by a goods producer. Servitization matters because it changes the cost and price of providing that service flow. Let E_X denote total expenditure allocated to the goods nest. The goods

price index implied by the CES aggregate is

$$P_X = \left(\sum_{g \in \mathcal{G}} \varphi_g^{\varepsilon_X} P_g^{1-\varepsilon_X} \right)^{\frac{1}{1-\varepsilon_X}}. \tag{9}$$

Conditional demand for good line g is

$$x_g = \varphi_g^{\varepsilon_X} \left(\frac{P_g}{P_X} \right)^{-\varepsilon_X} \frac{E_X}{P_X}. \tag{10}$$

Thus the product scale of good g is generated directly by consumer demand. With homogeneous consumers and CES goods demand, every good with finite price is used in positive quantity. The baseline therefore does not introduce a separate customer-count variable. What matters for servitization is the size of existing product use, x_g: larger x_g means more product interactions, more data, and more opportunities to build product-based diagnostics. Similarly, let E_S denote total expenditure allocated to the service nest. The service price index is

$$P_S = \left(\sum_{j \in \mathcal{J}} \psi_j^{\varepsilon_S} p_j^{1-\varepsilon_S} \right)^{\frac{1}{1-\varepsilon_S}}. \tag{11}$$

Conditional demand for service line j is

$$s_j = \psi_j^{\varepsilon_S} \left(\frac{p_j}{P_S} \right)^{-\varepsilon_S} \frac{E_S}{P_S}. \tag{12}$$

This is the exact service-side counterpart of the goods demand equation above. When evaluating service demand at a counterfactual posted price p, write $s_j(p)$ for the right-hand side of this equation with $p_j = p$. The own-price elasticity of demand for service line j is ε_S. The CES nest determines demand across service lines. It does not make standalone and bundled provision two differentiated varieties within the same line. To make the within-line Bertrand margin explicit, momentarily restore the consumer index. Consumers value goods and services separately. For service line j,

$$s_{ij} = s_{ij}^S + s_{ij}^B, \tag{13}$$

where S denotes standalone service provision and B denotes bundled service provision through the associated good. The consumer wants the service flow s_{ij}; standalone and bundled provision are alternative suppliers of that same service line, not separate varieties inside the CES service nest. With homogeneous consumers and unit mass, this becomes $s_j = s_j^S + s_j^B$ in aggregate.

The baseline is a static switching model. Consumers use x_g units of good g and may demand service line j. If service j is attached to good g, later denoted $g(j)$, the bundled add-on is supplied along that existing use of the good. The first pass compares incremental service prices, taking x_g as given, in order to isolate the service-provider switching margin. The good has already been purchased on the goods margin, so this comparison does not charge the consumer twice for the good. A richer package-choice closure would compare buying the good plus the service add-on against buying standalone service without that good. Later sections discuss this richer closure in which the service add-on raises the value of the good and feeds back into x_g.

Within a service line, standalone and bundled provision are perfect substitutes along the existing use of good g. Providers compete in prices, and service demand is served by the lowest-price supplier. If a standalone provider posts p_j^S and the goods producer posts p_{gj}^B, the transaction price for service line j on that margin is

$$p_j = \min\{p_j^S, p_{gj}^B\}. \tag{14}$$

If $p_j^S < p_{gj}^B$, the standalone provider serves $s_j(p_j^S)$ and the goods producer serves zero. If $p_{gj}^B < p_j^S$, the goods producer serves $s_j(p_{gj}^B)$ and the standalone provider serves zero. This is the within-service-line Bertrand or cheapest-provider margin that generates service-provider switching.

For a simple posted-price benchmark, provider offer prices can be written as a common markup over marginal cost:

$$p = \mu_S c, \qquad \mu_S = \frac{\varepsilon_S}{\varepsilon_S - 1}. \tag{15}$$

With a common markup, the provider with lower marginal cost posts the lower price and captures the service line. This is a simple pricing-rule benchmark that keeps the CES demand system transparent. In a fully strategic Bertrand version with perfect substitutes, the low-cost provider may instead limit-price just below the rival's price; the service-provider switching threshold is still governed by relative costs, but the profit expression changes. Departures such as captive add-on pricing, switching costs, or tie-in can be added later by allowing provider-specific demand elasticities.

4.3 Standalone Service Production

Start with the service sector alone. A service provider produces service line j using two inputs:

1. information or diagnostic input, denoted I_j;

2. execution input, denoted H_j.

Diagnostic input is the part of the service that figures out what needs to be done. Execution input is the part that carries out the service. For repair, diagnostics identify the failure and execution fixes it. For health care, diagnostics include measurement, monitoring, and assessment, while execution includes treatment and patient-facing time. For logistics, diagnostics include routing, prediction, and tracking, while execution includes movement, handling, and delivery. For software support, diagnostics identify the user problem or system state, while execution resolves it.

Service output is

$$q_j^S = Z_j \left(I_j\right)^{\alpha_j} \left(H_j\right)^{1-\alpha_j}, \qquad \alpha_j \in [0,1]. \tag{16}$$

The statistic α_j is central. High-α_j services are information intensive: diagnosis, monitoring, matching, codification, or prediction is a large part of production. Low-α_j services are execution intensive: physical presence, direct human time, or manual delivery is a large part of production.

For a standalone service provider, diagnostic input and execution input are both produced with service time. Let L_j denote total service time and let $\tau_j \in [0,1]$ be the share allocated to diagnostics. Then

$$I_j = \tau_j L_j, \qquad H_j = (1-\tau_j) L_j. \tag{17}$$

Thus standalone production can be written as

$$q_j^S = Z_j \tau_j^{\alpha_j} (1-\tau_j)^{1-\alpha_j} L_j. \tag{18}$$

Let the wage be w. The standalone unit cost of service line j is

$$c_j^S = \min_{\tau \in [0,1],\, L \geq 0} wL \quad \text{s.t.} \quad Z_j \tau^{\alpha_j} (1-\tau)^{1-\alpha_j} L = 1. \tag{19}$$

Equivalently,

$$c_j^S = \min_{\tau \in [0,1]} \frac{w}{Z_j \tau^{\alpha_j} (1-\tau)^{1-\alpha_j}}. \tag{20}$$

For an interior service line, $\alpha_j \in (0,1)$, the optimal standalone diagnostic share is

$$\tau_j^{S,*} = \alpha_j, \tag{21}$$

and therefore

$$c_j^S = \frac{w}{Z_j \alpha_j^{\alpha_j} (1-\alpha_j)^{1-\alpha_j}}. \tag{22}$$

Holding w and Z_j fixed, this standalone unit cost is not globally increasing in diagnostic intensity:

$$\frac{\partial \log c_j^S}{\partial \alpha_j} = \log\left(\frac{1-\alpha_j}{\alpha_j}\right). \tag{23}$$

Thus c_j^S rises with α_j when $\alpha_j < 1/2$, is maximized at $\alpha_j = 1/2$, and falls when $\alpha_j > 1/2$. This is a feature of the symmetric time-allocation benchmark: diagnostic and execution time are equally costly, so changing α_j changes the optimal time split rather than making diagnostic-intensive services mechanically more expensive. If we want high-α_j services to have higher standalone costs, we can let diagnostic productivity or Z_j vary with α_j, or allow diagnostic time to be more costly than execution time.

The equality simply normalizes the unit-cost problem to one unit of service output. The time-allocation formulation makes the interpretation of α_j especially transparent: absent product-generated information, a high-α_j service is one where the provider spends a large share of time diagnosing rather than executing.

4.4 Goods Lines, Product Demand, and Product Scale

Now introduce goods producers. A producer of good line g sells a good with product appeal or quality φ_g. Demand for the good is the CES demand x_g derived above. The baseline scale of good line g is simply x_g. High-φ_g goods have larger product scale through CES demand, and larger product scale makes product-based service provision more data rich. This is the object we need for the servitization mechanism in the representative-consumer benchmark.

The static model treats x_g as inherited by the service decision. The good price P_g is paid on the goods margin, while the service-switching margin below is an incremental service-price comparison along the existing use of the good. This keeps the baseline close to the representative-consumer CES structure: all goods are used in positive amount, and servitization changes who provides the associated service flow.

4.4.1 Extension: Heterogeneous Goods Adoption

A richer extension can add an explicit extensive margin and a literal customer base. Suppose consumers i have heterogeneous tastes for goods, so only some consumers buy or use good g. Then

$$N_g = \int \mathbf{1}\{i \text{ uses good } g\}\, di \tag{24}$$

is a true customer count. In that version, product scale can depend on both the number of customers and their usage intensity. The service add-on can make the good more attractive: consumers who would otherwise buy standalone service may choose the good because it

provides access to a cheaper or better bundled service. Formally, this can be represented by an effective product appeal φ_g^{eff} that rises when the service add-on is available, or by a full package choice comparing standalone service with buying good g plus the service add-on.

That extension creates an additional feedback loop: servitization can expand demand for the good, which raises N_g or x_g, improves product-generated diagnostics, and increases the return to further service provision. The current baseline shuts down this extensive margin so that the baseline isolates the intensive service-provider switching mechanism.

4.4.2 Goods Production

The baseline model does not need rich goods-side heterogeneity. A simple production function is enough:

$$q_g^G = Z^G \ell_g^G, \tag{25}$$

and one can normalize $Z^G = 1$. In this version, differences in product scale come from φ_g, not from a separate firm productivity term.

This keeps the model disciplined. The same object, φ_g, can be interpreted as product quality, reputation, design, ecosystem value, or consumer appeal. These forces determine product demand x_g. If later we need to separate low production cost from high product appeal, we can introduce line-specific goods productivity Z_g^G. For now, adding both Z_g^G and φ_g would create two sources of goods-side heterogeneity when only one is needed.

4.5 Product-Service Amenability

Each service line has at most one potential associated goods line. Let $g(j)$ denote the good line that could generate information relevant for service j. Some services have no meaningful product channel; for them, $g(j)$ can be left inactive and $m_j = 0$. Let $m_j \in \{0, 1\}$. If $m_j = 1$, service j is amenable to servitization through the associated good line $g(j)$. If $m_j = 0$, the service is not amenable. The good may have many users, and the producer may be technologically sophisticated, but product-generated information does not help produce this service.

This binary version is deliberately stark. It separates two ideas:

1. α_j asks how diagnostic intensive the service is;
2. m_j asks whether the relevant diagnostics can be generated by a product-based information layer.

Both are needed. A service can be diagnostic intensive but not product amenable. Another service can be product amenable but mostly execution intensive, so the cost reduction is

small.

4.6 Goods Producers Without the Intangible Layer

A goods producer that does not invest in an intangible layer simply produces the good:

$$q_g^G = Z^G \ell_g^G. \tag{26}$$

If $m_j = 0$, the associated goods producer cannot turn the good into a service-delivery technology for line j. If $m_j = 1$ but the goods producer has not invested in the intangible layer, it also does not provide the service add-on. Consumers who want service j buy it from standalone service providers at price p_j^S.

Thus production of a goods line is not enough. The goods producer needs both an amenable service line and an intangible layer that turns product use into diagnostic input. Without that layer, there is no bundled service price and no service-provider switching margin.

4.7 The Intangible Layer

A producer of good line g can invest in an intangible layer. This layer includes software, sensors, data infrastructure, monitoring systems, customer accounts, service routines, analytics, and organizational capability. The layer is costly to build, but it lets the goods producer use product demand and product use, x_g, to generate diagnostic information.

Let A_t denote the economy-wide intangible or IT productivity shock. A higher A_t means that data storage, software, cloud infrastructure, sensors, communication, and codification make a given intangible investment more effective. This is the place where the IT revolution enters the model: mainly through the productivity of K_g in generating product-based diagnostics, and, separately, through line-specific setup costs if desired.

Let K_g denote the producer's intangible investment. This investment is not free. Let $C_g^K(K_g)$ denote the cost of building the stock of intangible capability used across service lines. We can think of this as the cost of software, data systems, sensors, customer-account infrastructure, and organizational routines. It is increasing and convex in K_g. A simple reduced form is

$$C_g^K(K_g) = \bar{c}_g^K \frac{K_g^{1+\eta}}{1+\eta}, \qquad \eta > 0. \tag{27}$$

For service line j, product-generated diagnostic input is

$$\Omega_{gj}(A_t, K_g, x_g) = m_j \lambda_j A_t K_g^{\kappa} x_g^{\nu}, \qquad \kappa, \nu \geq 0. \tag{28}$$

The term m_j is the technological switch. If $m_j = 0$, then

$$\Omega_{gj}(A_t, K_g, x_g) = 0. \tag{29}$$

The service cannot be servitized through product-generated information. If $m_j = 1$, then the diagnostic input supplied by the product and intangible layer rises with the IT shock A_t, the firm's intangible investment K_g, and product scale x_g. This object is the bundled counterpart of diagnostic input I_j: in bundled provision, the good supplies diagnostics and human labor supplies execution.

The x_g term captures a data and scale channel. Larger product use can generate more observations, more monitoring information, more repeated interactions, and more opportunities to amortize fixed service routines. This is also where the model begins to point toward concentration and market power: large goods producers may have a stronger information advantage precisely because their products are used at larger scale.

In the homogeneous unit-measure economy, x_g has two equivalent readings. It is the representative consumer's use intensity of good g, and because the mass of consumers is one, it is also aggregate product use. This is why connected-product examples fit the notation. If a consumer wears a smartwatch more consistently, the product generates a richer diagnostic history for that consumer's health or monitoring service. In the homogeneous baseline, the same x_g also measures the product scale available to the firm. A heterogeneous-consumer version would separate individual use a_{ig} from aggregate product scale X_g^{use}, where $X_g^{\text{use}} = \int a_{ig}\, di$. The baseline deliberately collapses both into x_g.

4.7.1 Line-Specific Fixed Cost of Adding a Service

Introducing a service add-on requires a fixed cost:

$$F_{gj}^I(A_t, x_g). \tag{30}$$

This cost includes software development, data engineering, service-process design, compliance, integration with the product, and the organizational cost of becoming a service provider. It is separate from $C_g^K(K_g)$: $C_g^K(K_g)$ is the cost of the firm-level intangible stock that can be shared across service lines, while $F_{gj}^I(A_t, x_g)$ is the line-specific setup cost for adding service j to good g.

A useful reduced form is

$$F_{gj}^I(A_t, x_g) = \bar{F}_{gj} A_t^{-\zeta} x_g^{-\xi}, \qquad \zeta, \xi \geq 0. \tag{31}$$

The parameter ζ captures the IT revolution: higher A_t lowers the line-specific setup cost of adding the service. The parameter ξ captures the possibility that large product scale makes deployment easier or lowers the effective fixed cost of creating a scalable service system. If we want to be more conservative, we can set $\xi = 0$ and let x_g matter only through diagnostic information quality.

4.8 Bundled Service Production

When a goods producer with $m_j = 1$ invests in the intangible layer, it can produce a service add-on using the same service production function as standalone service providers. The difference is the source of diagnostic input.

For bundled provision, the product and intangible layer provide the diagnostic input:

$$I_{gj}^B = \Omega_{gj}(A_t, K_g, x_g). \tag{32}$$

Human labor is used only for execution:

$$H_{gj}^B = L_{gj}. \tag{33}$$

Service output is therefore

$$q_{gj}^B = Z_j \Omega_{gj}(A_t, K_g, x_g)^{\alpha_j} L_{gj}^{1-\alpha_j}. \tag{34}$$

For one unit of bundled service, the normalized production condition is

$$Z_j \Omega_{gj}(A_t, K_g, x_g)^{\alpha_j} L_{gj}^{1-\alpha_j} = 1. \tag{35}$$

The same production function is used in both sectors. The difference is that the goods producer's intangible layer supplies the diagnostic input. The goods producer still needs execution time. Servitization replaces human diagnostic labor with product-generated information; it does not eliminate the execution part of service work. If Ω_{gj} is very low, bundled production is extremely costly or infeasible.

4.9 Cost Comparison

The standalone service unit cost is

$$c_j^S = \min_{\tau \in [0,1]} \frac{w}{Z_j \tau^{\alpha_j} (1-\tau)^{1-\alpha_j}} = \frac{w}{Z_j \alpha_j^{\alpha_j} (1-\alpha_j)^{1-\alpha_j}}. \tag{36}$$

The goods producer's bundled service unit cost, conditional on having invested in the intangible layer and offering the service add-on, is

$$c_{gj}^{B}(\Omega_{gj}) = \min_{L \geq 0} wL \quad \text{s.t.} \quad Z_j \Omega_{gj}^{\alpha_j} L^{1-\alpha_j} = 1. \tag{37}$$

For $\Omega_{gj} > 0$ and $\alpha_j \in (0,1)$, this is

$$c_{gj}^{B}(\Omega_{gj}) = w \left(\frac{1}{Z_j \Omega_{gj}^{\alpha_j}} \right)^{\frac{1}{1-\alpha_j}}. \tag{38}$$

This is an operating unit cost. It deliberately does not include the cost of choosing K_g. The cost of K_g enters the firm's entry and investment problem below. Equivalently, the equation above asks: given the intangible capability already in place, how much execution labor is needed to deliver one unit of the bundled service?

If $m_j = 0$, or if the goods producer does not build the intangible layer, this bundled technology is not available. The service is supplied by standalone providers. For notation, set $c_{gj}^{B}(0) = +\infty$ and $p_{gj}^{B}(0) = +\infty$. If $m_j = 1$, $K_g > 0$, and $A_t > 0$, then $\Omega_{gj} > 0$. When Ω_{gj} is very small, bundled production is extremely costly. As Ω_{gj} rises, bundled cost falls:

$$\frac{\partial \log c_{gj}^{B}(\Omega_{gj})}{\partial \log \Omega_{gj}} = -\frac{\alpha_j}{1-\alpha_j} < 0. \tag{39}$$

The cost saving is

$$\Delta c_{gj} = c_j^{S} - c_{gj}^{B}(\Omega_{gj}). \tag{40}$$

It is increasing in the product-generated diagnostic input:

$$\frac{\partial \Delta c_{gj}}{\partial \Omega_{gj}} > 0. \tag{41}$$

Since

$$\Omega_{gj}(A_t, K_g, x_g) = m_j \lambda_j A_t K_g^{\kappa} x_g^{\nu}, \tag{42}$$

the cost saving is larger when the product-service pair is amenable, $m_j = 1$, and when product-generated diagnostic input rises through higher intangible productivity A_t, larger intangible investment K_g, or larger product scale x_g. Diagnostic intensity α_j governs how valuable Ω_{gj} is in service production.

This is the core production-side result. The IT shock does not make all services cheaper for goods producers. It lowers the cost of services whose diagnostic component can be partly produced by a product-based intangible layer.

Proposition: Cost Threshold and Service-Provider Switching. Fix a service line j and its potential associated good $g(j)$. Suppose $m_j = 1$, $\alpha_j \in (0,1)$, and the bundled service can be supplied along existing use of good $g(j)$. Let

$$p_j^S = \mu_S c_j^S, \qquad p_{gj}^B(\Omega_{gj}) = \mu_S c_{gj}^B(\Omega_{gj}), \qquad \mu_S = \frac{\varepsilon_S}{\varepsilon_S - 1}. \tag{43}$$

Within service line j, consumers choose the lower posted price:

$$p_j(\Omega_{gj}) = \min\{p_j^S, p_{gj}^B(\Omega_{gj})\}. \tag{44}$$

Because $c_{gj}^B(\Omega_{gj})$ is decreasing in Ω_{gj}, there is a threshold level of product-generated diagnostic input $\bar{\Omega}_{gj}$ defined by

$$\mu_S c_{gj}^B(\bar{\Omega}_{gj}) = \mu_S c_j^S, \qquad \text{or equivalently} \qquad c_{gj}^B(\bar{\Omega}_{gj}) = c_j^S. \tag{45}$$

For $\Omega_{gj} < \bar{\Omega}_{gj}$, bundled service is infeasible or too costly, so the standalone provider posts the lower price and serves $s_j(p_j^S)$, while the goods producer serves zero. For $\Omega_{gj} > \bar{\Omega}_{gj}$, the goods producer posts the lower price and serves $s_j(p_{gj}^B(\Omega_{gj}))$, while the standalone provider serves zero. This is the service-provider switching margin. The good price P_g is not included in this threshold because the baseline comparison is conditional on the existing use x_g. Since

$$\Omega_{gj} = m_j \lambda_j A_t K_g^{\kappa} x_g^{\nu}, \tag{46}$$

the threshold can be crossed by an increase in the economy-wide IT/intangible shock A_t, the firm's intangible investment K_g, or product scale x_g.

The same switching logic can be drawn in (α_j, A_t) space. Holding K_g and x_g fixed, define the switching boundary $A_j^*(\alpha_j; K_g, x_g)$ implicitly by

$$p_{gj}^B\Big(\Omega_{gj}(A_j^*(\alpha_j; K_g, x_g), K_g, x_g); \alpha_j\Big) = p_j^S(\alpha_j). \tag{47}$$

For $m_j = 1$, the goods producer wins the within-line Bertrand comparison whenever

$$A_t \geq A_j^*(\alpha_j; K_g, x_g). \tag{48}$$

This fixed-K_g version is the cleanest way to see the comparative statics. Higher α_j makes product-generated diagnostics more important in bundled production, but the exact effect of α_j on the switching threshold also depends on the standalone cost benchmark $c_j^S(\alpha_j)$ and the diagnostic productivity term λ_j. The figure therefore illustrates the empirically relevant

Figure 12
Switching threshold

The horizontal axis is product-generated diagnostic input, $\Omega_{gj} = m_j \lambda_j A_t K_g^{\kappa} x_g^{\nu}$. Higher A_t, K_g, or x_g moves the producer to the right. Higher α_j makes the bundled price more sensitive to Ω_{gj}; in the illustrated case, diagnostic-intensive services cross the Bertrand switching threshold earlier. To the left of the threshold, standalone providers serve service demand; to the right, the goods producer serves it.

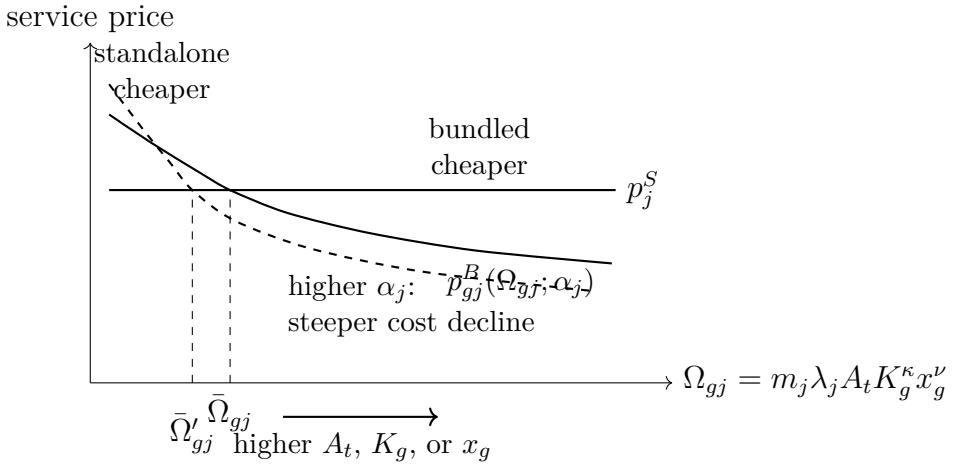
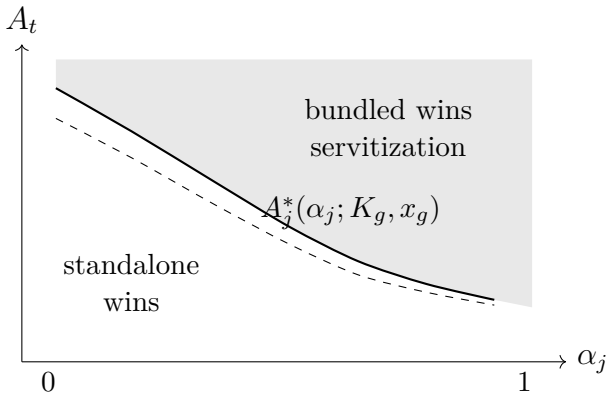


Figure 13
Switching region

Switching region in (α_j, A_t) space for fixed K_g and x_g. The shaded area is where the goods producer's bundled service price is below the standalone service price, so service demand switches to the goods producer. The downward-sloping boundary is the case in which diagnostic-intensive, product-amenable services are especially exposed to the IT/intangible shock. The dashed boundary illustrates how higher K_g or x_g shifts the switching threshold down.



case in which diagnostic-intensive, product-amenable services require a smaller IT/intangible shock to switch. Higher K_g or x_g shifts the whole boundary down. If K_g is chosen optimally, the boundary becomes an entry boundary defined by the maximized payoff below; that version is more appropriate for quantitative work.

4.10 Entry into Service Lines

Let p_j^S be the standalone service price and $p_{gj}^B(\Omega_{gj})$ the bundled service price:

$$p_j^S = \mu_S c_j^S, \qquad p_{gj}^B(\Omega_{gj}) = \mu_S c_{gj}^B(\Omega_{gj}). \tag{49}$$

Using the CES service demand $s_j(p)$, write

$$\Omega_{gj}^K \equiv \Omega_{gj}(A_t, K_g, x_g). \tag{50}$$

Define the bundled-provider indicator

$$\mathcal{B}_{gj}(K_g; A_t, x_g) = \mathbf{1}\left\{p_{gj}^B(\Omega_{gj}^K) \leq p_j^S\right\}. \tag{51}$$

This indicator is one only when the goods producer wins the within-line Bertrand comparison. For a given intangible stock K_g, the operating contribution of service line j is

$$\pi_{gj}(K_g; A_t, x_g) = \left[p_{gj}^B(\Omega_{gj}^K) - c_{gj}^B(\Omega_{gj}^K)\right] s_j(p_{gj}^B(\Omega_{gj}^K))\mathcal{B}_{gj}(K_g; A_t, x_g). \tag{52}$$

If the firm were considering only service line j, a pairwise entry condition would be

$$\max_{K_g \geq 0}\left\{\pi_{gj}(K_g; A_t, x_g) - C_g^K(K_g)\right\} \geq F_{gj}^I(A_t, x_g). \tag{53}$$

This pairwise condition is useful for intuition, but it is not the full scope problem because the intangible stock K_g can be shared across service lines. The firm-level formulation lets the goods producer choose both K_g and the set of service lines it offers. Let $e_{gj} \in \{0, 1\}$ indicate whether good producer g offers service line j. The joint scope problem is

$$\Pi_g^*(A_t, x_g) = \max_{K_g \geq 0,\ \{e_{gj}\}_{j \in \mathcal{J}}} \left\{\sum_{j \in \mathcal{J}} e_{gj}\left[\pi_{gj}(K_g; A_t, x_g) - F_{gj}^I(A_t, x_g)\right] - C_g^K(K_g)\right\}. \tag{54}$$

In this version, $C_g^K(K_g)$ is paid once at the firm level, while $F_{gj}^I(A_t, x_g)$ is paid for each service line that enters the firm's scope. A service line is added when its marginal contribution is positive at the chosen K_g. This is the object that should discipline quantitative work on multi-line scope and concentration.

These conditions summarize the servitization margin. Entry is more likely when:

1. the service is amenable to product-generated diagnostics, $m_j = 1$;

2. the service is diagnostic intensive, especially when high α_j is associated with high stan-

dalone diagnostic costs or high product-diagnostic productivity λ_j;

3. the IT shock A_t makes a given K_g more productive and may lower line-specific setup costs;

4. the goods producer has large product scale x_g;

5. service demand is large or price elastic;

6. the common CES markup μ_S converts a cost advantage into operating profits.

Product scale enters through the production side. It can improve diagnostic technology, and it can lower the effective setup cost if scalable intangible systems are easier to build or deploy for high-use products. In the common-markup benchmark, bundled goods producers do not earn higher profits because they charge a higher service markup. They earn higher profits because lower bundled service costs imply lower service prices and larger service quantities. In the current static model, x_g is predetermined for the service decision. The heterogeneous-adoption extension above is where service availability can feed back into goods demand and generate a true customer-count expansion.

To visualize the aggregate entry margin, define a servitization share

$$\mathcal{S}(A_t) = \frac{1}{GJ} \sum_{g \in \mathcal{G}} \sum_{j \in \mathcal{J}} e^*_{gj}(A_t, x_g), \tag{55}$$

where $e^*_{gj}(A_t, x_g)$ is the service-line choice from the joint scope problem above. As A_t rises, Ω_{gj} rises for any given K_g, the marginal return to choosing K_g rises, the optimal K_g may rise, and $F^I_{gj}(A_t, x_g)$ may fall. More product-service pairs cross the entry condition, so the servitization share increases.

4.11 Implications

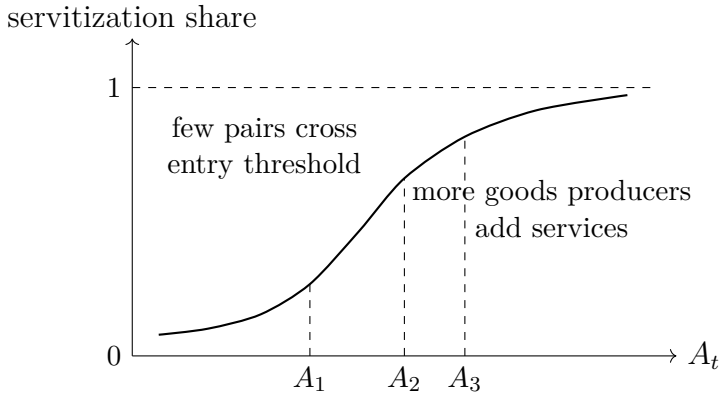
This version of the model gives a clean order of implications.

4.11.1 Scope Expansion

After an increase in A_t, goods producers are more likely to add service lines. Entry is concentrated in service lines with $m_j = 1$, high product-diagnostic productivity λ_j, and diagnostic intensity α_j that makes product-generated information valuable. It is also concentrated among goods producers with large x_g, because high product scale gives them a stronger information base.

Figure 14
Aggregate entry

A rise in the IT/intangible shock A_t raises product-generated diagnostics and can lower service-line setup costs. As more product-service pairs satisfy the entry condition, the servitization share increases.



4.11.2 Business Stealing and Market Expansion

Before servitization, service line j is supplied by the standalone provider:

$$Q_j^0 = s_j(p_j^S). \tag{56}$$

After servitization, if

$$p_{gj}^B(\Omega_{gj}) < p_j^S, \tag{57}$$

the goods producer wins the within-line Bertrand comparison and supplies

$$Q_j^1 = s_j(p_{gj}^B(\Omega_{gj})). \tag{58}$$

This quantity can be decomposed as

$$Q_j^1 = \underbrace{s_j(p_j^S)}_{\text{business stealing from standalone}} + \underbrace{\left[s_j(p_{gj}^B(\Omega_{gj})) - s_j(p_j^S)\right]}_{\text{market expansion}}. \tag{59}$$

The first term is not new service use. It is the pre-existing service demand that moves from standalone providers to the goods producer once the bundled price is lower. The second term is new service use created by the price decline. In the homogeneous baseline, this is an intensive-margin expansion: each representative consumer uses more of service j.

4.11.3 Tradeability and Access

The same accounting gives a tradeability interpretation once consumers or locations differ. In the baseline above, servitization expands service use through the intensive margin $s_j(p)$. In a richer geography or heterogeneous-preference version, product-generated diagnostics can also expand the set of consumers for whom service j is feasible. The idea is that the diagnostic component becomes more codified, remotely monitored, or embedded in the product, while execution may still require human service time.

Let $\mathcal{A}_j^S$ be the set of consumers or locations that can access standalone service j. Let $\mathcal{A}_{gj}^B$ be the set that can access the bundled, product-enabled version. This bundled access set depends on A_t, K_g, and x_g. If product-generated diagnostics make the service more tradable,

$$\mathcal{A}_j^S \subseteq \mathcal{A}_{gj}^B(A_t, K_g, x_g). \tag{60}$$

Then, when the bundled provider has the lower delivered service price,

$$\begin{aligned} Q_{gj}^{B,1} = & \underbrace{\int_{\mathcal{A}_j^S} s_{ij}(p_j^S)\, di}_{\text{business stealing from standalone}} + \underbrace{\int_{\mathcal{A}_j^S} \left[s_{ij}(p_{gj}^B(\Omega_{gj})) - s_{ij}(p_j^S) \right] di}_{\text{intensive market expansion}} \\ & + \underbrace{\int_{\mathcal{A}_{gj}^B(A_t,K_g,x_g)\setminus\mathcal{A}_j^S} s_{ij}(p_{gj}^B(\Omega_{gj}))\, di}_{\text{access and tradeability expansion}}. \end{aligned} \tag{61}$$

Thus the current model captures market expansion through lower prices and higher use. The tradeability extension adds an extensive/access margin: consumers who were previously constrained by local availability, distance, time, or diagnostic frictions can now be served because the product carries or produces the information needed to deliver the service.

4.11.4 Concentration and Market Power

Large product scale raises the return to intangible investment and may improve diagnostic technology. This can generate a feedback loop:

$$\begin{aligned} \text{large } x_g \quad &\Rightarrow \quad \text{better product-generated diagnostics} \\ &\Rightarrow \quad \text{lower service cost} \\ &\Rightarrow \quad \text{more service scope} \quad \Rightarrow \quad \text{larger ecosystem.} \end{aligned} \tag{62}$$

This loop can also run through demand for the good itself. A lower bundled service price, or a better service made possible by product-generated diagnostics, can raise the effective

value of good g. Then consumers demand more x_g, which further improves Ω_{gj}:

$$\begin{aligned} A_t \uparrow \text{ or } K_g \uparrow \quad &\Rightarrow \quad p^B_{gj} \downarrow \text{ or service quality rises} \\ &\Rightarrow \quad x_g \uparrow \\ &\Rightarrow \quad \Omega_{gj}(A_t, K_g, x_g) \uparrow \quad \Rightarrow \quad \text{further service scope.} \end{aligned} \tag{63}$$

Atomistic consumers need not internalize their effect on the firm's aggregate cost. They respond to the private value of the connected product and the posted bundled service terms. The firm internalizes the product-scale effect when it chooses K_g, service scope, and pricing. In the baseline above this feedback is held fixed; in a consumer-closure version it becomes part of the equilibrium determination of x_g. Under the common-markup posted-price benchmark, the service markup is common across standalone and bundled providers:

$$\mu_S = \frac{p^S_j}{c^S_j} = \frac{p^B_{gj}(\Omega_{gj})}{c^B_{gj}(\Omega_{gj})} = \frac{\varepsilon_S}{\varepsilon_S - 1}. \tag{64}$$

Thus the bundled producer's advantage is not a higher service markup. It is lower cost. Conditional on winning the within-line Bertrand comparison, and holding the service price index fixed, bundled operating profit is

$$\left[p^B_{gj}(\Omega_{gj}) - c^B_{gj}(\Omega_{gj})\right] s_j(p^B_{gj}(\Omega_{gj})) \propto \left[c^B_{gj}(\Omega_{gj})\right]^{1-\varepsilon_S}. \tag{65}$$

Because $\varepsilon_S > 1$, lower bundled cost raises operating profit even though the markup is unchanged. A large x_g raises the return to investing in K_g by improving product-generated diagnostics and, if $\xi > 0$, lowering effective setup costs. In this benchmark, market power shows up less as higher service markups and more as scope and concentration: large goods producers can spread the intangible layer over more product use, enter more service lines, and potentially make their ecosystem more valuable. The welfare gain from the IT shock depends on pass-through. With the CES markup fixed, some of the cost decline is passed through into lower service prices. Consumers benefit from lower service prices and from expanded service use on the intensive margin, while producers benefit from product-scale advantages in diagnostics and setup. Additional consumer costs of market power would require a departure from the common-markup CES benchmark, such as captive add-on pricing, switching costs, or tying between the good and the service.

4.12 Quantification

TBD

Figures

Figure 15
Service Items Frequency in Trademark Filings

The figure presents the word cloud for individual service items covered in the trademark data. Word size is proportional to frequency in the full USPTO trademark database. Service items are extracted from goods/services descriptions in Nice Classes 35–45.

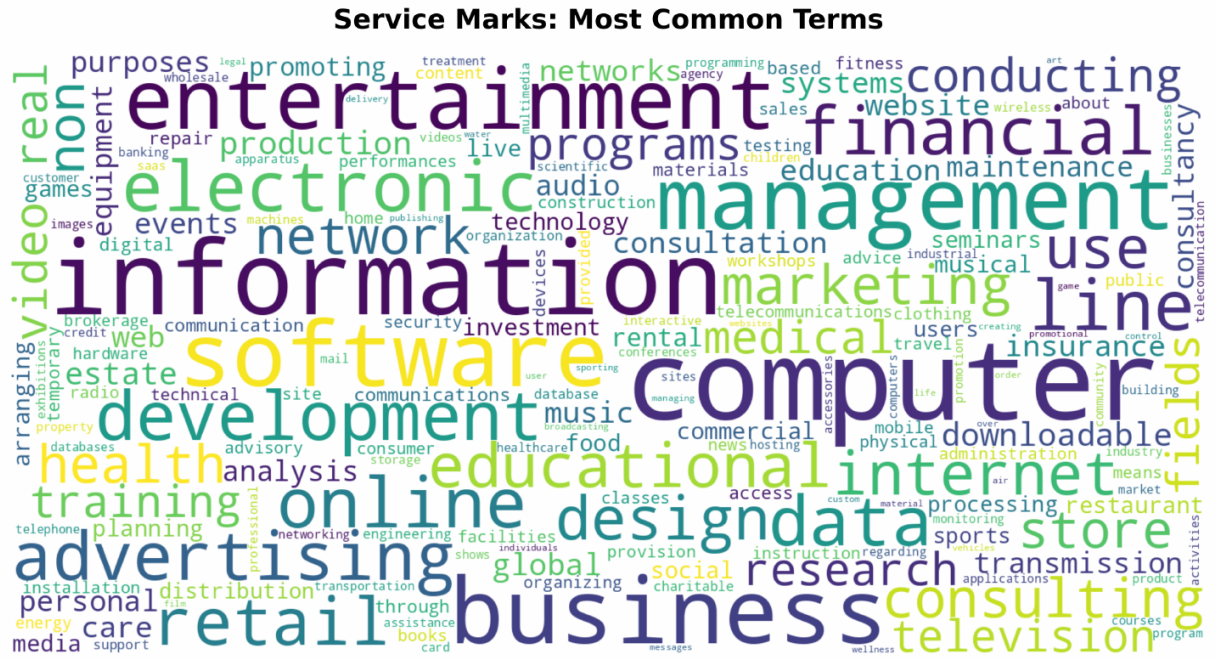


Figure 16
Trademark and Patent Innovation Indices

The figure presents standardized time-series indices for the KPSS patent innovation index (Kogan et al., 2017) and the total trademark innovation index. Both series are expressed as $(x_t - \bar{x})/sd(x)$ and therefore have mean zero and standard deviation one over the sample series.

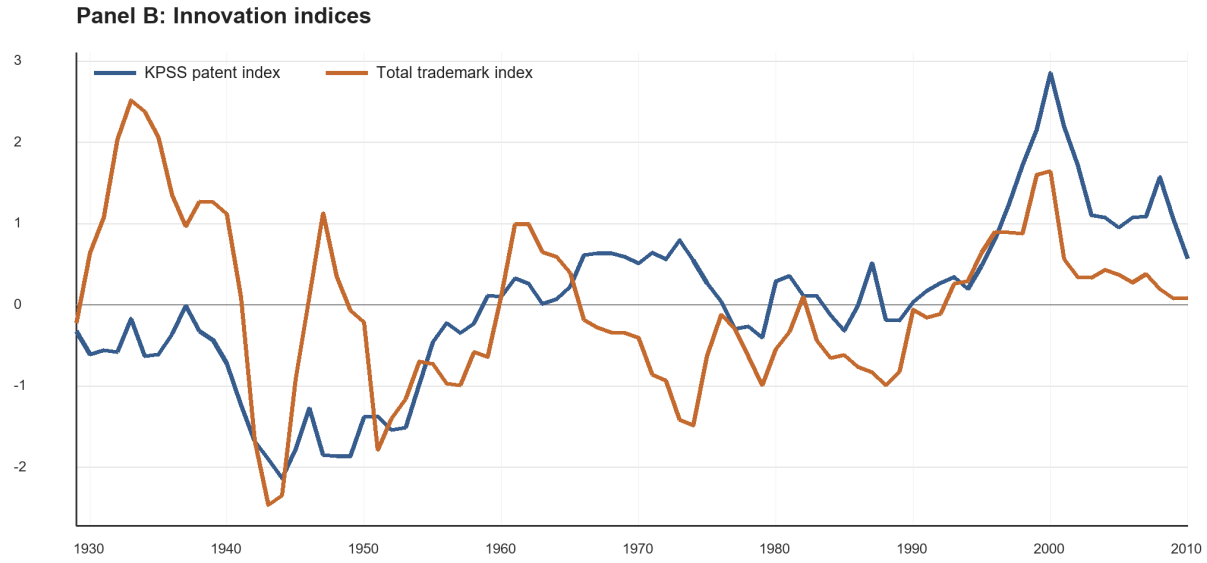


Figure 17
Trademark Use over Time

The figure presents the share of firms holding at least one trademark, by year, overall and by sector. The denominator is all firms in the NETS panel that ever reach at least three employees. The sample period is 1990 to 2020.

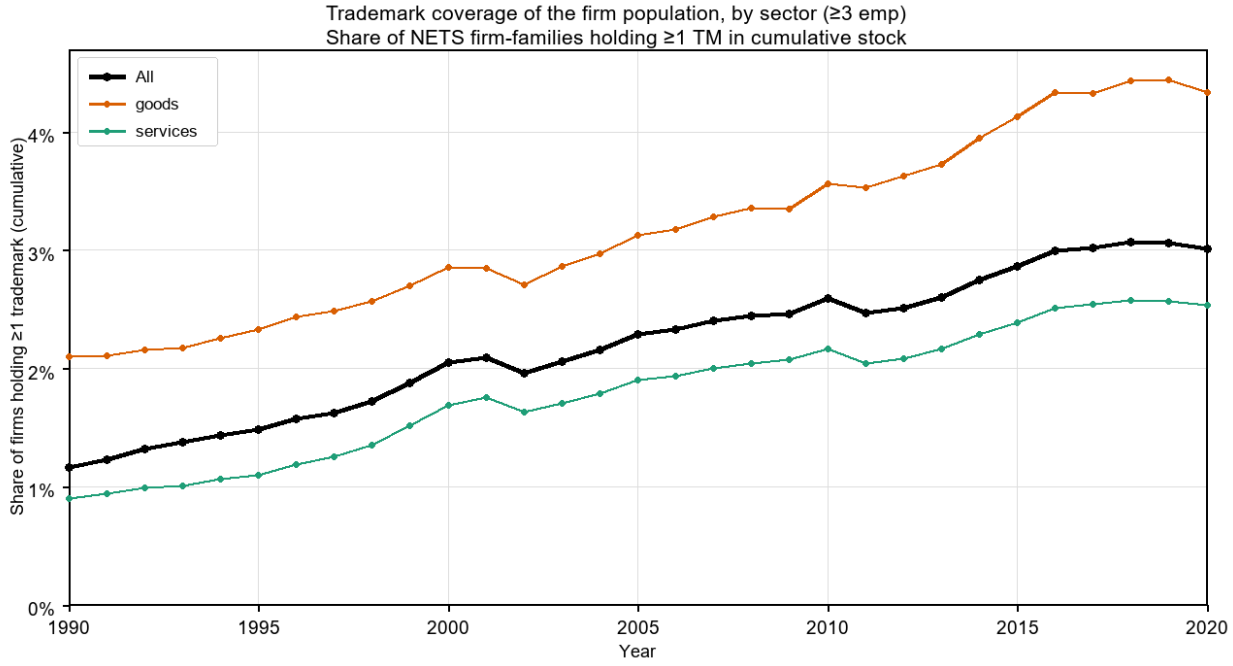


Figure 18
Trademark Use by Employment Bucket

The figure presents the share of firms holding at least one trademark by firm employment bucket and year. Buckets are defined on contemporaneous employment in NETS. The sample period is 1990 to 2020.

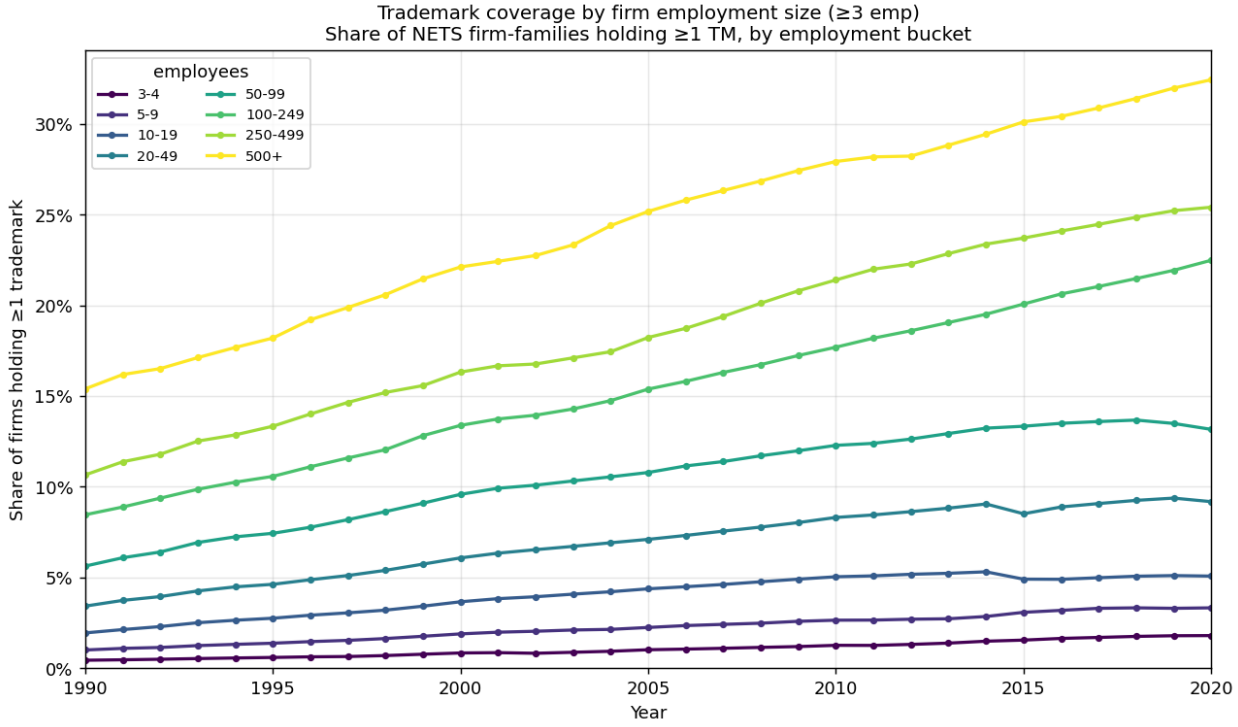


Figure 19
Firm Switching Goods and Services Examples

The figure presents the service share and its change before and after strategic events for illustrative firms. *Service share* is defined as the share of service or bundle (goods and service) trademarks to the total number of trademarks filed by firm i in year t. The pre and post periods are centered on the event year. The pre period includes all sample years up to the event year for which the company exists in the data, that is, from the IPO year of the firm, and the post period includes all sample years after the event year. The sample period is 1980 to 2019.

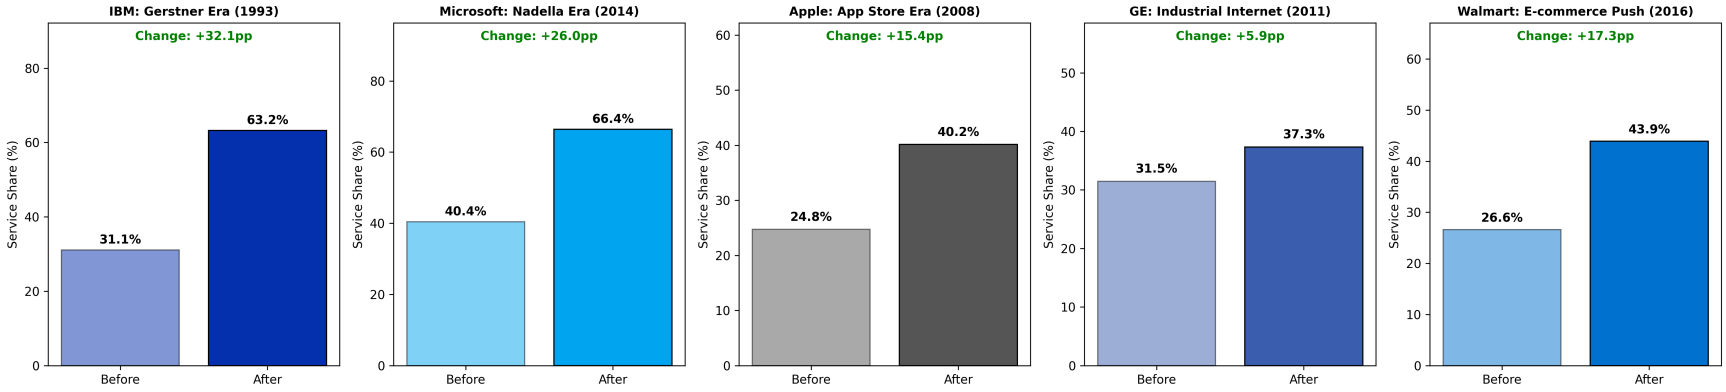


Figure 20
Illustrative Trademark Composition Time Series

The figure presents the time series of trademark filing shares for eleven illustrative examples of service-native and servitizing firms. Trademarks are categorized in three mutually exclusive groups: *Pure goods* includes all trademarks where all items are goods, *Pure services* includes all trademarks where all items are services, and *Bundle* includes all trademarks where there are both goods and services items in the trademark. All ratios are with respect to the total number of trademarks filed by firm i in year t. The red dashed line represents strategic event years at the firm level. The sample includes IBM, Microsoft, Apple, Amazon, Google, Salesforce, GE, Caterpillar, P&G, J&J, and Walmart. The sample period is 1980 to 2019.

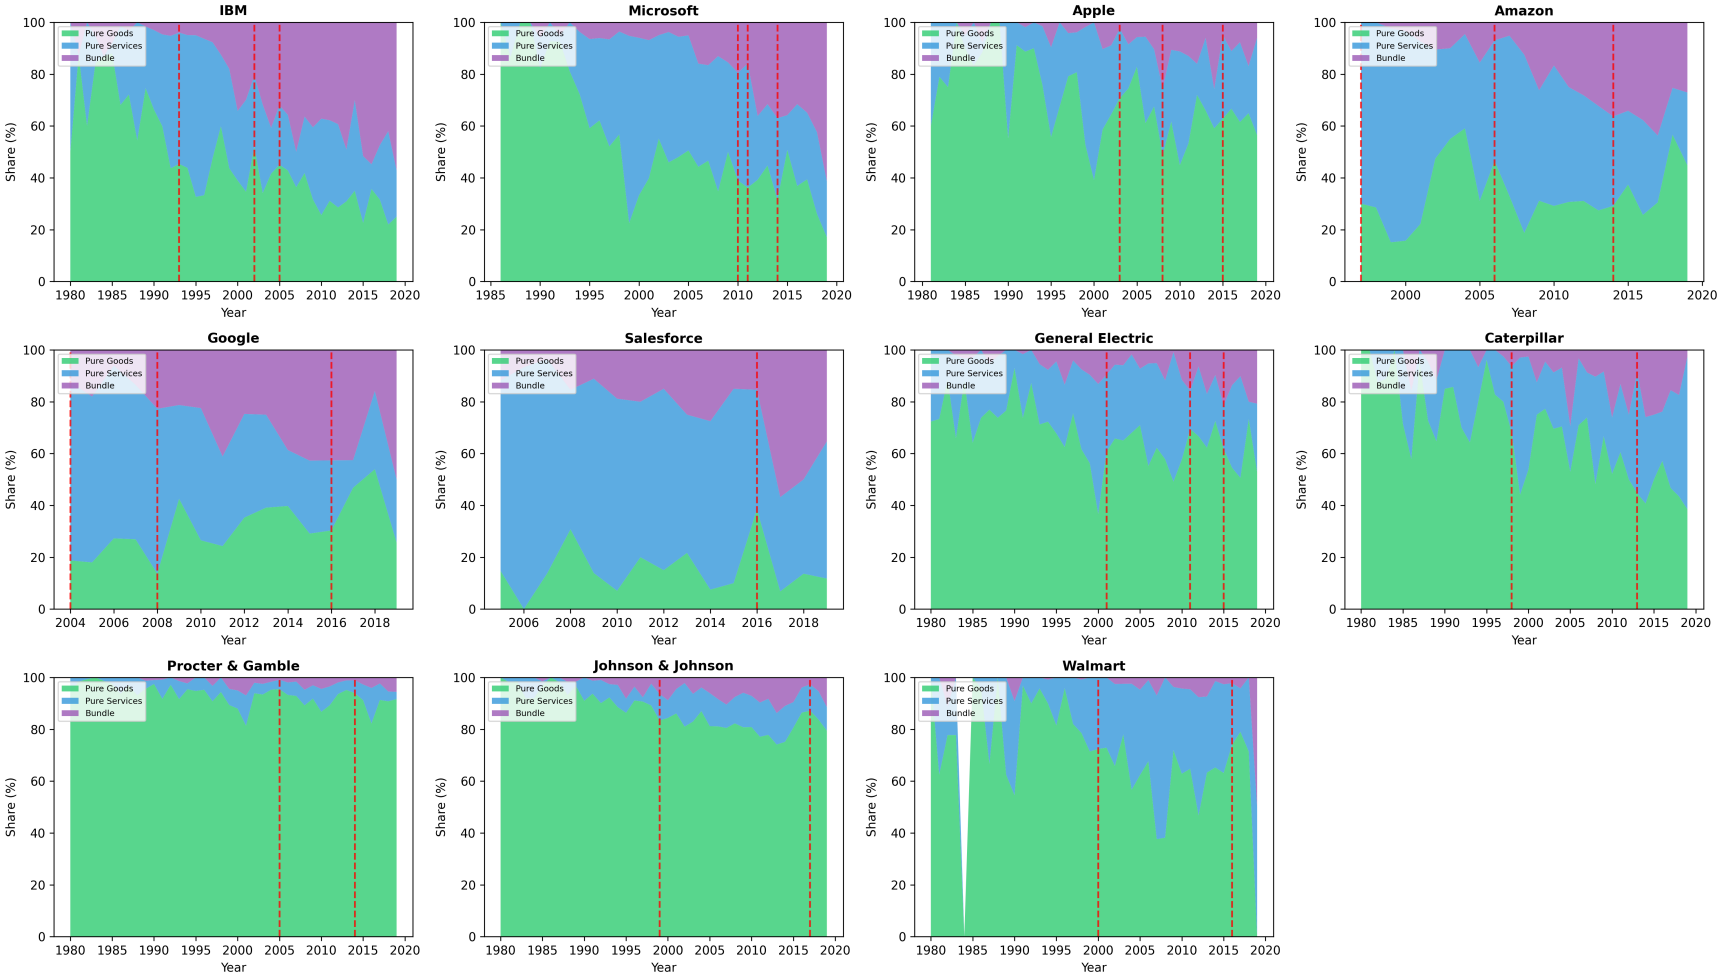
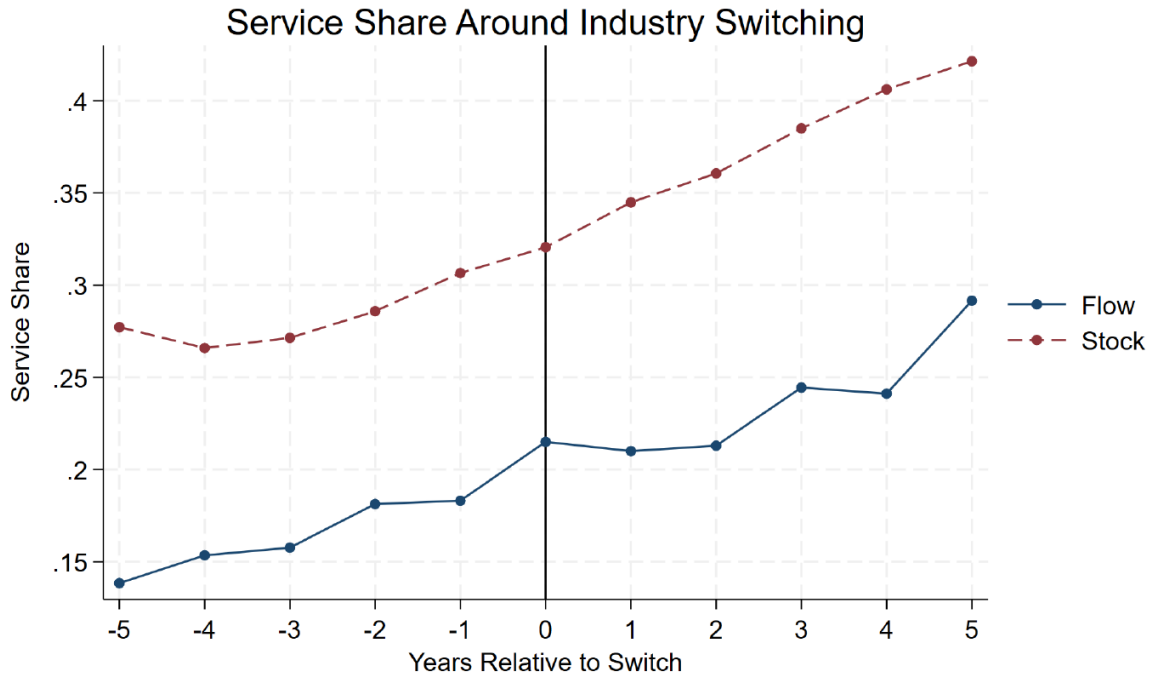


Figure 21
Event Study around Industry Switching

The figure presents an event study around the switch from goods-producing to service industry classifications, for the flow service share of trademark filings and the service share stock. Event time zero is the year of reclassification. The sample period is 1980 to 2019.

Figure 2
Event study around industry switching

The figure presents an event study around the switch from 'manufacturing' to 'service' industries and the service share stock.



Tables

Table 2
Trademark-Compustat Matched Sample

The table presents the trademark sample available and used in the analysis. Items are discrete goods or service descriptions extracted by parsing the semicolon-delimited goods/services statement within each Nice class. Panel A presents the full USPTO file; Panels B to D present the CCM-USPTO matched sample. The sample includes CRSP-Compustat firms matched to USPTO trademark data using the procedure described in Section 2. All firms in the sample must have a PERMNO, SHRCD=10,11 and total assets>0. Trademark types are defined by item composition: *Pure goods* = only items in Classes 1–34; *Pure service* = only items in Classes 35–45; *Bundle* = items in both goods and service classes.

Panel A. Full USPTO	
Unique trademarks	10,820,369
Total items	51,709,419
Goods items (Classes 1–34)	37,541,058 (72.6%)
Service items (Classes 35–45)	14,168,361 (27.4%)
Average items per trademark	4.78
Panel B. CCM Sample Coverage	
Firm-year observations	195,100
Unique Compustat firms	19,519
Sample period	1980–2019
Panel C. Matched Trademarks	
Total trademarks matched	977,439
Pure goods trademarks	638,501 (65.3%)
Pure service trademarks	287,515 (29.4%)
Bundle trademarks	51,423 (5.3%)
Panel D. Matched Items	
Total items	3,309,592
Goods items (Classes 1–34)	2,228,200 (67.3%)
Service items (Classes 35–45)	1,081,392 (32.7%)
Average items per trademark	3.39

Table 3
Examples of Item-Level Extraction from Trademark Filings

The table presents examples of how goods and services descriptions are parsed into discrete items and classified as goods (Classes 1–34) or services (Classes 35–45). Items within each trademark are separated by semicolons in the original USPTO filing.

Industry	Item description (parsed)	Class	Type
Example 1: Industrial equipment manufacturer — bundle trademark			
	Industrial machinery for manufacturing	7	Good
	Replacement parts for industrial machinery	7	Good
	Installation of industrial machinery	37	Service
	Maintenance and repair of industrial equipment	37	Service
	Technical consulting in the field of manufacturing	42	Service
Example 2: Consumer electronics company — bundle trademark			
	Smartphones	9	Good
	Tablet computers	9	Good
	Computer software for mobile devices	9	Good
	Online retail store services	35	Service
	Cloud computing services	42	Service
	Technical support services	42	Service
Example 3: Pharmaceutical company — bundle trademark			
	Pharmaceutical preparations for treating diabetes	5	Good
	Medical diagnostic test kits	5	Good
	Providing health information via website	44	Service
	Medical monitoring services	44	Service
Example 4: Business services firm — pure service trademark			
	Business management consulting	35	Service
	Human resources consulting	35	Service
	Executive recruitment services	35	Service

Table 4
Summary Statistics

The table presents summary statistics for the core servitization variables and the Compustat controls used in the validation analysis. The sample includes CRSP-Compustat firms matched to USPTO trademark data using the procedure described in Section 2, with PERMNO, SHRCD=10,11 and total assets>0. Financial variables are winsorized at the 1st and 99th percentiles. The sample period is 1980 to 2019.

Variable	Mean	Median	Std. dev.	N
Panel A. Servitization measures				
Service share flow	0.18	0.00	0.34	195,100
Service share stock	0.30	0.11	0.36	195,100
Bundle items (goods and services), flow	3.68	0.00	61.41	195,100
Pure goods items flow	9.31	0.00	116.41	195,100
Pure service items flow	3.97	0.00	23.94	195,100
Bundle trademarks flow	0.26	0.00	1.47	195,100
Pure goods trademarks flow	3.27	0.00	15.11	195,100
Pure service trademarks flow	1.47	0.00	6.40	195,100
Panel B. Financial characteristics				
Ln(Assets)	5.34	5.23	2.30	195,100
Ln(Market capitalization)	4.98	4.85	2.23	193,223
Ln(Employment)	1.01	0.56	1.13	186,627
CAPX/AT	0.06	0.03	0.07	195,100
PPENT/AT	0.25	0.17	0.24	195,100
ROA (OIBDP/AT)	0.06	0.10	0.57	178,526
Leverage	0.23	0.18	0.22	195,100
Tobin's Q	2.09	1.28	37.01	186,000
Panel C. Volatility and risk				
Asset growth	0.14	0.06	0.41	178,526
Sales volatility	0.38	0.21	0.37	179,071
Profit volatility	0.81	0.30	1.08	179,071

Table 5
Validation Case Study

The table presents the service and bundle shares of case study firms from different categories and the known events that have spurred their service introductions. *Category* describes the type of firm. *Years* provides the CCM available sample years for the firm. *Change* is the difference in the outcome variable between 1980 and 2019. Columns 1980 to 2010 present the share of services or bundles for the decade. *Key events* provides a list of years with strategic changes at the firm level. The sample includes IBM, Microsoft, Apple, Amazon, Google, Salesforce, GE, Caterpillar, P&G, J&J, and Walmart. Panel A presents *Service share*, defined as the total share of service and bundle (goods and service) trademarks to the total number of trademarks filed by firm i in year t. Panel B presents *Bundle share*, defined as the share of bundle (goods and service) trademarks to the total number of trademarks filed by firm i in year t.

SIC2	Firm	Category	Years	1980	1990	2000	2010	Change	Key events
					Panel A. Service share (%)				
73	IBM	Tech - Servitization	1980-2019	27.41	52.40	60.10	71.19	43.78	1993: Gerstner becomes CEO, "Services" strategy; 2002: Acquires PwC Consulting ($3.5B); 2005: Sells PC division to Lenovo
73	Microsoft	Tech - Servitization	1986-2019	3.33	32.66	55.16	63.92	60.59	2010: Azure cloud launched; 2011: Office 365 launched; 2014: Nadella CEO, "Cloud-first"
36	Apple	Tech - Servitization	1981-2019	13.81	26.33	37.07	39.24	25.43	2003: iTunes Store launched; 2008: App Store launched; 2015: Apple Music, services push
59	Amazon	Tech - Service Native	1997-2019		75.44	64.01	65.75	−9.69	1997: IPO as online retailer; 2006: AWS cloud launched; 2014: Echo/Alexa launched
73	Google	Tech - Service Native	2004-2019			75.48	64.91	−10.57	2004: IPO, Gmail and Maps launched; 2008: Android OS, Chrome browser; 2016: Pixel phone, hardware push Chromebook
73	Salesforce	Tech - Service Native	2005-2019			85.37	84.81	−0.56	1999: Founded as SaaS CRM; 2004: IPO; 2016: Einstein AI launched
37	General Electric	Manufacturing	1980-2019	24.61	27.88	40.83	37.82	13.20	2001: Immelt CEO, services focus; 2011: "Industrial Internet"; 2015: GE Digital formed
35	Caterpillar	Manufacturing	1980-2019	16.16	24.39	34.11	51.61	35.45	1998: Cat Financial services expand; 2013: Connected assets, telematics
28	Procter & Gamble	Consumer Goods	1980-2019	3.23	6.19	8.52	9.40	6.17	2005: Gillette acquisition; 2014: Brand divestiture
28	Johnson & Johnson	Healthcare Products	1980-2019	3.39	10.34	17.25	19.74	16.35	1999: Healthcare services expand; 2017: Actelion acquisition
53	Walmart	Retail	1980-2019	17.44	16.34	37.64	40.94	23.50	2000: Walmart.com launched; 2016: Jet.com acquisition; 2020: Walmart+ launched
					Panel B. Bundle share (%)				
	IBM	Tech - Servitization	1980-2019	4.52	7.13	35.18	38.44	33.92	
	Microsoft	Tech - Servitization	1986-2019	0.00	4.25	9.31	23.68	23.68	
	Apple	Tech - Servitization	1981-2019	1.41	2.28	8.81	14.64	13.23	
	Amazon	Tech - Service Native	1997-2019		0.78	10.23	30.54	29.76	
	Google	Tech - Service Native	2004-2019			15.67	39.81	24.14	
	Salesforce	Tech - Service Native	2005-2019			9.68	29.69	20.01	
	General Electric	Manufacturing	1980-2019	1.60	5.38	6.63	16.01	14.41	
	Caterpillar	Manufacturing	1980-2019	3.21	1.85	9.51	17.41	14.20	
	Procter & Gamble	Consumer Goods	1980-2019	0.13	1.32	3.02	3.68	3.56	
	Johnson & Johnson	Healthcare Products	1980-2019	0.48	3.58	6.38	9.24	8.76	
	Walmart	Retail	1980-2019	1.23	1.07	2.37	3.80	2.57	

Table 6
Industry Classification Switching Prediction

The table presents the regression analysis of the switch from goods-producing to service industries. The dependent variable is an indicator equal to one if the firm switches from a goods to a service industry within k years,

$$\mathbf{1}\{\text{Goods} \to \text{Service SIC}\}_{j,t+k} = \alpha + \beta \cdot \text{Filed any service TM}_{jt} + \mathbf{X}'_{jt}\gamma + \delta_i + \delta_t + \varepsilon_{jt}.$$

Filed any service TM is an indicator for filing at least one trademark covering service classes in year $t-1$. The narrow definition classifies SICH1 codes 6–8 as services; broad adds Transport and Utilities (SICH1=4); paper adds Trade (SICH1=5). Temporary switches that reverse within two years are excluded. All specifications include year and industry fixed effects. Standard errors clustered at the firm level are presented in parentheses. $^{*}(p < 0.10)$, $^{**}(p < 0.05)$, $^{***}(p < 0.01)$. The sample period is 1980 to 2019.

Table 7
Industry Classification Switching Prediction

The table presents the regression analysis of the switch from manufacturing to service industries. Dependent variable is an indicator equal to one if the firm switches from a goods to a service industry within k years.

$$\mathbf{1}\{\text{Goods} \to \text{Service SIC}\}_{j,t+k} = \alpha + \beta \cdot \text{Filed any service TM}_{jt} + \mathbf{X}'_{jt}\gamma + \delta_i + \delta_t + \varepsilon_{jt}, \tag{8}$$

"Filed any service TM" is an indicator for filing at least one trademark covering service classes in year $t-1$. Narrow definition classifies SICH1 codes 6–8 as services; Broad adds Transport and Utilities (SICH1=4); Paper adds Trade (SICH1=5). Temporary switches that reverse within two years are excluded. Standard errors clustered at the firm level are presented in parentheses. All specifications include year and industry fixed effects. $^{*}(p < 0.10)$, $^{**}(p < 0.05)$, $^{***}(p < 0.01)$.

	Narrow			Broad			Paper		
	(1) k=1	(2) k=3	(3) k=5	(4) k=1	(5) k=3	(6) k=5	(7) k=1	(8) k=3	(9) k=5
Filed any service TM (t-1)	0.001**	0.000	0.001	0.001*	0.001*	0.001	0.002***	0.000	0.001**
	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Ln(Assets)	-0.001**	-0.000*	-0.000	-0.000	-0.000	-0.000	-0.001**	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Ln(Employment)	-0.000	-0.000	-0.000	-0.001*	-0.001*	-0.001	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
ROA	-0.003***	-0.001*	-0.002*	-0.003***	-0.002**	-0.004***	-0.002**	-0.002**	-0.003**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.006***	0.006***	0.005***	0.007***	0.006***	0.006***	0.008***	0.006***	0.006***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	108,100	88,811	73,580	94,494	77,470	64,012	91,489	75,355	62,576

Table 7
Servitization by Industry

The table presents the service and goods shares by one-digit SIC industry, sorted by service share. The sample includes CRSP-Compustat firms matched to USPTO trademark data using the procedure described in Section 2. All firms in the sample must have a PERMNO, SHRCD=10,11 and total assets>0. *Pure goods* includes all trademarks where all items are goods, *Pure services* includes all trademarks where all items are services, and *Bundle* includes all trademarks where there are both goods and services items in the trademark. All ratios are with respect to the total number of trademarks in the industry over the sample 1980 to 2019.

Panel A. Full sample industry distribution

					Service (%)			Share (%)		
SIC1	Industry	Obs	Firms	No. TMs	Firms	Min	Max	P. Goods	P. Services	Bundle
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
6	Finance	33,562	3,396	94,223	96%	0.4%	100%	28%	66%	6%
8	Services (Prof.)	7,810	857	18,806	95%	0.3%	100%	33%	58%	8%
4	Transport/Utilities	16,282	1,401	88,683	96%	0.6%	100%	40%	55%	5%
7	Services (Bus.)	23,593	2,873	123,924	88%	0.2%	100%	43%	47%	10%
1	Mining/Construct.	11,433	1,232	14,921	89%	0.3%	100%	49%	45%	6%
5	Trade	19,391	1,962	96,525	87%	0.3%	100%	60%	35%	6%
9	Public Admin.	2,200	389	13,353	78%	1.0%	100%	73%	23%	5%
0	Agriculture	598	69	2,570	64%	1.7%	100%	76%	21%	3%
3	Manuf. (Heavy)	50,641	4,529	275,753	63%	0.1%	100%	83%	13%	5%
2	Manuf. (Light)	29,590	2,811	248,681	66%	0.2%	100%	85%	11%	3%

		Service (%)			Within Industry Share (%)			Total Share (%)			
SIC1	Industry	Firms	Min	Max	P. Goods	P. Services	Bundle	P. Goods	P. Services	Bundle	Industry
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	

Panel B. Industry distribution 1980

6	Finance	82%	6.7%	100%	35%	64%	1%	2%	3%	0%	5%
8	Services (Prof.)	100%	16.7%	100%	50%	48%	2%	0%	0%	0%	0%
4	Transport/Util.	75%	2.9%	100%	61%	37%	2%	3%	2%	0%	5%
7	Services (Bus.)	62%	16.7%	100%	45%	50%	4%	2%	2%	0%	4%
1	Mining/Construct.	64%	22.2%	100%	62%	37%	1%	1%	1%	0%	2%
5	Trade	59%	1.4%	100%	70%	27%	3%	5%	2%	0%	7%
9	Public Admin.	50%	8.1%	100%	84%	12%	4%	1%	0%	0%	1%
0	Agriculture	0%	0.0%	0%	100%	0%	0%	0%	0%	0%	0%
3	Manuf. (Heavy)	21%	1.1%	100%	92%	6%	1%	34%	2%	0%	36%
2	Manuf. (Light)	26%	0.7%	100%	94%	5%	0%	37%	2%	0%	39%

Panel C. Industry distribution 2019

6	Finance	92%	1.1%	100%	29%	57%	14%	4%	7%	2%	12%
8	Services (Prof.)	90%	1.3%	100%	29%	56%	15%	1%	1%	0%	2%
4	Transport/Util.	92%	3.0%	100%	44%	47%	9%	5%	5%	1%	11%
7	Services (Bus.)	92%	7.0%	100%	25%	55%	20%	3%	7%	3%	13%
1	Mining/Construct.	70%	1.8%	100%	49%	39%	12%	1%	1%	0%	1%
5	Trade	83%	0.3%	100%	56%	33%	11%	6%	4%	1%	11%
9	Public Admin.	100%	20.0%	53%	61%	30%	8%	1%	0%	0%	1%
0	Agriculture	75%	16.7%	100%	44%	44%	11%	0%	0%	0%	0%
3	Manuf. (Heavy)	66%	0.3%	100%	71%	18%	12%	17%	4%	3%	24%
2	Manuf. (Light)	63%	0.4%	100%	77%	13%	10%	19%	3%	2%	25%

References

Baumol, W. J. 1967. Macroeconomics of unbalanced growth: The anatomy of urban crisis. *American Economic Review* 57:415–26.

Bernard, A. B., S. J. Redding, and P. K. Schott. 2010. Multiple-product firms and product switching. *American Economic Review* 100:70–97.

Buera, F. J., and J. P. Kaboski. 2012. The rise of the service economy. *American Economic Review* 102:2540–69. doi:10.1257/aer.102.6.2540.

Crozet, M., and E. Milet. 2017. Should everybody be in services? The effect of servitization on manufacturing firm performance. *Journal of Economics & Management Strategy* 26:820–41.

Cusumano, M. A., S. J. Kahl, and F. F. Suarez. 2015. Services, industry evolution, and the competitive strategies of product firms. *Strategic Management Journal* 36:559–75.

Dinlersoz, E. M., N. Goldschlag, A. Myers, and N. Zolas. 2018. An anatomy of U.S. firms seeking trademark registration. CES Working Paper 18-22.

Farboodi, M., and L. Veldkamp. 2023. Data and markets. *Annual Review of Economics* 15:23–40.

Gerstner, L. V. 2002. *Who says elephants can't dance? inside IBM's historic turnaround*. HarperBusiness.

Graham, S. J. H., G. Hancock, A. C. Marco, and A. F. Myers. 2013. The USPTO trademark case files dataset: Descriptions, lessons, and insights. *Journal of Economics & Management Strategy* 22:669–705.

Herrendorf, B., R. Rogerson, and Á. Valentinyi. 2014. Growth and structural transformation. In P. Aghion and S. N. Durlauf, eds., *Handbook of Economic Growth*, vol. 2, chap. 6, 855–941. Elsevier.

Hoberg, G., and G. M. Phillips. 2025. Scope, scale, and concentration: The 21st century firm. *The Journal of Finance* 80:415–66.

Jones, C. I., and C. Tonetti. 2020. Nonrivalry and the economics of data. *American Economic Review* 110:2819–58–.

Jordà, Ò. 2005. Estimation and inference of impulse responses by local projections. *American Economic Review* 95:161–82.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman. 2017. Technological innovation, resource allocation, and growth*. *The Quarterly Journal of Economics* 132:665–712.

Leo Liu, E. S., and W. W. Tham. 2026. Commercialization margins and uncertainty. Working Paper.

Mayer, T., M. J. Melitz, and G. I. P. Ottaviano. 2014. Market size, competition, and the product mix of exporters. *American Economic Review* 104:495–536.

Ngai, L. R., and C. A. Pissarides. 2007. Structural change in a multisector model of growth. *American Economic Review* 97:429–43.

Nordhaus, W. D. 2008. Baumol's diseases: A macroeconomic perspective. *The B.E. Journal of Macroeconomics* 8:1–39.

Wise, R., and P. Baumgartner. 1999. Go downstream: The new profit imperative in manufacturing. *Harvard Business Review* 77:133–41.

Online Appendix

(For Online Publication)

Table A1
Nice Classification System: Goods and Service Classes

The table present the main classes in the Nice Classification. Nice Classification is an international system for categorizing goods and services for trademark registration. Classes 1–34 cover goods; Classes 35–45 cover services. A single trademark may be registered in multiple classes.

Class	Category	Examples
Panel A: Goods Classes (1–34)		
1	Chemicals	Industrial chemicals, adhesives, fertilizers
5	Pharmaceuticals	Medicines, medical preparations, dietary supplements
7	Machinery	Machines, motors, engines, machine tools
9	Electronics	Computers, software, scientific instruments, smartphones
11	Environmental Control	Heating, cooling, lighting, water purification
12	Vehicles	Automobiles, aircraft, boats, bicycles
25	Clothing	Apparel, footwear, headgear
28	Games/Toys	Toys, sporting goods, video game apparatus
Panel B: Service Classes (35–45)		
35	Business Services	Advertising, retail services, business management consulting
36	Financial Services	Banking, insurance, real estate, investment services
37	Construction/Repair	Building construction, installation and repair services
38	Telecommunications	Broadcasting, data transmission, telephone services
39	Transportation	Shipping, logistics, travel arrangement, storage
40	Material Treatment	Custom manufacturing, food processing, printing
41	Education/Entertainment	Training, publishing, entertainment production
42	Technology Services	Software development, IT consulting, R&D services
43	Food/Accommodation	Restaurants, hotels, catering services
44	Medical Services	Healthcare, veterinary services, therapy, pharmacies
45	Legal/Security	Legal services, security consulting, licensing

Table A2
United States Classification System: Goods and Service Classes (Prior U.S. Schedule)

The table presents selected classes in the *prior United States* trademark classification schedule. Under the USPTO rules, this prior U.S. schedule applies (for statutory purposes) to registrations resulting from applications filed on or before August 31, 1973, unless amended to adopt international (Nice) classification. Goods classes are numbered 1–52; service classes are numbered 100–107. A single trademark may be registered in multiple classes.

Class	Category	Examples
Panel A: Goods Classes (selected from 1–52)		
1	Raw materials	Raw or partly prepared materials (e.g., basic industrial inputs)
6	Chemicals	Chemicals and chemical compositions
12	Construction materials	Building and construction materials
18	Pharmaceuticals	Medicines and pharmaceutical preparations
19	Vehicles	Vehicles and related goods
21	Electrical apparatus	Electrical apparatus, machines, and supplies
22	Games/Toys	Games, toys, and sporting goods
23	Machinery/Tools	Cutlery, machinery, and tools (and parts thereof)
26	Scientific instruments	Measuring and scientific appliances
34	Environmental control	Heating, lighting, and ventilating apparatus
39	Clothing	Clothing (e.g., apparel)
46	Foods	Foods and ingredients of foods
51	Cosmetics	Cosmetics and toilet preparations
Panel B: Service Classes (100–107)		
100	Miscellaneous	Miscellaneous services not otherwise classified
101	Business services	Advertising and business (e.g., advertising, business management)
102	Financial services	Insurance and financial (e.g., banking, insurance)
103	Construction/repair	Construction and repair services
104	Communications	Communication services (e.g., broadcasting, telecom)
105	Transportation/storage	Transportation and storage (e.g., shipping, logistics, warehousing)
106	Material treatment	Material treatment (e.g., printing, processing, custom manufacturing)
107	Education/entertainment	Education and entertainment services