A Godsend Bug:

Demand-Driven Upgrading in India's IT Sector*

BooKang Seol†

August 3, 2026

Click here for the latest version

Abstract

For firms in developing countries, innovation often takes the form of upgrading. Exporting is one prominent route, but foreign orders bundle buyer knowledge, quality demands, relationships, and scale. Because these arrive together and access persists, exporter growth does not reveal which of them changes the firm. This paper asks whether a large but temporary export opportunity can bring organizational upgrading forward and leave firms more capable after it disappears. The Year 2000 (Y2K) problem created one for Indian software firms: a surge of foreign demand to repair legacy software before an immovable deadline. In a newly digitized industry panel spanning 1992–2003, firms whose pre-existing technology stacks matched this work emerged persistently larger after the deadline, with an employment differential of 0.425 log points (roughly 53%). They exhibited a further, durable widening in software-process and quality-management certification, while reported client counts and foreign-market reach show no comparably robust persistent differentials. A simple model formalizes this scale–capability complementarity, where capability allows firms to operate reliably at scale, and the prospect of greater scale, albeit temporary, can make it worthwhile to build capability sooner.

*I am grateful to Jonas Hjort, Eric Verhoogen, Miguel Urquiola, Diana Weinhold, Joana Naritomi, Sandra Sequeira, Apoorv Gupta, Suresh Naidu, and Xavier Jaravel for their feedback. Seminar participants at Columbia, Northwestern Annual Conference, SNU provided helpful comments. Deepak Mishra, director at the ICRIER, provided generous support and opportunities to connect with industry experts in India. Suma Athreye at IIT Delhi and Aarti Reddy at ICRIER helped collect the data. Eva Nappert provided excellent research assistance. LSE provided research funding.

†Department of International Development, London School of Economics. b.seol@lse.ac.uk

1 Introduction

For firms in developing countries, innovation often takes the form of upgrading: adopting technologies, organizational practices, and quality standards already used at the frontier (Gerschenkron 1962; Verhoogen 2023). Exporting is one prominent route to such upgrading, but identifying which features of exporting change the firm remains difficult. More productive firms select into foreign markets (Melitz 2003), and foreign orders bundle knowledge from buyers, demand for higher quality, relationships that support future sales, and the scale over which fixed investments can be recovered (Verhoogen 2008; Atkin et al. 2017a; Lileeva and Trefler 2010; Bustos 2011). Because these forces often arrive together and foreign-market access tends to persist, subsequent exporter growth does not reveal which part of the bundle changed the firm or whether continuing access itself sustains later activity. The role of scale is especially difficult to assess because a permanent market expansion may both induce a fixed investment and continue to justify it. This paper asks whether a large but temporary export opportunity can instead bring organizational upgrading forward and leave firms more capable after the opportunity disappears. In the late 1990s, the Year 2000 (Y2K) computer problem created such an opportunity for Indian software firms through a surge of foreign demand for repairing legacy software before a known, immovable deadline on January 1, 2000. I find that firms whose pre-existing technology stacks matched this work emerged persistently larger after the deadline and exhibited a further, durable widening in software-process and quality-management certification, while reported client counts and foreign-market reach show no comparably robust persistent differentials. A simple model formalizes this scale–capability complementarity: capability allows firms to operate reliably at scale, but the prospect of greater scale can also make it worthwhile to build capability sooner.

The Year 2000 computer problem, or Y2K, originated in a simple memory-saving convention. When computer memory was costly, programmers often encoded calendar years using two digits rather than four. Programs embedded in payroll, banking, insurance, billing, and other essential systems could therefore interpret the year 2000 as 1900. Because this convention was embedded throughout decades of program logic, no common patch could eliminate the problem. Remediation instead required programmers to locate affected routines, rewrite them in the original languages and computing environments, and test the surrounding systems. The task created concentrated demand for engineers already fluent in COBOL, IBM mainframes, and related legacy technologies. US public companies alone disclosed approximately $52 billion in remediation spending in the filings collected for this

paper (U.S. Securities and Exchange Commission 1998; Loughran and McDonald 2011). Commercial remediation activity accelerated sharply in 1998, peaked in 1999, and collapsed after January 1, 2000, a deadline no buyer, supplier, or government could move. For a brief period, expertise in technologies that new software development had largely left behind became unusually valuable because essential installed systems still depended on them.

Four features make the episode unusually revealing. First, demand was large and market-generated, arising on commercial terms rather than through an experiment or targeted program. Its scale made investments with substantial fixed organizational costs a plausible response. Second, exposure depended on pre-existing technical specialization. The COBOL and mainframe expertise that positioned firms for remediation had been accumulated through earlier business, before the late surge in Y2K spending. Measuring that specialization in 1995, and more stringently in 1992, limits the concern that firms acquired it to select into the later opportunity. Third, the opportunity rewarded proficiency in legacy rather than frontier technologies. Although the work could still transmit buyer knowledge, standards, or organizational practices, its technical content did not itself confer a continuing frontier advantage once remediation ended. Finally, direct remediation demand expired on a known date. This expiry creates an expiration test unavailable when market access expands permanently. Any Y2K-induced response that persists after the deadline must be carried by a state the opportunity left behind rather than by continuing remediation demand itself.

Persistent activity alone cannot identify the durable state that carries it. A temporary opportunity may push a firm across a market-participation threshold, leaving behind a buyer relationship, approved-vendor status, reputation, or foreign presence that supports subsequent sales (Baldwin and Krugman 1989; Roberts and Tybout 1997; Das et al. 2007). Alternatively, it may induce investment in an organizational capability that remains useful in ordinary business. These mechanisms need not be mutually exclusive. Either can leave employment higher after the original orders end. Participation sustains demand for the firm's services, while organizational capability expands the scale and complexity of work the firm can execute reliably. I therefore treat persistent scale as the phenomenon to be explained and compare organizational and participation outcomes to assess which durable margins changed alongside it.

The model formalizes organizational upgrading as an investment-timing problem. As a firm takes on more projects, scarce organizational attention is spread more thinly, raising the costs of coordination, testing, and rework. A documented process and quality-management system expands the scale the organization can operate reliably, so the return to installing

it rises with project demand. Installing it is a sunk investment, however, and a firm can postpone the cost. Without Y2K, a firm may prefer to upgrade later—or not within the relevant horizon—rather than incur that cost immediately. Because ordinary and remediation projects use the same delivery organization, Y2K adds a finite premium to having the system in place during the remediation window. The shock induces early upgrading only when that premium makes adoption now preferable to the firm's best no-Y2K alternative. The premium need not cover the sunk cost on its own because the system also earns returns in ordinary business.

The model's implications concern who upgrades and what persists. The upgrading response is concentrated near the no-Y2K adoption margin: firms that would upgrade early anyway are unaffected, firms far from adoption continue to wait, and intermediate firms bring the investment forward. Because the upgraded system continues to serve ordinary projects, an obsolete technical match can predict modern organizational capability and the scale it supports after the match itself has lost direct value—an obsolete-match reversal. Comparison firms may later reach their own upgrading dates, so the differential can narrow. Persistence through the sample horizon is therefore an empirical result, not an implication of the setup.

Measuring both the opportunity and the durable states it may have left behind requires observing firms' technologies, scale, organization, and market participation in the same panel. I digitize and link the seven surviving editions of the *NASSCOM Directory of Indian Software & Service Companies*, spanning 1992 to 2003, into a panel of 1,248 firms. Each profile reports technology stacks, employment and registered capital, revenue and exports, certifications, services, named clients, and foreign-market reach. The main design follows the 235 incumbent firms listed in 1995. Firm histories are reconstructed across name changes, acquisitions, and re-entry using directory evidence, news archives, and Ministry of Corporate Affairs records. This process records 150 within-panel M&A events, which enter as time-varying controls, and helps distinguish directory nonappearance from legal exit. An independently assembled corpus of 110,082 US regulatory filings measures the magnitude, timing, and client-sector composition of disclosed Y2K spending from the buyer side rather than from Indian firms' own reports. For a selected subset of firms, annual audited Prowess statements provide longer histories for financial outcomes, while Software Technology Parks of India reports date the rollout of export infrastructure.

I measure exposure using the technology stacks firms reported in 1995, before remediation spending accelerated. Of the 235 incumbents, 135 reported at least one qualifying

COBOL- or mainframe-family technology; a stricter measure distinguishes an isolated legacy skill from a coherent production stack. The baseline match predicts reported Y2K engagement: matched firms were 32.7 percentage points more likely to report offering Y2K services in 1998 or 1999, against a sample mean of 47 percent. A similarly constructed index of modern technologies has no additional predictive power once the legacy match is included. Event studies trace the matched–comparison differential wave by wave relative to 1995, while split-window difference-in-differences summarize relative changes during the 1998–99 remediation window and the 2001–03 post-deadline period against the pooled 1992–95 pre-period. Firm fixed effects absorb permanent differences across firms, and entry-cohort-by-year effects allow firms first observed in different directory waves to follow different common paths.

The legacy match was not randomly assigned, so identification rests on relative changes and a conditional parallel-trends assumption rather than equality in baseline levels. Two related limitations matter. First, pre-shock time variation in the directories is limited. The single wave before 1995 can reveal an earlier differential but cannot establish a long pre-trend. Because Y2K was already known, I reconstruct exposure solely from the 1992 directory as the strictest check that the measured technical match preceded the late commercial surge. Annual Prowess accounts extend the pre-period for the small, selected subset of firms that can be linked to legacy-based exposure. The full Prowess universe places IT firms against other industries, but differential sectoral pre-trends make that exercise corroborative rather than a second identification design. Second, matched firms were already older, larger, and more foreign-oriented. Within-firm comparisons remove these permanent level differences but not differential changes associated with them. I therefore allow predetermined measures of firm quality, prior market position, and communications capacity to predict separate remediation and post-deadline changes, and account separately for exposure to the broader US IT boom. Balanced-panel, reweighting, and legal-survival checks assess whether later directory nonappearance drives the results. The maintained identifying assumption remains that, absent Y2K, matched and comparison firms would have experienced similar changes around the remediation window and deadline. Under that assumption, the estimates identify the differential consequence of carrying the predetermined legacy match when the demand surge arrived among observed 1995 incumbents.

The evidence delivers two main findings. First, legacy-matched firms exhibit a persistent expansion in scale, with the clearest evidence coming from employment. During the 1998–99 remediation period, their employment differential relative to comparison firms is 0.172 log points and statistically indistinguishable from zero. After the deadline, it rises to 0.425 log

points—approximately 53%, or roughly 50 employees at the 1995 comparison-group mean—and the post-minus-remediation contrast is statistically significant. Registered capital rises earlier, during remediation, although its magnitude is more sensitive to functional form. The post-deadline estimates for software exports and revenue are positive but too imprecise to provide a precise output counterpart to the employment result. The employment differential remains positive across five exposure constructions and under PPML estimation in levels. Matched audited accounts show employee compensation and fixed assets moving in the same direction, while corporate-registry evidence suggests higher legal survival. Because the audited sample is selected and the survival estimate is coding-sensitive, I treat this evidence as corroborative. Taken together, legacy-matched firms emerged from the Y2K episode operating at a persistently larger scale, most clearly in employment.

Second, legacy-matched firms exhibit a further, durable widening in reported process and quality-management certification. On a 0–4 ladder running from no certification through an ISO-family credential to progressively higher Capability Maturity Model status, legacy-matched firms rise 0.319 points more than comparison firms during remediation and 1.066 points more after the deadline. The corresponding post-deadline differentials are 32.6 percentage points for reported CMM status, the outcome most closely connected to software-process maturity, and 24.2 percentage points for an ISO-family credential. An important qualification is that the aggregate ladder already exhibits differential movement between 1992 and 1995, although the pre-period divergence is substantially smaller for CMM. Section 8 examines how sensitive the later ladder differential is to alternative continuations of that earlier movement. Two additional patterns nevertheless sharpen the organizational-upgrading interpretation. Among the 206 firms reporting no certification in 1995, legacy-matched firms subsequently climb 0.783 more ladder points than comparison firms, so the result is not confined to firms already certified. The response is also concentrated among firms with coherent legacy production stacks, which gain an additional 0.685 post-deadline ladder points relative to firms with incomplete stacks. These patterns overall make organizational upgrading the most natural interpretation of the durable post-deadline certification gap.

By comparison, participation-related outcomes measured in the same directories move less clearly. Reported client counts and foreign-country reach show no comparably robust persistent differentials, although the estimates are imprecise and their confidence intervals permit economically meaningful positive responses. On a standardized scale, the post-deadline certification differential exceeds the foreign-country differential by 0.82 standard deviations, a contrast that remains statistically distinguishable from zero after accounting

for initial firm size; the certification–client contrast is less robust under the same comparison. Other participation measures also change little. The share of classified clients that is foreign, already approximately 0.80 among legacy-matched firms before Y2K, does not rise, and reported service breadth changes little. A small, survivor-conditional transition nevertheless suggests that participation may have complemented upgrading: among firms observed in both 1998 and 2001 and reporting no foreign presence in 1998, 60% of legacy-matched firms reported one by 2001, compared with 33% of comparison firms. These comparisons suggest that the durable change beyond scale occurred on an organizational-capability and quality-assurance margin, with broader market participation potentially playing a complementary role.

Conditioning on observable pre-existing advantages attenuates the certification estimate but does not eliminate either main pattern. The most demanding specification interacts predetermined firm quality, prior foreign engagement, and communications capacity with the remediation and post-deadline periods, so that firms with these observable strengths may change differentially in each window. It therefore asks whether the legacy-match differential remains after accounting for period-specific changes associated with observable pre-shock strength. The certification estimate falls from 1.066 to 0.571 ladder points, while the employment estimate remains 0.476 log points. Certification gains are substantial among firms with both limited and broad prior foreign reach, and the employment estimates are positive in every tercile of initial size. Certification also remains positive when exposure is reconstructed from technologies reported in 1992. Within the main 1995 design, the findings remain stable under alternative outcome and exposure definitions, panel constructions, attrition adjustments, and difference-in-differences estimators. Additional tests address two specific competing explanations. The certification differential remains after accounting for firms' exposure to broader US IT spending, and its timing follows the January 2000 Y2K deadline rather than the rise and subsequent collapse of the dot-com boom. Domestic certification policies also do not readily explain the main response: one program targeted small firms, whereas legacy-matched firms were generally larger, and the other did not reward the higher CMM levels that contribute to the ladder widening. Overall, the employment and certification patterns are robust to the observable controls and alternative specifications, and the timing and policy evidence weighs against broader IT demand and domestic certification incentives as complete explanations.

The paper makes three contributions. First, it advances work on firm upgrading and learning by exporting by highlighting how a temporary increase in foreign demand can bring

organizational upgrading forward when it raises the scale over which fixed organizational costs can be recovered.[1] Existing evidence shows that firms can be held back by missing knowledge, contracting frictions, financing constraints, or lack of access to buyers who reward higher quality. Randomly assigned export orders among Egyptian rug producers improved productivity and quality through buyer-conveyed knowledge, with no evidence that monetary or time investments explain those quality gains (Atkin et al. 2017a). Ilzetzki (2024) shows that surges in government demand led capacity-constrained aircraft plants to improve their production methods. Permanent expansions in foreign-market access can make fixed productivity and technology investments profitable (Lileeva and Trefler 2010; Bustos 2011), but because the larger market remains available, the investment is never observed apart from the continuing access that justifies it. This paper studies the case those designs leave open: can a finite, market-generated export opportunity bring a sunk organizational investment forward and leave an organizational differential after the inducing market disappears? Y2K provides such a setting. The opportunity was temporary, its incidence depended on firms' pre-existing technical specialization, and the durable response appears on a certification margin consistent with organizational investment. Firms may already know about a better organizational system and still postpone adopting it because its fixed overhead does not pay at their ordinary scale, and a finite commercial opportunity can change that calculation. Foreign demand may therefore transform firms not only by conveying new knowledge, but also by changing when it becomes worthwhile to formalize their organization. The episode does not isolate project scale from buyer standards, learning, financing, or relationships, but it shows why the scale bundled inside exporting deserves separate attention.

Second, the paper contributes to work on temporary shocks and hysteresis by asking which firm state changes when the shock ends.[2] Persistence is not the default outcome of temporary demand. Procurement evidence illustrates both possibilities: among Ecuadorian construction firms, the sales and input responses to randomly allocated contracts dissipate within roughly 14 months (Carrillo et al. 2023), whereas Brazilian procurement winners re-

[1] Verhoogen (2023) surveys firm-level upgrading in developing countries. On buyer-conveyed knowledge, quality upgrading, and supplier relationships, see Verhoogen (2008), Atkin et al. (2017a), and Alfaro-Ureña et al. (2022); on market access and fixed investment, see Lileeva and Trefler (2010) and Bustos (2011); on subsidized market access, see Ali et al. (2025); on demand-induced product innovation, see Jaravel (2019); on organizational and managerial adoption barriers, see Atkin et al. (2017b), Bloom et al. (2013), and Bloom et al. (2020); and on the option value of waiting to undertake an irreversible investment, see McDonald and Siegel (1986).

[2] On trade hysteresis through sunk export-entry and re-entry costs, see Baldwin and Krugman (1989), Roberts and Tybout (1997), and Das et al. (2007). On durable industrial effects of temporary procurement, protection, and other dated shocks, see Ferraz et al. (2015), Juhász (2018), Hanlon (2017), Hanlon (2020), Lane (2025), and Mitrunen (2025).

main larger beyond the contracts partly because they win subsequent auctions and enter other markets (Ferraz et al. 2015). Where persistence arises, trade-hysteresis models commonly represent its carrier as a participation state—a relationship, reputation, or foothold in the market—that is inferred from persistent activity. The Y2K panel instead records observable counterparts to both classes of state: certification for organizational capability or assurance, and client counts, foreign-market reach, and foreign presence for participation. Y2K's known expiry makes the comparison informative because continuing remediation demand cannot directly explain firms' later scale. The capability-related margin moves more robustly than the observed participation margins. This comparison does not establish that certification mediated employment or rule out participation as a complementary mechanism. It instead places capability accumulation alongside participation assets as an empirically visible margin associated with demand-shock persistence.

Third, the paper provides firm-level evidence on capability formation in a tradable service and a microeconomic account of a widely cited turning point in India's software industry.[3] Earlier work documents the importance of reputation and certification in software contracting and traces the industry's movement toward more sophisticated service provision. Aggregate software exports, however, rose smoothly through Y2K and cannot reveal which firms changed or what changed inside them. The firm panel locates the differential response among incumbents holding the inherited legacy match and connects their pre-shock technical position to an observable organizational margin after the deadline. Y2K did not create India's software industry, and the design does not identify its contribution to the national export takeoff. Its narrower contribution is evidence that the episode accelerated organizational maturation among firms positioned to exploit it. The resulting obsolete-match reversal gives that history microeconomic content. Proficiency in technologies the frontier had abandoned created access to demand from the installed base on which the frontier still depended, and the organizational capability accumulated around that demand outlasted it. A comparative advantage in a modern service industry was built, in part, on a legacy capability whose commercial premium was scheduled to expire.

The paper proceeds as follows. Section 2 describes the data. Section 3 presents the background and stylized facts. Section 4 develops the model. Section 5 sets out the research design. Section 6 reports the scale and upgrading results. Section 7 examines observable prior

[3]On the development of India's software industry and its evolving service capability, see Arora and Athreye (2002) and Athreye (2005); on reputation, contracting, and certification, see Banerjee and Duflo (2000) and Arora and Asundi (1999); on human-capital and migration responses associated with the sector's expansion, see Ghose (2026) and Khanna and Morales (2025); and on the economy-wide consequences of services-led growth, see Fan et al. (2023).

advantage and enabling conditions. Section 8 evaluates alternative explanations. Section 9 concludes.

2 Data

This paper draws on nine data sources. The spine is a digitized firm panel built from the *NASSCOM Directory of Indian Software & Service Companies*, which records India's IT firms and the capabilities each held. The other eight sources play one of two roles: some corroborate the directory's self-reports, and others supply evidence it does not collect. Together they let me observe the Y2K demand shock, the firms exposed to it, whether those firms upgraded, and the channels and confounds that bear on how I interpret the results. Appendix A details the full construction and validation record, from extraction and entity resolution to variable construction, cross-dataset integration, and known limitations.

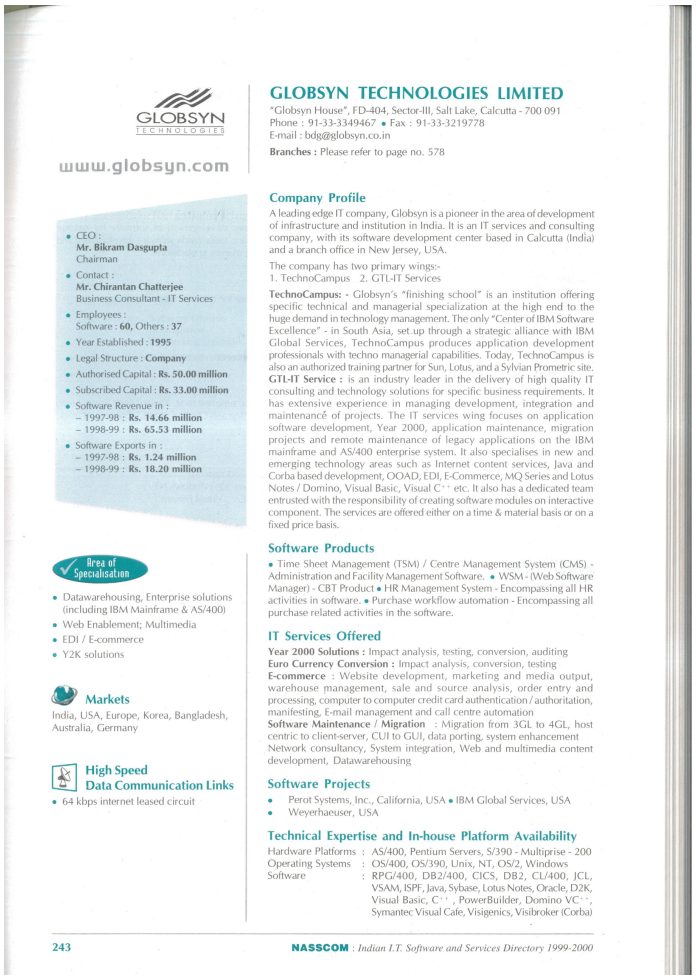
2.1 Data sources

Table 1 lists the nine sources; I describe each below.

NASSCOM directories. The National Association of Software and Service Companies (NASSCOM), the industry trade body for export-oriented IT, published an annual member directory through the 1990s and early 2000s. These directories record a detailed, structured profile for each member firm. I digitize the seven surviving editions (1992, 1995, 1998, 1999, 2001, 2002, and 2003) into a panel of 1,248 firms and 3,203 firm-year records. Each profile reports firm size, software revenue and exports, ownership and age, the service and industry mix, named clients, the programming-language and platform stack, and quality certifications. Figure 1 shows a representative page. From the 1998 edition onward, the directory adds two columns, *Year 2000 Solutions* and *Euro Currency Solutions*, in which firms report whether they undertake Year 2000 remediation and Euro-conversion work.

The directory is a voluntary member roster, not a population register. Its often-cited 95% coverage figure refers to industry revenue, which a handful of large exporters dominate, not to the count of firms. The economically significant universe is therefore captured densely and the small-firm tail sparsely (Arora and Athreye 2002). I build the paper's primary Y2K exposure measures from the technology stacks firms reported in the 1995 directory, before the shock became salient. The binary L_i^{bin} marks firms whose COBOL or mainframe stack matched Y2K remediation requirements. The coherence measure D_i^{Y2K} refines this,

grading how complete and internally coherent that legacy match was. Section 5 sets out the construction of both.



GLOBSYN TECHNOLOGIES

www.globsyn.com

GLOBSYN TECHNOLOGIES LIMITED

"Globsyn House", FD-404, Sector-III, Salt Lake, Calcutta - 700 091
Phone : 91-33-3349467 • Fax : 91-33-3219778
E-mail : bdg@globsyn.co.in

Branches : Please refer to page no. 578

- CEO : **Mr. Bikram Dasgupta** Chairman
- Contact : **Mr. Chirantan Chatterjee** Business Consultant - IT Services
- Employees : Software : **60**, Others : **37**
- Year Established : **1995**
- Legal Structure : **Company**
- Authorised Capital : **Rs. 50.00 million**
- Subscribed Capital : **Rs. 33.00 million**
- Software Revenue in : – 1997-98 : **Rs. 14.66 million** – 1998-99 : **Rs. 65.53 million**
- Software Exports in : – 1997-98 : **Rs. 1.24 million** – 1998-99 : **Rs. 18.20 million**

Area of Specialisation

- Datawarehousing, Enterprise solutions (including IBM Mainframe & AS/400)
- Web Enablement; Multimedia
- EDI / E-commerce
- Y2K solutions

Markets

India, USA, Europe, Korea, Bangladesh, Australia, Germany

High Speed Data Communication Links

- 64 kbps internet leased circuit

Company Profile

A leading edge IT company, Globsyn is a pioneer in the area of development of infrastructure and institution in India. It is an IT services and consulting company, with its software development center based in Calcutta (India) and a branch office in New Jersey, USA.

The company has two primary wings:-
1. TechnoCampus 2. GTL-IT Services

TechnoCampus: - Globsyn's "finishing school" is an institution offering specific technical and managerial specialization at the high end to the huge demand in technology management. The only "Center of IBM Software Excellence" - in South Asia, set up through a strategic alliance with IBM Global Services, TechnoCampus produces application development professionals with techno managerial capabilities. Today, TechnoCampus is also an authorized training partner for Sun, Lotus, and a Sylvian Prometric site. **GTL-IT Service :** is an industry leader in the delivery of high quality IT consulting and technology solutions for specific business requirements. It has extensive experience in managing development, integration and maintenancé of projects. The IT services wing focuses on application software development, Year 2000, application maintenance, migration projects and remote maintenance of legacy applications on the IBM mainframe and AS/400 enterprise system. It also specialises in new and emerging technology areas such as Internet content services, Java and Corba based development, OOAD, EDI, E-Commerce, MQ Series and Lotus Notes / Domino, Visual Basic, Visual C++ etc. It also has a dedicated team entrusted with the responsibility of creating software modules on interactive component. The services are offered either on a time & material basis or on a fixed price basis.

Software Products

• Time Sheet Management (TSM) / Centre Management System (CMS) - Administration and Facility Management Software. • WSM - (Web Software Manager) - CBT Product • HR Management System - Encompassing all HR activities in software. • Purchase workflow automation - Encompassing all purchase related activities in the software.

IT Services Offered

Year 2000 Solutions : Impact analysis, testing, conversion, auditing
Euro Currency Conversion : Impact analysis, conversion, testing
E-commerce : Website development, marketing and media output, warehouse management, sale and source analysis, order entry and processing, computer to computer credit card authentication / authoritation, manifesting, E-mail management and call centre automation
Software Maintenance / Migration : Migration from 3GL to 4GL, host centric to client-server, CUI to GUI, data porting, system enhancement Network consultancy, System integration, Web and multimedia content development, Datawarehousing

Software Projects

- Perot Systems, Inc., California, USA • IBM Global Services, USA
- Weyerhaeuser, USA

Technical Expertise and In-house Platform Availability

Hardware Platforms	:	AS/400, Pentium Servers, S/390 - Multiprise - 200
Operating Systems	:	OS/400, OS/390, Unix, NT, OS/2, Windows
Software	:	RPG/400, DB2/400, CICS, DB2, CL/400, JCL, VSAM, ISPF, Java, Sybase, Lotus Notes, Oracle, D2K, Visual Basic, C++ , PowerBuilder, Domino VC++, Symantec Visual Cafe, Visigenics, Visibroker (Corba)

Figure 1: A representative NASSCOM directory profile page.

Notes: One firm's profile from the digitized 1999 directory. Each member firm occupies a structured entry reporting headcount, software and total revenue, exports, ownership, the service and industry mix, named clients, the platform and language stack, certifications, and, from 1998, the Year 2000 and Euro self-report columns.

SEC–EDGAR Y2K disclosures. I measure the foreign Y2K demand shock from the client side, using US corporate disclosures of remediation spending. In 1998 the Securities and Exchange Commission required public companies to disclose their Year 2000 readiness, costs, and risks (U.S. Securities and Exchange Commission 1998). I extract these disclosures from 110,082 10-K and 10-Q filings spanning 1993 to 2005, drawn from the cleaned EDGAR corpus of Loughran and McDonald (2011). Aggregating disclosed dollars to nine sectors — the eight NASSCOM industry groups plus a residual — yields a $52 billion client-demand

universe, within the contemporaneous \$100–115 billion US total. This forms the Bartik demand shifter $g_{s,t} = g_s \times G_t$, whose time path rises sharply in 1998 and 1999 and then falls to near zero once the deadline passes. The sector ranking, with telecom and finance far above retail, reproduces the contemporaneous Senate and Gartner most-vulnerable lists, and the largest disclosed programs match the corporate Y2K budgets reported in the period's press.

Prowess. Prowess, a firm-level financial database maintained by the Centre for Monitoring Indian Economy (CMIE), provides an audited counterpart to the self-reported directory and extends the analysis beyond IT. It adds audited financial statements and balance-sheet detail that the directory does not collect, and it covers listed and larger firms across the whole economy, not only in IT. I use Prowess in three ways. First, I match 457 of the 1,248 panel firms to Prowess to benchmark the directory against audited accounts and to check robustness on the large-firm tail. Across overlapping firm-years, audited and self-reported revenue line up at the median—their median $\log_{10}$ ratio is -0.008—though agreement across firm-years is heterogeneous, with 41% agreeing within 10%. Second, because Prowess spans every sector, I use it to compare IT firms with non-IT firms around the Y2K deadline, a comparison the IT-only directory cannot support. Third, I use it on its own to build a within-IT regional exposure covering the IT entrants the matched sample cannot reach (Section 6.3).

STPI annual reports. The annual reports of the Software Technology Parks of India (STPI) record the rollout of export connectivity across India, which I use to identify the infrastructure channel. I digitize the 13 surviving reports, spanning 1991 to 2006, into a panel of 58 centers with opening-year and capacity information, alongside a national software-export series that runs through FY2004. STPI gave firms more than connectivity. It offered a single-window export regime: bonded premises, dedicated high-speed satellite and leased-line links, customs and regulatory clearance, and administration of the export tax holidays. A firm could set up an export operation inside an STPI park without assembling these pieces itself. By 1998, 92% of the panel's firms sat in a district with an operating STPI center, up from 49% in 1992.

SHRUG / Economic Census. District-level economic activity comes from the Economic Censuses, accessed through the SHRUG platform. SHRUG, the Socioeconomic High-resolution Rural-Urban Geographic platform, harmonizes the Economic and Population Cen-

suses and other Indian administrative data onto a common, time-consistent geographic key (Asher et al. 2021). I take district IT employment from the Economic Censuses of 1990, 1998, 2005, and 2013, and join it to the firm panel through a 2001-Census district key. The firm geography lines up with the census IT counts, with a district-level rank correlation of 0.83. These data supply the local economic context and a district-level exposure measure for the geographic analysis.

SEMIS and NSS. Two education datasets test whether the demand shock induced a local schooling response. The first is the Secondary Education Management Information System (SEMIS), an annual administrative census of India's secondary and higher-secondary schools, which I downloaded from the SEMIS portal in 2018 while the data was still available online. It records each school's founding year, management type, medium of instruction, and level, covering roughly 270,000 schools a year. Its founding-year field lets me build secondary-school entry rates by district and year, the margin a demand-driven supply response could move. The second is the National Sample Survey (NSS), a repeated household survey representative at the state level. I use six rounds spanning 1993 to 2017, which carry schooling and consumption schedules. Appendix B.10 uses SEMIS for the final education-supply check; the NSS outcomes are exploratory and do not support the paper's conclusion.

BEA. The Y2K deadline falls inside the dot-com investment boom, the most obvious confound. To separate the two, I use Bureau of Economic Analysis (BEA) Fixed Assets and GDP-by-industry data to build a measure of IT capital deepening, g^{IT}, that tracks the dot-com investment wave. It enters as a control alongside the Y2K legacy-remediation shifter, so the dot-com wave and remediation demand compete to explain the same outcomes. Compustat revenue denominators and federal Y2K outlays from OMB and GAO complete the shifter construction.

2.2 The firm-year panel

Building a firm-level panel from the directories requires resolving three threats to firm identity across editions: name changes, mergers and acquisitions, and exits and re-entries. I address them in two stages, both detailed in Appendix A. First, a language model reads each directory profile twice and distills it into structured fields. Second, I assign each firm a stable identifier and trace its history across editions with an LLM-based research procedure, checking its name changes, acquisitions, and continued existence against public sources, in-

Table 1: Data sources

Source	**Level**	**Coverage**	**Key variables**	**Use**
NASSCOM directories	Firm	7 editions, 1992–2003; 1,248 firms / 3,203 firm-years	Size, exports, certifications, clients, legacy stack, Y2K/Euro self-reports	Headline panel: outcomes and the predetermined exposure (L_i^{bin})
SEC–EDGAR Y2K corpus	US filing	110,082 10-K/10-Q, 1993–2005	Size, timing, sector mix of foreign Y2K demand ($52B disclosed)	The Bartik shifter $g_{s,t}$; the shock's magnitude and deadline
Prowess (CMIE)	Firm (audited)	457 matched firms + the full cross-industry universe; FY1987–2018	Audited financials vs. the self-reports (median $\log_{10}$ ratio -0.008)	External validation (matched firms); cross-industry and within-IT regional corroboration
STPI annual reports	Center / national	58 centers; reports 1991–2006, exports to FY2004	Connectivity rollout timing and the aggregate export trajectory	Infrastructure channel; macro context
SHRUG / Econ. Census	District	EC 1990/1998/2005/2013	Local IT employment; district exposure (EC-IT Spearman 0.83)	District outcomes; geographic exposure
SEMIS	School	∼270k secondary schools; 13 waves	Private/English school founding	Education-supply check
NSS	Household / state	6 rounds, 1993–2017	Attendance, private schooling, expenditure	Exploratory only
BEA	Industry	Fixed Assets; GDP-by-industry	IT capital deepening (the dot-com wave)	Dot-com control g^{IT}
Compustat; OMB/GAO	Firm; federal	Fundamentals; Y2K outlays	Revenue denominators; government Y2K spend	Shifter denominators and supplement

Notes: Together the sources let me observe the shock (EDGAR), the firms exposed to it (NASSCOM), whether they upgraded (NASSCOM certifications), and the channels and confounds (STPI, SHRUG, SEMIS/NSS, BEA); Prowess corroborates the directory's financial outcomes in audited accounts. Construction and validation detail for each source is in Appendix A.

cluding news archives and the Ministry of Corporate Affairs (MCA) registry. The procedure keeps an acquirer and its target separate, so that M&A-driven growth can be netted out from organic growth, and it records 150 within-panel mergers and acquisitions that I use as time-varying controls. The result is a panel of 1,248 firms and 3,203 firm-years. The same procedure validates the roster. Registry checks adjudicate the status of three-quarters of the firms, and their self-reported founding years agree with the registry to within one year in 82% of cases.

The analysis runs on two samples. Descriptive evidence and the district-level analyses use the full panel of 1,248 firms. The firm-level upgrading tests use a narrower estimation universe: the 235 firms listed in the 1995 directory, followed across 1,015 firm-years. The two exposure measures partition this universe differently. The binary L_i^{bin} splits it into 135 firms with the legacy match and 100 without. The coherence gate D_i^{Y2K} refines this, separating the 55 firms with a coherent legacy production stack from the 170 without; 10 firms with no readable stack leave the gate undefined and are dropped, leaving a 225-firm common sample. Roster attrition reflects directory nonappearance rather than firm exit and tracks firm size, not exposure, so it does not threaten the comparisons, and inverse-probability-weighted checks confirm this. A firm carries an exposure measure only if it was listed in 1995, so firms first observed later have none and appear in the post-1995 entrants column. Table 2 reports the sample by wave (Panel A) and the key outcomes over the estimation universe (Panel B).

Exposed firms differ from comparison firms in expected ways. Because the legacy stack that exposed a firm to Y2K demand was typically built up by established companies, exposed firms are on average larger, older, and more oriented to banking and finance (Appendix Table B.1). Identification does not rest on these firms being comparable at baseline. It rests instead on two features of the setting: the Y2K shock was externally dated and temporary, and a firm's exposure was its programming stack as of 1995, fixed before the shock arrived. The empirical strategy therefore reads upgrading off within-firm changes timed to the deadline, not off the cross-sectional gap between exposed and comparison firms, and Section 5 develops the argument. One data feature does bear on the dollar outcomes. Exposed firms report revenue and exports more often than comparison firms, so those variables are observed on a differentially selected sample, a point I return to when interpreting them.

Table 2: The analysis sample

Panel A: composition by directory wave

		Legacy match L_i^{bin}		Coherence gate D_i^{Y2K}			Post-1995
Wave	Observed	Exposed	Control	Coherent	Non-coh.	No stack	entrants
1992	124	56	35	31	60	0	33
1995	235	135	100	55	170	10	0
1998	408	106	73	44	131	4	229
1999	486	92	68	43	113	4	326
2001	657	85	59	38	102	4	513
2002	687	67	47	30	80	4	573
2003	606	57	35	25	65	2	514

Panel B: median [mean] of key outcomes, estimation universe

Wave	Employment	Software exports ($m)	Software revenue ($m)	Capital ($m)
1992	80 [434]	0.4 [3]	0.5 [2]	2.5 [43]
1995	65 [216]	0.5 [4]	0.6 [4]	7.0 [44]
1998	115 [402]	1.4 [10]	1.4 [9]	14.2 [49]
1999	151 [538]	2.2 [15]	2.4 [13]	23.7 [84]
2001	217 [712]	3.4 [28]	3.9 [30]	45.0 [99]
2002	255 [954]	5.0 [34]	5.3 [38]	51.8 [119]
2003	380 [1,248]	5.2 [44]	5.9 [48]	66.5 [192]

Notes: The compiled firm–year panel (1,248 firms, 3,203 firm–years). Panel A splits each wave's firms by the two headline exposure measures, each measured on the 1995 estimation universe: the legacy-match binary L_i^{bin} (legacy-matched vs. comparison) and the deeper coherence gate D_i^{Y2K} (a coherent legacy production stack vs. not; "No stack" are firms with no readable pre–shock stack listing, for which the gate is undefined and which are dropped from the depth specifications, leaving the 225–firm common sample). Each split sums to the estimation universe; the gate is a near–subset of L_i^{bin} (54 of the 55 coherent firms are also L_i^{bin}–exposed, one is a 1992–listing pickup). "Post-1995 entrants" are firms first listed after 1995, which carry no pre–shock exposure measure; in the 1992 row the column instead counts the 33 firms listed in 1992 but not in 1995, which also carry no exposure measure. Panel B reports the median and, in brackets, the mean of each outcome over the estimation–universe firms reporting it in that wave; the gap between median and mean reflects the right tail (the largest firms are unwinsorized). Coverage by exposure group and the 1995 covariate comparison are in Appendix Table B.1; the legacy-matched-vs-comparison trajectories are in Figure 4.

3 Background and Stylized Facts

3.1 Brief History of the Y2K Shock

The Millennium Bug, or Y2K problem, originated in a decades-old programming convention. To save memory at a time when storage was expensive, mainframe software written from the 1960s through the 1980s routinely stored calendar years as two digits rather than four. A program built this way could not distinguish the year 2000 from the year 1900. On January 1, 2000, any date calculation, sort, or comparison that relied on the year field risked returning the wrong answer or failing outright. The risk was concentrated in the systems that could least afford to fail — payroll, banking, insurance, and utility billing (United States Senate, Special Committee on the Year 2000 Technology Problem 2000) — and it activated on a single calendar date that no institution could move. The fixed date separates Y2K from an ordinary technology-adoption problem. Compliance could not be phased in, and remediation had to be substantially complete before the deadline.

Fixing the problem was not a matter of installing a new package or a patch. Because the two-digit convention was embedded throughout each program's logic rather than confined to a single field or module, remediation meant locating every affected date calculation and rewriting it, one line at a time. Each fix had to be written in the same language and on the same platform as the original code — overwhelmingly COBOL, running on mainframe hardware. The work was therefore labor-intensive rather than capital-intensive, and a firm's ability to do it turned on a specific technical match: fluency in COBOL and related legacy languages, and familiarity with mainframe operating environments. A supplier that had never worked in this stack could not retrain in time for the deadline; a supplier that already had could take on the work immediately. That match between an old and otherwise unremarkable technical specialization and a sudden, precisely dated foreign demand for it is the source of the exposure variation this paper uses, and Section 5 builds on it directly.

3.2 Three Stylized Facts

Three patterns in the data motivate the empirical strategy. First, the Y2K shock was large and precisely dated. Second, it concentrated in the mainframe-reliant sectors and rewarded the Indian firms that already held the matching legacy stack, which then diverged from the firms without it. Third, the gains took a form that persisted past the deadline.

Magnitude and timing. The Y2K remediation market was large and ran on a hard deadline. US public companies disclosed about $52 billion in Y2K-remediation spending, within the contemporaneous $100–115 billion US total. Figure 2 traces this demand in their filings. Panel (a) plots the share of 10-K and 10-Q filings carrying a substantive Y2K disclosure. The rate is near zero through 1997, jumps in early 1998, and climbs to its peak near 87% in 1999 once the SEC's disclosure mandate takes hold. It collapses once the January 2000 deadline passes. Panel (b) shows the same deadline in dollars and in language. Cumulative disclosed spending reaches roughly $46 billion by the first quarter of 2000, and the share of cost disclosures reporting an actual rather than an estimated figure flips from near zero to near one in that quarter. The demand was real, and it expired on schedule. On the Indian side the same shock left no comparable signature in the aggregate. National software exports climbed more than tenfold across these years, but as a smooth secular rise. Y2K remediation was one strand of a broader offshoring boom, so the national series shows no discrete bump at the deadline and no collapse after it. I therefore identify the Y2K channel from firm-level exposure, not from the national time series.

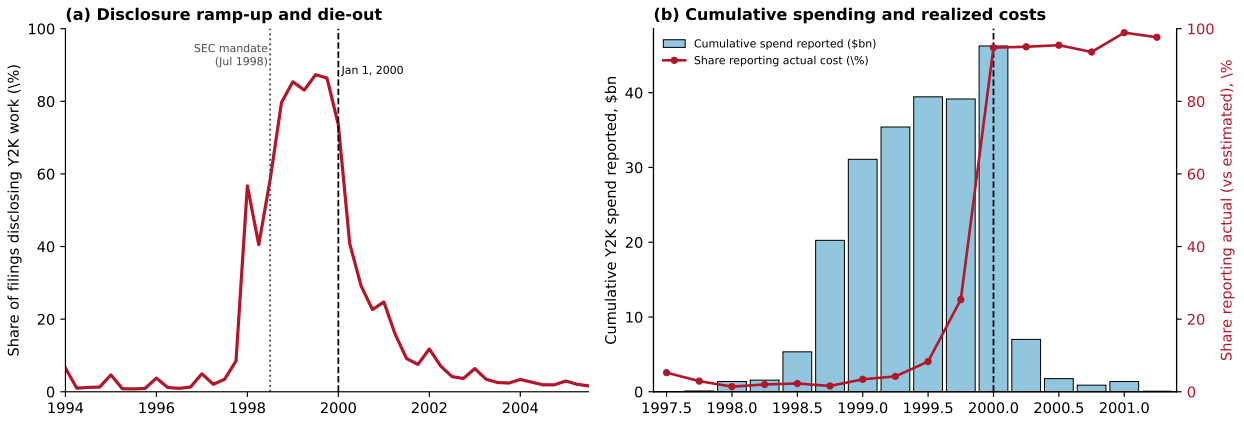


Figure 2: The foreign Y2K demand was large and expired on schedule.

Notes: From the SEC–EDGAR corpus of US 10-K/10-Q filings. (a) Share of all filings carrying a substantive Y2K-remediation disclosure, by quarter; the grey dotted line marks the SEC's July 1998 disclosure mandate and the dashed line the January 2000 deadline. (b) Bars: cumulative Y2K spending reported in filings, by quarter of filing ($bn); line (right axis): the share of cost disclosures reporting an actual rather than an estimated figure. The post-2000 fall in panel (b) is reporting die-out, not a reversal of spending.

Exposure and divergence. The demand for that legacy match did not come evenly across the US sectors that generated it. Figure 3 shows the median Y2K bill per firm by US sector (panel a) alongside the share of each sector's filings that mention mainframe systems (panel b). The heaviest-spending sectors are utilities, insurance carriers, and securities, and the

sectors that spent more were generally the more mainframe-reliant, so the demand tracked the legacy systems that had to be remediated. The work rewarded firms that already spoke the legacy languages. In 1995, 135 of the 235 firms in the analysis sample (about 57%) reported a legacy-stack match in the COBOL and mainframe technologies the remediation work demanded. I use this binary indicator, L_i^{bin}, as the exposure measure. A stricter measure, the coherence gate D_i^{Y2K}, asks not whether a firm reported legacy technology but whether it had assembled a coherent legacy production stack; 55 firms clear it, on the 225-firm common sample. Section 5 constructs both.

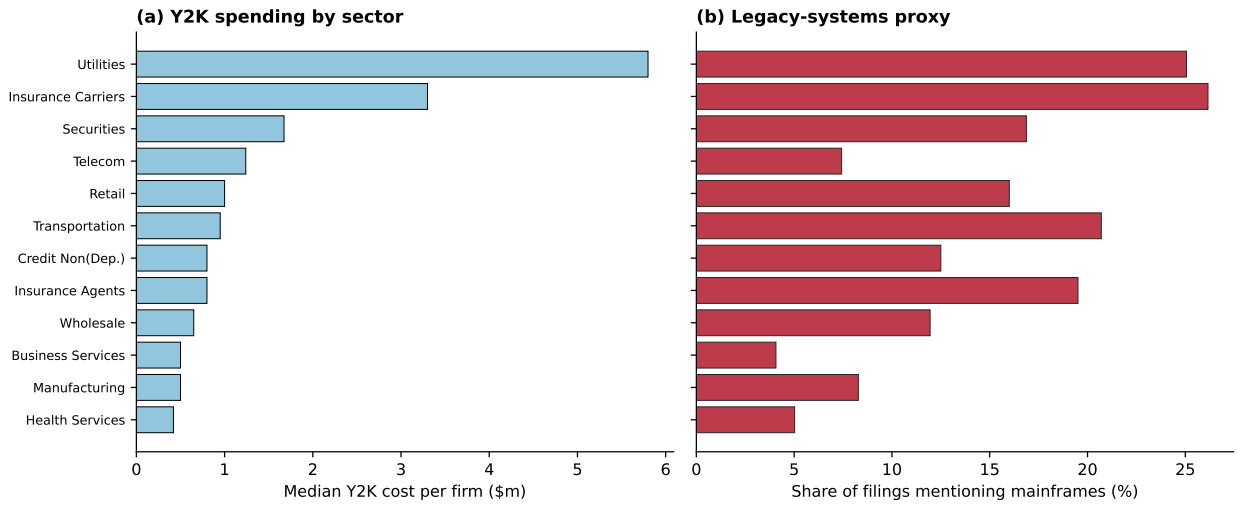


Figure 3: The Y2K shock fell on the mainframe–reliant sectors.

Notes: (a) Median disclosed Y2K cost per firm, by sector ($m), among firms reporting Y2K content in the SEC–EDGAR corpus (the finer EDGAR filing sectors, not the nine client-industry groups used in the shift-share; sectors with at least 40 such firms, the top twelve by median cost). (b) Share of each sector's filings that mention mainframe systems; panel (b) keeps panel (a)'s order. Banking falls outside the top twelve here, its per-firm median held down by many small community banks, though in aggregate it is a large Y2K spender.

Exposed and comparison firms then followed visibly different paths. Within the analysis sample the gap between the two groups is roughly stable through 1995 and widens sharply from the shaded Y2K window onward, on software exports, software revenue, software and engineering headcount, and total employment alike, with no closing after the deadline (Figure 4). The figure is only descriptive. Exposed firms start larger, especially on the dollar margins, survivor composition changes across waves, and the divergence could be confounded by many factors. I isolate the upgrading channel in Section 6.2.

Persistence. The temporary demand left a durable mark. Exposed firms upgraded, and the upgrading did not fade when the demand ended. Certification adoption diverges between

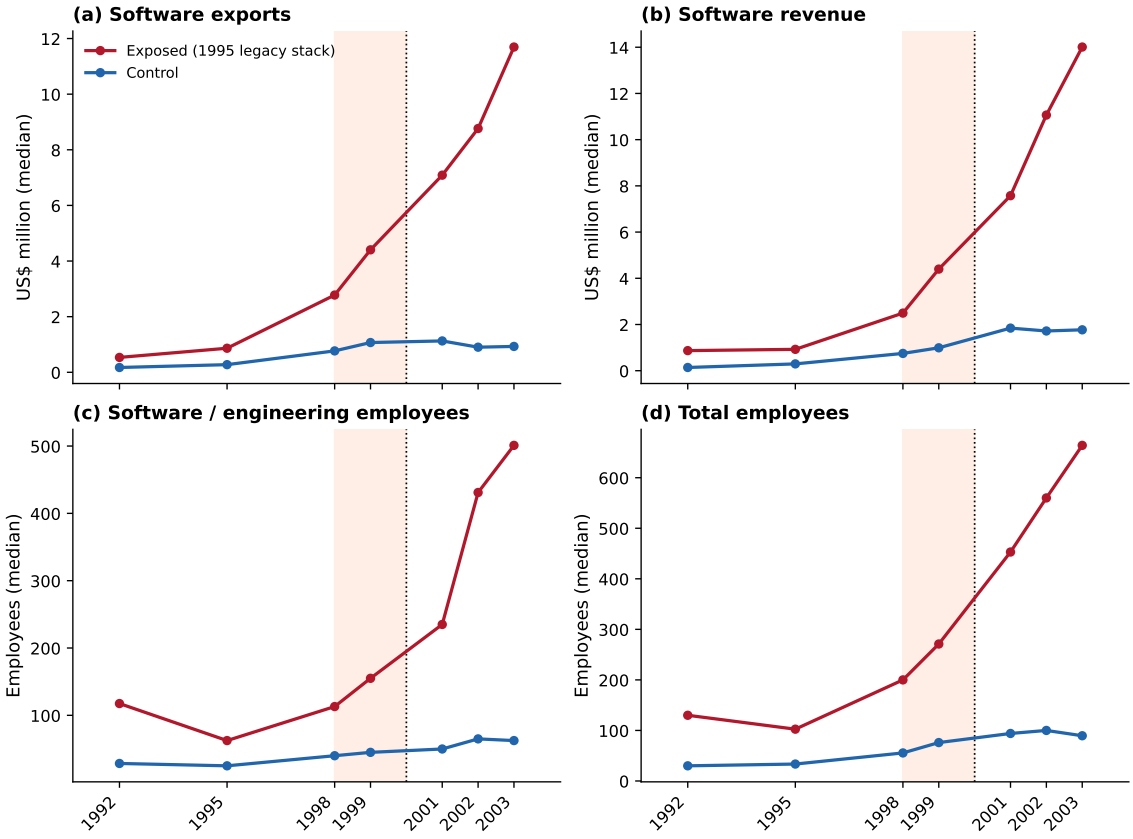


Figure 4: Raw trajectories, exposed vs comparison firms (analysis sample; descriptive).

Notes: Median (a) software exports, (b) software revenue, (c) software/engineering employees, and (d) total employees, by directory wave and 1995 exposure; the shaded band marks the Y2K window (1998–2000) and the dotted line the January 2000 deadline. Unconditional within the 235-firm analysis sample; the exposed cell falls from 135 firms in 1995 to 57 in 2003, so the 1992–1995 movement in the headcount panels partly reflects which firms report. The level persistence shown here is the scale margin tested formally in Section 6; it is not the upgrading-vs-expansion decomposition, which is Section 6.2.

exposed and comparison firms from the mid-1990s, widens through the deadline, and remains wide well past it, still growing on certification depth through 2003 (Figure 5). Both groups certify, with the comparison firms catching up late, but the pattern is a continued differential adoption. The small exposed lead visible in 1992 is the early-certifier pre-trend I disclose and address in the formal result (Section 6.2). That the gap persists rather than reverts points to a durable capability upgrade, not a reversible expansion.

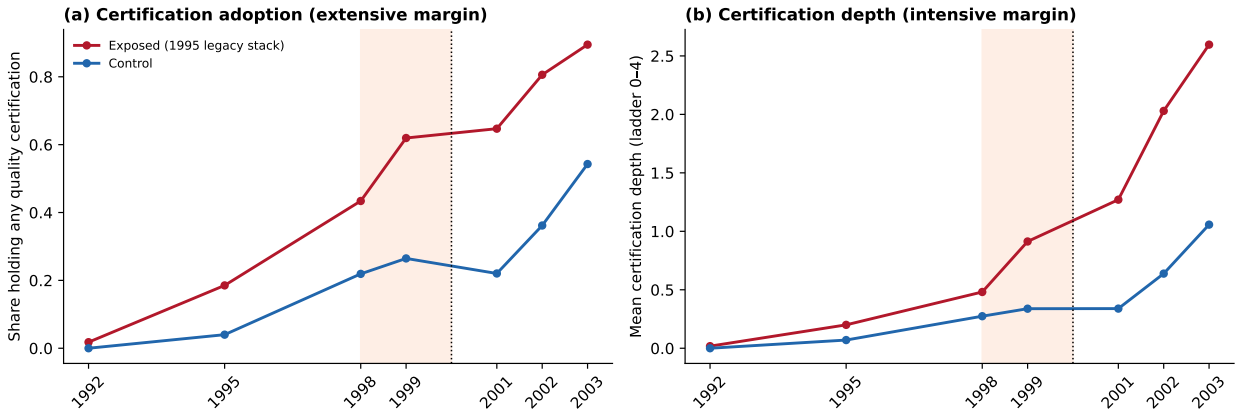


Figure 5: Certification adoption, exposed vs comparison firms.

Notes: (a) Share holding any quality certification; (b) mean certification depth (ladder 0–4), by directory wave and 1995 exposure, within the 235-firm analysis sample. Descriptive and unconditional (the exposed cell falls 135 → 57 across waves). The formal within-firm fingerprint and the no-persistent-differential test on the expansion margins are Section 6.2.

Together these patterns guide the empirical strategy of the paper: (1) an externally dated, temporary demand shock, (2) firm-level divergence in exposure, and (3) a persistent upgrading margin.

4 A Simple Model of Demand-Driven Upgrading

Temporary export opportunities can leave persistent effects through several channels. Access to a larger market can make fixed productivity or technology investments profitable (Lileeva and Trefler 2010; Bustos 2011). Exporting can also generate learning through production experience and knowledge transferred by foreign buyers (Atkin et al. 2017a), while foreign demand for quality can reward internal upgrading (Verhoogen 2008; Verhoogen 2023). Firms may also acquire relationships, reputation, approved-vendor status, or foreign presence that can support persistent activity because incumbents avoid the sunk entry or re-entry costs faced by firms without those assets (Baldwin and Krugman 1989; Roberts and Tybout 1997;

Das et al. 2007). These mechanisms can operate together, and persistent activity alone does not reveal which mechanism is responsible.

Y2K is revealing for two reasons. First, the remediation opportunity was temporary and externally dated. Direct Y2K demand ended at a known deadline, so continued remediation work cannot explain persistence. Within the model, persistence therefore requires a state that outlives the window. Second, Y2K assigned the opportunity through a technologically specific requirement that was conceptually distinct from general productivity and modern technological leadership. A matching legacy stack was necessary to perform remediation work. Among matched firms, however, prior foreign reach and how widely buyers searched as the deadline approached could still affect the amount of work received. The access extension formalizes this distinction, and Section 7 asks whether observed responses vary with pre-shock reach.

Receiving Y2K work bundled several experiences. It increased project load, connected firms to buyers and global delivery networks, and could expose them to new standards, feedback, and production knowledge. The model takes project load as its baseline because it provides a transparent account of the two outcomes the paper observes most directly: employment and certification, with certification proxying latent organizational capability. Ordinary and remediation projects use the same delivery organization. As more work is spread across a finite stock of organizational attention, coordination, testing, and rework losses rise. A process and quality-management system expands the effective attention the organization can deploy, so its value rises with the scale of the operation. Temporary demand therefore generates two responses. The *expansion response* is the increase in operating scale while remediation demand lasts, holding durable states fixed. The *upgrading response* is earlier investment in the process and quality-management system itself.

One benchmark and two substantive results give this mechanism empirical content. The expansion benchmark is temporary by construction: the fixed-state increase in scale vanishes when remediation demand ends. The first substantive result is that the upgrading response occurs only near the no-Y2K upgrading margin: firms with strong scale–capability complementarity upgrade early anyway, firms with weak complementarity continue to wait, and intermediate firms change their timing because of Y2K. The second is the obsolete-match reversal. Once acquired, the capability remains useful in ordinary work, so an obsolete technical match can predict organizational capability and the scale it supports after the match has lost direct value. The predetermined match does not separately identify project load, access, buyer standards, learning, and financing. The project-load baseline is therefore

one mechanism, not an identified decomposition of this bundled experience. Prior access, buyer assurance, and persistent foreign-market participation enter as distinct extensions or alternatives.

The first subsection derives the firm's scale choice and the return to organizational capability. The second solves the upgrading-timing problem and characterizes persistence. The final subsection states the model's empirical implications and interpretation boundary.

4.1 Environment and the firm problem

Firms begin with predetermined technology stacks and ongoing non-remediation business. The Y2K shock creates additional demand for services that a firm can provide only if its stack matches the legacy technologies the work requires. Each firm makes two choices: in every period it chooses its operating scale, and it decides whether and when to incur a sunk cost to upgrade its organizational capability. The model represents this upgrade as a move from the firm's baseline organizational system to a durable process and quality-management system. Such a system reduces the coordination, testing, and rework costs of operating at scale. The direct implication that more demand raises current scale is immediate. The substantive choice is when to pay the sunk upgrading cost. A firm with a valuable upgrade may still rationally wait, and temporary Y2K demand changes behavior only when its finite return overcomes that value of waiting. Once induced, the capability can support ordinary-market scale after remediation ends.

Time is discrete, and the end date T is known from the outset. The remediation window is $\mathcal{W} \equiv \{0, 1, \ldots, T\}$. Let $L_i \in \{0, 1\}$ denote firm i's predetermined legacy-technology match. It equals one when the firm's pre-shock technology stack includes qualifying COBOL or mainframe technologies relevant to Y2K remediation, and zero otherwise. Let $S_t \geq 0$ denote the common intensity of the global Y2K demand shock, with $S_t > 0$ only for $t \in \mathcal{W}$. Let y_{it} index the incremental Y2K demand facing that firm. Let $\mu_{it} > 0$ index demand for firm i's ordinary, non-remediation services. The baseline writes these demand indices as

$$y_{it} = L_i S_t, \qquad m_{it} = \mu_{it} + y_{it}. \tag{1}$$

Thus y_{it} isolates the temporary remediation-demand component, while m_{it} is the total project-demand index facing the firm. The baseline assumes that ordinary and remediation projects use the same delivery organization, so their combined demand determines one operating-scale choice. Equation (1) gives technical match a specific role: L_i activates the

potential Y2K demand. The baseline abstracts from differences in access to Y2K work among matched firms; Appendix C allows prior foreign reach and deadline search to scale the firm-level Y2K demand index y_{it}.

Given these demand paths, $x_{it} \geq 0$ denotes the volume of project or delivery activity the firm chooses to operate in period t. Employment is its closest empirical counterpart. Scarce organizational attention does not impose a hard upper bound on x; it makes expansion increasingly costly. The organizational-capability state $q_{it} \in \{0, 1\}$ distinguishes the firm's baseline organizational system from an upgraded process and quality-management system. A firm in state $q = 0$ may pay $K_i > 0$ once to enter state $q = 1$, with documented workflows, project monitoring, testing protocols, and repeatable quality controls. The organizational state is latent; certification is an observed proxy for it. These routines are in the same family as the codified monitoring and quality-control systems studied experimentally by Bloom et al. (2013). Once a firm upgrades, the baseline keeps it in $q = 1$. This is a durability assumption through the post-deadline years observed in the paper, not a claim of literal permanence, and it is consistent with evidence that management practices can persist (Bloom et al. 2020).

The scale–capability complementarity follows from allocating scarce organizational attention across operating units. An operation of scale $x > 0$ consists of a continuum of project or delivery units $j \in [0, x]$. Let $z_j > 0$ denote the effective organizational attention assigned to unit j. In organizational state q, the firm can allocate total attention $A_{iq} > 0$ across its operation. The upgraded system allows a given managerial organization to coordinate more activity, so $A_{i1} > A_{i0}$. This capacity gain is the organizational-technology primitive; it does not by itself determine whether or when the firm pays K_i. Effective organizational attention need not be literal manager-hours. It summarizes the coordination capacity created by documented workflows, monitoring systems, and quality controls. This interpretation treats those practices as an organizational technology (Bloom et al. 2013) and follows the broader idea that managerial attention and production knowledge are scarce organizational resources (Geanakoplos and Milgrom 1991; Garicano 2000).

At a unit receiving attention z_j, let expected coordination, testing, and rework loss be $\ell_i(z_j)$. The units are symmetric by assumption, and the loss function is decreasing and strictly convex: $\ell_i'(z) < 0$ and $\ell_i''(z) > 0$. The same loss function applies in both organizational states; upgrading changes effective capacity A_{iq} rather than the shape of ℓ_i. The firm allocates its organizational capacity to minimize total operating losses:

$$C_{iq}(x) = \min_{\{z_j > 0\}_{j \in [0,x]}} \int_0^x \ell_i(z_j)\, dj \qquad \text{subject to} \qquad \int_0^x z_j\, dj \leq A_{iq}. \tag{2}$$

Because the loss falls with attention, the capacity constraint binds. Total attention is therefore A_{iq}, and average attention across the x operating units is A_{iq}/x. Because the units are symmetric and the loss is strictly convex, Jensen's inequality implies that total loss is minimized by assigning this average to every unit. Thus

$$z_j^* = \frac{A_{iq}}{x} \quad \text{for every } j, \qquad C_{iq}(x) = x\ell_i\left(\frac{A_{iq}}{x}\right). \tag{3}$$

To obtain closed-form scale and upgrading choices, assume that the per-unit loss function is $\ell_i(z) = \theta_i/(2z)$. The parameter $\theta_i > 0$ scales the coordination, testing, and rework loss caused by insufficient organizational attention. Holding attention z fixed, a higher θ_i therefore means a larger expected operating loss. Equation (3) then gives

$$C_{iq}(x) = x\frac{\theta_i}{2(A_{iq}/x)} = \frac{\theta_i}{2A_{iq}}x^2, \qquad C_{iq}(0) = 0. \tag{4}$$

At a fixed scale $x > 0$, the operating-cost saving from upgrading is

$$C_{i0}(x) - C_{i1}(x) = \frac{\theta_i x^2}{2}\left(\frac{1}{A_{i0}} - \frac{1}{A_{i1}}\right) > 0, \qquad \frac{d}{dx}\left[C_{i0}(x) - C_{i1}(x)\right] = \theta_i x\left(\frac{1}{A_{i0}} - \frac{1}{A_{i1}}\right) > 0. \tag{5}$$

Equation (5) separates the fixed-scale return to upgrading into three components. The parameter θ_i measures the severity of attention-related losses, x^2 captures the strain created by operating at scale, and $(1/A_{i0} - 1/A_{i1})$ captures how much the upgrade relaxes the capacity constraint. The x^2 term has two sources: a larger operation contains more units, and a fixed capacity leaves less attention for each unit. Holding the other terms fixed, the cost saving is therefore larger when attention-related losses are more severe, scale is greater, or the upgrade produces a larger reduction in inverse capacity. This is a fixed-scale comparison; the firm has not yet been allowed to adjust x in response to demand.

The attention-allocation problem determines operating cost at any given scale. The firm, however, chooses scale in response to project demand. The gross return generated at scale x is normalized to $m_{it}x$, so a larger demand index raises the return to operating at scale. Period-t operating profit conditional on q is

$$\pi_{it}(x, q; m_{it}) = m_{it}x - \frac{\theta_i}{2A_{iq}}x^2 = m_{it}x - \frac{c_{iq}}{2}x^2, \qquad c_{iq} \equiv \frac{\theta_i}{A_{iq}}, \tag{6}$$

where $A_{i1} > A_{i0}$ implies $c_{i1} < c_{i0}$. The first term is the gross return from serving project demand at scale x; the second is the organizational operating cost derived in equation (4).

The firm chooses $x \geq 0$ to maximize this expression. Because the objective is strictly concave, the first-order condition $m - c_{iq}x = 0$ gives optimal scale and operating profit:

$$x^*_{it}(q; m) = \frac{A_{iq}}{\theta_i} m = \frac{m}{c_{iq}}, \qquad \Pi^*_{it}(q; m) = \frac{A_{iq}}{2\theta_i} m^2 = \frac{m^2}{2c_{iq}}. \tag{7}$$

Equation (7) also shows how the organizational-capability state changes the scale response to demand. In state q, a one-unit increase in m raises optimal scale by $1/c_{iq}$. Define one-half of the difference between the two state-specific slopes as

$$\psi_i \equiv \frac{1}{2}\left(\frac{1}{c_{i1}} - \frac{1}{c_{i0}}\right) = \frac{A_{i1} - A_{i0}}{2\theta_i} > 0. \tag{8}$$

Equivalently, upgrading steepens the scale response to demand by $1/c_{i1} - 1/c_{i0} = 2\psi_i$. This normalization also makes ψ_i the coefficient on m^2 in the optimized return to capability:

$$\Delta_i(m) \equiv \Pi^*_{it}(1; m) - \Pi^*_{it}(0; m) = \psi_i m^2. \tag{9}$$

Thus ψ_i is a firm parameter measuring the strength of the scale–capability complementarity, while $\Delta_i(m)$ is the optimized return that this parameter generates at demand m. In capacity terms, it is the increase in effective organizational capacity generated by upgrading, divided by twice the loss-severity parameter.

Equations (5) and (9) answer different questions. Equation (5) holds operating scale x fixed and changes the organizational state. Equation (9) holds project demand m fixed, changes the organizational state, and lets the firm reoptimize scale in each state. The distinction matters for θ_i. A higher θ_i magnifies the saving from upgrading at a given scale, but it also contracts optimal scale in both states. After scale is reoptimized, and holding demand and both capacity levels fixed, the return $\Delta_i(m)$ falls with θ_i.

Because $\Delta_i'(m) = 2\psi_i m > 0$, optimized profit has increasing differences in capability and project demand. Process and quality-management capability is more valuable when applied across a larger operation. Moreover, $\Delta_i''(m) = 2\psi_i$, so ψ_i governs how sharply the return to capability rises with demand after scale has been optimized. The inverse-attention assumption provides the closed forms used in the body. Appendix C.1 proves the increasing and strictly convex return for a decreasing and strictly convex attention-loss function common to both organizational states, with interior choices. Appendix C.2 gives sufficient conditions outside the structured attention family. This complementarity parallels models in which access-induced revenue expansion raises the return to fixed-cost productivity

or technology investment (Lileeva and Trefler 2010; Bustos 2011).

4.2 Temporary demand, upgrading timing, and persistence

The preceding subsection showed that organizational capability raises efficient scale and that its return rises with total project demand. This subsection asks whether temporary Y2K demand changes only current scale or also when the firm pays the sunk upgrading cost—and, if so, for which firms. The analysis first isolates the return that Y2K adds to capability, then compares that finite return with the option to wait.

Holding organizational capability fixed, temporary Y2K demand changes optimal scale by

$$x^*_{it}(q; \mu_{it} + y_{it}) - x^*_{it}(q; \mu_{it}) = \frac{y_{it}}{c_{iq}}. \tag{10}$$

Because $c_{iq} > 0$ and $y_{it} = L_i S_t$, the right-hand side is positive for a matched firm whenever $S_t > 0$, and it is zero for an unmatched firm or after the shock ends. This holding-fixed comparison is the *expansion response*. It is not separately observed in employment: firms choose q in response to the same shock, so observed employment combines the flexible scale response with any scale change associated with upgrading.

Ordinary demand makes organizational capability valuable both during and after remediation. Those ordinary-market returns, together with the sunk cost, determine when the firm would upgrade even without Y2K. To isolate what the shock adds to that decision, compare the one-period return to capability under total demand with its return under ordinary demand alone. The difference is the flow Y2K upgrading premium:

$$\begin{aligned} \phi_{it} &\equiv \Delta_i(\mu_{it} + y_{it}) - \Delta_i(\mu_{it}) \\ &= \psi_i \left(2\mu_{it} y_{it} + y_{it}^2\right). \end{aligned} \tag{11}$$

For $y_{it} > 0$,

$$\frac{\partial \phi_{it}}{\partial y_{it}} = 2\psi_i(\mu_{it} + y_{it}) > 0, \qquad \frac{\partial \phi_{it}}{\partial \mu_{it}} = 2\psi_i y_{it} > 0. \tag{12}$$

Greater Y2K demand facing the firm raises this premium. Holding y_{it} and ψ_i fixed, greater ordinary demand raises it as well. The cross term $2\mu_{it} y_{it}$ is the pooled-capacity channel: the same Y2K increment creates a larger upgrading payoff when layered onto a stronger ordinary-demand base because both demand components use the same delivery organization. This comparative static concerns the latent ordinary-demand index μ_{it}, holding y_{it} and ψ_i fixed. It is not an unconditional prediction by observed employment, which also reflects the firm's

organizational state and operating-cost parameters.

The flow premium shows what capability already in place earns from Y2K in one period. It does not yet show whether the firm should pay K_i at $t = 0$. Upgrading is an optimal-timing problem, not a test of whether capability has a positive gross return. By waiting, the firm pays K_i later and reduces the cost's present value. If ordinary demand is expected to grow, the early flow returns $\psi_i \mu_{it}^2$ may also be small relative to later returns; for some firms, the discounted return never covers the cost. With constant ordinary demand, by contrast, the firm either upgrades immediately or never, so finite later upgrading requires ordinary capability returns to rise sufficiently over time.

Following the dynamic sunk-entry tradition (Baldwin and Krugman 1989; Roberts and Tybout 1997; Das et al. 2007), the body treats the remediation window as one investment-decision phase. An initially unupgraded firm either upgrades at $t = 0$ or remains at $q = 0$ throughout the window and then upgrades at its best later date or never. Let Ω_i denote the present-value advantage, when Y2K demand is absent, of that best delayed-or-never plan over upgrading at $t = 0$. A positive Ω_i is the value of waiting. It incorporates the sunk cost and all ordinary-market capability returns, including returns after the remediation window.

Y2K adds a finite return to having the upgraded system in place during the window. Let $\lambda \in (0, 1)$ be the discount factor and define

$$\begin{aligned} \mathcal{R}_i &\equiv \sum_{t \in \mathcal{W}} \lambda^t \left[(\mu_{it} + y_{it})^2 - \mu_{it}^2 \right] = \sum_{t \in \mathcal{W}} \lambda^t \left(2\mu_{it} y_{it} + y_{it}^2 \right), \\ \Gamma_i &\equiv \sum_{t \in \mathcal{W}} \lambda^t \phi_{it} = \psi_i \mathcal{R}_i. \end{aligned} \tag{13}$$

The window-demand index $\mathcal{R}_i$ is the discounted increase in squared project demand due to Y2K. It is larger when the firm faces more remediation demand or when that demand is layered onto a larger ordinary-demand base. Holding the squared-demand increment fixed across dates, it is also larger when the increment arrives earlier. It is not a measure of total Y2K profit. The premium Γ_i prices this demand opportunity at the firm's scale–capability complementarity: $\mathcal{R}_i$ describes the opportunity, while ψ_i describes the firm. In the project-load baseline, Y2K changes $\mathcal{R}_i$, not ψ_i.

The operating return earned at $q = 0$ under Y2K is common to upgrading now and waiting through the window, so it cancels from their comparison. Y2K therefore changes the advantage of upgrading at $t = 0$ by Γ_i, and the net advantage is $\Gamma_i - \Omega_i$. I resolve indifference

in favor of earlier upgrading. The shock then changes the firm's choice exactly when

$$0 < \Omega_i \leq \Gamma_i. \tag{14}$$

The three cases are immediate. A firm with $\Omega_i \leq 0$ upgrades at $t = 0$ even when Y2K demand is absent. Y2K induces early upgrading when $0 < \Omega_i \leq \Gamma_i$. A firm with $\Omega_i > \Gamma_i$ continues to wait. An unmatched firm cannot be induced through this baseline because $L_i = 0$ implies $\mathcal{R}_i = \Gamma_i = 0$. Holding Ω_i and ψ_i fixed, the comparative statics in equation (12) imply that greater remediation demand, or greater ordinary demand during remediation, raises Γ_i.

Appendix C.3 writes the two branches formally, proves the cutoff, and gives the now-or-never special case. It also allows unrestricted upgrading dates by optimizing the Y2K and no-Y2K branches separately, then adds implementation and certification lags.

What happens after the deadline is distinct from why the firm upgraded. Once $t > T$, $y_{it} = 0$, so the direct remediation component of scale disappears. If Y2K induced the firm to enter $q = 1$ while its no-Y2K state remains $q = 0$, ordinary-market scale differs by

$$x^*_{it}(1; \mu_{it}) - x^*_{it}(0; \mu_{it}) = \mu_{it}\left(\frac{1}{c_{i1}} - \frac{1}{c_{i0}}\right) = 2\psi_i\mu_{it} > 0 \qquad \text{for } \mu_{it} > 0. \tag{15}$$

The upgraded system is not left with idle organizational capacity. At a common ordinary-demand index, optimal attention per operating unit is $z^*_i(\mu) = \theta_i/\mu$ in both states, while optimal scale is proportional to capacity. Hence $x^*_{it}(1; \mu)/x^*_{it}(0; \mu) = A_{i1}/A_{i0}$: under fixed parameters, the capability gap produces a constant proportional scale difference and a level difference that grows with ordinary demand. This is a firm-counterfactual statement, not a decomposition of the observed employment coefficient. The baseline holds the ordinary-demand path fixed across the Y2K and no-Y2K states. If Y2K instead raises later demand through a durable buyer relationship or foreign-market presence, that is the separate participation mechanism introduced below.

Because the quadratic form makes the capability return proportional to ψ_i, it is possible to characterize the organizational types whose upgrading date Y2K changes. Let $\underline{\psi}^Y_i$ be the minimum complementarity at which upgrading at $t = 0$ weakly dominates every plan that delays upgrading until after the window or never upgrades when Y2K demand is present. Let $\overline{\psi}^N_i$ be the corresponding minimum when Y2K demand is absent. The exact max-over-date expressions are in Appendix C.3.

The two thresholds expose both margins of the response. A firm with $\psi_i \geq \overline{\psi}^N_i$ finds early upgrading worthwhile on ordinary demand alone; it upgrades at $t = 0$ with or without

Y2K. A firm with $\psi_i < \underline{\psi}_i^Y$ does not find early upgrading worthwhile even after adding the Y2K premium. Only between those cutoffs does the temporary opportunity change the decision.

Proposition 1 (Selection on complementarity) *For a matched firm with $\mathcal{R}_i > 0$, and holding the upgrading cost and ordinary-demand path fixed, Y2K induces upgrading at $t = 0$ if and only if*

$$\underline{\psi}_i^Y \leq \psi_i < \overline{\psi}_i^N. \tag{16}$$

A larger $\mathcal{R}_i$ lowers $\underline{\psi}_i^Y$ but leaves $\overline{\psi}_i^N$ unchanged, so the induced interval expands toward firms with weaker complementarity. If $L_i = 0$, then $\mathcal{R}_i = 0$, the two thresholds coincide, and the interval is empty.

The now-or-never benchmark makes the interval especially transparent. Let $\mathcal{M}_i \equiv \sum_{t=0}^{\infty} \lambda^t \mu_{it}^2 < \infty$ summarize the discounted ordinary demand over which capability would pay. Then Y2K induces early upgrading precisely when

$$\frac{K_i}{\mathcal{M}_i + \mathcal{R}_i} \leq \psi_i < \frac{K_i}{\mathcal{M}_i}. \tag{17}$$

The general proposition does not require the alternative to be never upgrading: Y2K may advance an upgrade that would otherwise occur later or induce one that is absent from the no-Y2K path. The known deadline does not eliminate the mechanism. The finite window return need cover the value of waiting Ω_i, not the entire sunk cost, because capability also earns returns in ordinary business after remediation ends. Comparison firms may therefore upgrade on their own schedules as the ordinary offshore market grows.

Corollary 1 (The obsolete-match reversal) *Fix ordinary-demand paths, cost parameters, and upgrading policies when Y2K demand is absent. At any horizon $h > T$ with $\mu_{ih} > 0$, L_i can predict organizational capability and the scale that capability supports because it selected the firm into the temporary Y2K demand shock, even though the matched technology no longer has direct value. For an induced upgrader, the firm-level difference between its Y2K and no-Y2K states remains at h precisely when its no-Y2K upgrading date lies beyond h, including when it would never upgrade without Y2K.*

Durability within a firm whose upgrade Y2K induced does not make the matched–comparison differential mechanically permanent. Comparison firms may later reach their no-Y2K upgrading dates. Persistence through the sample horizon is therefore an empirical finding, not a tautology of the binary state.

4.3 Empirical implications and interpretation boundary

The model points to three empirical implications for the timing and persistence of firm scale and organizational upgrading. Employment is the closest proxy for operating scale x_{it}, while certification proxies latent organizational capability q_{it}. Certification is an outcome of the shock, not an identified mediator of employment, revenue, or survival. Predetermined match L_i and the path S_t correspond to the pre-shock legacy match and the dated remediation window. Firm-level Y2K orders, ordinary demand, upgrading costs, and the technology parameters are not separately observed. The empirical analysis therefore tests reduced-form timing, persistence, dose, and heterogeneity implications rather than estimating the structural parameters, upgrading thresholds, or a decomposition of the employment response. Section 5 gives the full measurement map.

First, at a fixed organizational state, a matched firm expands while remediation demand is present, and this direct increase in scale ends with the deadline. Observed employment is the total response, however, because the same shock may also change organizational capability. Second, Y2K should advance upgrading among firms that would not yet upgrade on ordinary demand alone but are close enough to the upgrading threshold that the temporary premium changes their decision. The model therefore does not imply that every matched firm upgrades or that unmatched firms never do. Holding other firm characteristics fixed, a stronger Y2K opportunity expands the induced-upgrading set toward firms with weaker scale–capability complementarity. The multi-tier extension in Appendix C.4 also allows a stronger opportunity to increase the depth of upgrading. Comparison firms can later upgrade as ordinary demand grows. Third, after the deadline, remediation work itself can no longer support higher scale. If Y2K induced an organizational upgrade, however, the upgraded firm can coordinate a larger ordinary business than it could have without the upgrade. Scale and certification differentials can therefore persist and can later narrow as comparison firms upgrade. Sections 6.2 and 7 examine these timing, persistence, dose, and heterogeneity patterns.

These implications do not make the model's latent firm types observable. Initial employment combines ordinary demand, the starting organizational state, and operating-cost parameters; it is not a direct measure of scale–capability complementarity. Size heterogeneity can therefore show whether the response varies with initial scale, but it cannot classify firms into the intermediate-complementarity interval in the proposition.

Beyond the multi-tier formulation, Appendix C develops three extensions that limit how the empirical patterns can be interpreted. First, the access extension separates technical

match from prior foreign reach. This differs from interventions that assign export orders or subsidize fixed market-access costs (Atkin et al. 2017a; Ali et al. 2025). The extension in Appendix C.5 lets prior reach and deadline urgency affect the amount of work a matched firm receives. The data do not observe firm-level orders or buyer search. Section 7 tests whether observed responses vary with pre-shock reach. A flat profile is bounded evidence, not proof of equal access or compressed search frictions.

Second, buyer demand for quality or assurance can raise the return to upgrading even without the baseline cost complementarity (Verhoogen 2008; Verhoogen 2023). The procurement-gate extension in Appendix C.7 formalizes this possibility, but the paper does not establish that Y2K contracts imposed such a gate. Buyer interaction can also transmit production knowledge (Atkin et al. 2017a); financing, standards, or learning could change upgrading costs or organizational technology rather than only the demand path. Those channels require evidence or states not present in the baseline.

Third, organizational capability is not the only persistent state. Appendix C.6 lets $e_{it} \in \{0, 1\}$ represent a buyer relationship, approved-vendor status, reputation, or foreign presence. A temporary shock can make firms cross a sunk participation threshold and leave later activity higher (Baldwin and Krugman 1989; Roberts and Tybout 1997; Das et al. 2007). Persistent activity is therefore consistent with q, e, or both. Evidence that certification moved establishes that a capability-related outcome changed; it does not quantify how much of persistent scale runs through capability, rule out an unobserved participation asset, or establish mediation.

5 Research Design and Identification

Section 4 yields three empirical implications. First, firms whose pre-shock technology stacks matched Y2K remediation requirements expand while remediation demand is active. Second, the temporary opportunity can accelerate upgrading among exposed firms for which ordinary demand alone would not yet justify the investment. Third, after remediation demand ends, the exposed–comparison difference in scale can persist if the shock durably changes organizational capability or foreign-market participation.

I test these implications using time variation created by the Y2K shock and cross-sectional variation created by the pre-shock legacy-technology match. Although the Y2K problem was known before 1995, corporate remediation activity rose sharply only in 1998–99 and ended at the January 2000 deadline. I estimate an event study that traces the exposed–

comparison difference directory wave by directory wave, and a split-window difference-in-differences that averages the same difference within the 1998–99 remediation window and the 2001–03 post-deadline window.

The binary comparison asks whether exposed firms responded. The model further implies that the response should rise with the remediation demand generated in a firm's own client sectors. I therefore complement the event study with a Bartik-style demand dose that combines predetermined legacy match, firms' pre-shock client-sector mix, and measured sectoral Y2K spending; it provides supporting evidence on exposure intensity, and Section 5.3 develops it together with the validation evidence. Section 7 then conditions directly on predetermined firm quality, prior market position, and communications capacity and asks whether the response is confined to firms with broad prior reach or large initial scale. The modern-stack placebo belongs to the validation and dot-com tests rather than to that conditioning exercise.

Table 3 links the model variables to the measures used in the empirical analysis. The predetermined technology stack and remediation window operationalize L_i and S_t; employment is the closest measure of x_{it}, certification is an observed proxy for q_{it}, and client and foreign-market outcomes speak to the participation extension. Firm-level Y2K orders are unobserved, and the structural parameters are not separately identified, so the analysis tests reduced-form timing, dose, and persistence rather than estimating the model.

I combine nine sources to measure exposure and outcomes and to examine competing explanations. The seven surviving NASSCOM directories (1992, 1995, 1998, 1999, 2001, 2002, and 2003) report technology stacks, employment, revenue and exports, certifications, clients, and foreign-market reach; they supply the primary firm-level exposure and outcomes. From SEC–EDGAR 10-K and 10-Q filings, I extract corporate Y2K spending and readiness disclosures. Together with Compustat denominators and OMB/GAO public-spending benchmarks, these disclosures measure the timing and client-sector intensity of remediation demand. Prowess, an audited-accounts database covering listed and larger unlisted Indian firms across all industries, supports the three complementary designs of Section 5.4: the matched design carries the NASSCOM exposure to audited outcomes for the crosswalked firms, the full-Prowess design compares IT firms with the rest of the economy, and the regional design builds a within-IT geographic exposure from Prowess alone. Offshore delivery ran on dedicated data links, supplied in 1990s India primarily through the Software Technology Parks of India (STPI), so infrastructure could gate the response; STPI annual reports record where and when that infrastructure operated, and BEA industry accounts measure

Table 3: Mapping model variables to the empirical analysis

Model variable	Variable in the empirical analysis	Measurement and interpretation
Predetermined match L_i	Predetermined legacy match L_i: L_i^{bin}, L_i^{count}, coherent-layer count, D_i^{Y2K}, and the standardized coherence dose	All are constructed from stacks reported before the demand surge. The binary match and category count use 1995; the coherence measures pool 1992 and 1995. They capture a match to Y2K requirements, not upgrading or technical quality in general.
Common Y2K demand S_t	$\mathbf{1}[t = u]$, Shock_t, and Post_t	Event time traces the full path. Shock_t marks the 1998–99 positive-demand window; Post_t marks 2001–03, when direct Y2K demand has ended and $S_t = 0$. SEC–EDGAR disclosures and public-spending benchmarks discipline the timing.
Incremental Y2K demand $y_{it} = L_i S_t$	$L_i \times \mathbf{1}[t = u]$, $L_i \times \text{Shock}_t$, $L_i \times \text{Post}_t$; reported 1998–99 engagement; supporting dose $z_i = L_i \sum_s \omega_{is} g_s$	Firm-level orders are unobserved. The event-time and shock interactions grade exposure while remediation demand is active. The post interaction measures the total differential after $y_{it} = 0$. Reported engagement validates the requirement match, while z_i adds supporting evidence on demand intensity. Neither supplies independent identification.
Operating scale x_{it}	Employment outcome Y_{it}	NASSCOM employment is the closest observed measure of operating scale. A post-deadline coefficient is the total persistent response, not a structural estimate of upgrading.
Organizational capability q_{it}	Certification outcome Y_{it}	The ladder records none (0), ISO 9000 (1), and CMM Levels 3–5 (2–4); ISO and CMM are also examined separately. Certification is an observed proxy and an outcome, not an identified mediator.
Participation state e_{it} in the extension	Client and foreign-market outcomes Y_{it}	Reported clients, foreign-market reach, foreign-client share, and foreign presence can reveal a persistent participation margin, but none is a one-to-one measure of the latent state.
Ordinary demand and structural objects μ_{it}, m_{it}, K_i, and ψ_i	No direct empirical variable	These objects are not separately observed or structurally estimated. The analysis instead tests reduced-form timing, persistence, dose, and heterogeneity implications.

Notes. Section 4 defines the model objects. Section 2 and Appendix A document source construction and coverage. Revenue, exports, capital, audited financial outcomes, and legal survival are broader responses rather than proxies for a single model variable. Legal survival is suggestive extensive-margin evidence because the model has no exit margin.

general United States IT spending, the dot-com alternative to Y2K-specific demand. District IT activity from the Economic Census, organized on the SHRUG data platform, measures local IT exposure for the geographic school-entry test. SEMIS records school foundings; NSS household rounds provide exploratory schooling outcomes not used in the preferred test.

5.1 Predetermined exposure to Y2K demand

I measure firm-level exposure to the Y2K demand surge using legacy technologies reported before the surge. The baseline binary incidence measure L_i^{bin} equals one when a firm listed at least one of thirteen qualifying COBOL- or mainframe-family technology categories in 1995. Among the 235 firms in the 1995 directory cohort, 135 meet this definition. This fixed incumbent cohort is followed as an unbalanced panel; inclusion does not require appearance in every later wave or continued legal existence through 2003.

I complement the binary measure with four measures of the breadth and coherence of the predetermined technical match. First, L_i^{count} counts the qualifying categories in the 1995 stack. Second, the coherent-layer count records distinct core production layers without requiring them to belong to the same ecosystem, separating production depth from raw category breadth. Third, the coherent-stack indicator D_i^{Y2K}—the coherence gate—equals one if either of the two headline IBM ecosystems satisfies its minimum production rule. The IBM-mainframe rule requires at least three of five layers, including a legacy language, an operating-system or hardware anchor, and a database or transaction/batch layer; the AS/400 rule requires at least three of four layers, including a language and an operating-system or hardware anchor. Fourth, the coherence dose $\max\{C_{i,\text{MF}}/5, C_{i,\text{AS400}}/4\}$ records the completeness of the firm's best-matched ecosystem, where each C counts the production layers filled in that ecosystem. The dose enters the regressions in standardized form.

The three coherence measures pool the 1992 and 1995 listings to reduce false negatives from intermittent stack reporting. Ten of the 235 firms in the 1995 incumbent cohort have no readable stack in either year and are treated as missing in the depth specifications, leaving a 225-firm common sample; 55 satisfy the coherent-stack rule. These specifications additionally account for pre-shock reporting verbosity, frontier-stack breadth, and stack-field observability. Appendix A.7 gives the token classification, single-year timing checks, and alternative constructions.

5.2 Specification, estimand, and identifying assumption

I first estimate the full path to distinguish pre-existing differential change from responses during remediation and after the deadline. The cross-sectional logic follows Rajan and Zingales (1998): a predetermined requirement match grades exposure to an external shock. The use of a temporary, dated shock to trace persistence is closer to Hanlon (2017). For

outcome Y_{it}, I estimate

$$Y_{it} = \alpha_i + \delta_{c(i)t} + \sum_{u \neq 1995} \beta_u \Big(L_i \cdot \mathbf{1}[t = u] \Big) + \gamma' M_{it} + \varepsilon_{it}, \tag{18}$$

where L_i is a predetermined legacy-match measure, 1995 is the omitted wave, and u indexes directory waves. Firm fixed effects α_i absorb time-invariant firm differences. Entry-cohort-by-year fixed effects $\delta_{c(i)t}$ allow firms first observed in different directory waves to follow different common paths. The vector M_{it} records merger and acquisition events, and standard errors are clustered by firm.

For tables, I compress the path into remediation and post-deadline windows:

$$Y_{it} = \alpha_i + \delta_{c(i)t} + \beta_1 \Big(L_i \cdot \text{Shock}_t \Big) + \beta_2 \Big(L_i \cdot \text{Post}_t \Big) + \gamma' M_{it} + \varepsilon_{it}, \tag{19}$$

where Shock_t marks 1998–99, the empirical counterpart of positive S_t, and Post_t marks 2001–03. The pooled pre-period combines 1992 and 1995; specifications normalized to 1995 and exposure measures built only from information available in 1992 or 1995 are reported as sensitivities.

For the baseline binary measure, the estimand compares changes over time between matched and unmatched firms within the 1995 NASSCOM incumbent cohort. Under parallel trends conditional on the fixed effects and controls, each β_u is the change from 1995 among firms whose pre-shock stack matched Y2K requirements minus the corresponding change among firms whose stack did not, and in the pooled specification β_1 and β_2 compare the remediation and post-deadline differentials with the pooled 1992–1995 pre-period. The coefficients measure the consequence of having the requirement match when the demand surge arrived. They are not intention-to-treat effects in the experimental sense because no treatment, eligibility rule, or encouragement was assigned. Exposure is the observed pre-existing stack match, and reported Y2K engagement enters the analysis only as a validation outcome. For employment, β_2 is the total post-deadline differential.

Estimates in logs and estimates of proportional changes in conditional means can disagree when outcomes are skewed and responses are heterogeneous, so functional form is a substantive choice. For employment, capital, software exports, and revenue, I therefore pair log-linear fixed-effects estimates with PPML estimates in levels on the same outcome-specific observations; holding the sample fixed, the comparison isolates functional form and weighting, not the inclusion of zeros. The reported-client and foreign-market counts are small nonnegative integers, for which PPML is the natural count specification. Indicator

outcomes and the 0–4 certification ladder remain linear. Under PPML, $\exp(\beta) - 1$ is the proportional difference in conditional means.

Identifying assumption and limitations. The causal interpretation requires that, conditional on firm and cohort-by-year fixed effects and the included controls, firms with and without the pre-shock match would not otherwise have experienced differential changes around the remediation window and post-deadline years. The assumption concerns untreated paths, not equality in baseline levels. It is demanding because the legacy match correlates with firm size, age, and prior foreign orientation, and fixed effects alone do not rule out differential trends associated with those characteristics. Section 7 therefore reports both the full legacy-match differential and the residual after allowing a broad set of predetermined characteristics to generate their own remediation and post-deadline paths.

Four limitations discipline the interpretation. First, Y2K was known before 1995, so a legacy stack reported in 1995 could in principle reflect early positioning for remediation work; the match measures never incorporate subsequent engagement reports, and exposure built from the 1992 directory alone is the stricter timing check. Second, NASSCOM supplies only one earlier directory wave, so the event study can disclose the 1992 differential but cannot establish a long pre-trend; the matched-Prowess design supplies the annual audited pre-period the directory lacks and asks whether the same implications hold there. Third, the measures that pool the 1992 and 1995 technology listings partly use information reported after 1992, so their 1992 outcome coefficients mix reporting timing; the single-year constructions address that timing directly. Fourth, technical matches can be acquired, workers can move, and work can be subcontracted across firms, so part of the shock can reach comparison firms; such spillovers push the exposed–comparison differential toward zero rather than creating one, and entry after 1995 lies outside the NASSCOM design, which the full and regional Prowess designs reach under their own assumptions.

The principal competing explanation is the general late-1990s IT boom, which could favor the same firms on a similar path. The dated window and deadline sharpen the comparison but do not eliminate this alternative; Section 8 evaluates it. The post-deadline coefficient is a persistent Y2K response only under the maintained parallel-trends assumption. A persistent post-deadline response could reflect organizational capability or a durable participation margin. Certification can reveal whether a capability-related outcome moved, but cannot identify that outcome as the mediator of scale.

5.3 Validating the legacy match and the Bartik demand dose

The difference-in-differences design grades firms by their technical match, but it does not by itself establish that the stack measures predict subsequent Y2K work or that responses line up with demand in a firm's client sectors. I assess both before turning to the employment and upgrading results. First, Table 4 regresses each firm's 1998–99 directory report of offering Y2K work—recorded either in the dedicated field or its profile narrative—on each predetermined exposure measure. A firm with $L_i^{\text{bin}} = 1$ is 0.327 more likely to report engagement (SE 0.076), against a sample mean of 0.47. The category count and coherence dose are at least as strongly related to engagement; the coherence gate is positive but somewhat weaker (0.294, SE 0.085; $t = 3.45$). A frontier-stack placebo, constructed with the same coherence logic on modern technologies, predicts engagement only weakly on its own and not once the legacy match is held fixed. The validation therefore links reported engagement to the legacy requirement match rather than to technical sophistication in general.

Second, I introduce client-side demand intensity. Define

$$z_i = L_i \cdot \underbrace{\sum_s \omega_{is} g_s}_{B_i}, \tag{20}$$

where ω_{is} is firm i's pre-shock client-sector weight and g_s is the sector's EDGAR remediation-spending intensity. The dose z_i asks whether responses are larger when the technical match coincided with heavier measured demand in a firm's pre-shock client sectors, an exposure used notably in the matched-Prowess design. It varies across firms through both predetermined legacy match and pre-shock client composition; it is distinct from predetermined district IT presence in the regional-Prowess design. The shares-only component B_i predicts engagement weakly ($t = 1.15$), so much of the interaction's relevance inherits the legacy match. With only nine sector groups, z_i refines the exposure comparison but does not supply an independent source of identification.

The primary estimand is the response to predetermined exposure, not the effect of reported engagement. The predetermined legacy match may affect unreported projects, client acquisition, hiring, or credibility as well as the directory engagement field. The exclusion restriction required to interpret a coefficient obtained by instrumenting reported engagement with the legacy match is therefore implausible. Such a coefficient rescales the reduced form; it does not identify a causal effect per engaged firm. Appendix Table B.4 reports the sample-specific first stages, the corresponding two-stage coefficients, and Anderson–Rubin sets as a

noncausal scaling sensitivity.

5.4 Complementary designs, selection, and inference

The complementary Prowess designs test the scale and persistence implications on audited financial data, under assumptions and coverage that differ from the NASSCOM design's. The matched design carries the same NASSCOM exposure onto annual audited outcomes and supplies longer leads. It is closest to the core estimand, but its sample is small and selected toward larger firms, and Prowess observes none of the capability and participation margins—certification, reported clients, and foreign-market reach remain directory-native. The full-Prowess design compares IT with non-IT firms and traces the sector-wide path that the within-IT comparison differences out; differential pre-trends prevent it from becoming a second identification pillar. The regional design uses predetermined local IT presence within IT and includes entrants without a 1995 NASSCOM exposure; few effective geographic cells and hub confounding limit it to a consistency check.

The gates exercises use the NASSCOM design to condition on observed firm quality, market position, and communications capacity, test heterogeneity by prior reach and size,

Table 4: Predetermined exposure and subsequent reported Y2K engagement

Dependent variable: reported Y2K engagement, 1998–99				
Exposure measure	Estimate	Robust s.e.	t	Firms
Panel A. Requirement-match validation and placebo				
Any legacy technology, L_i^{bin}	0.327***	(0.076)	4.33	161
Legacy-category count	0.074***	(0.016)	4.69	161
Coherent legacy stack, D_i^{Y2K}	0.294***	(0.085)	3.45	158
Coherence dose (per s.d.)	0.176***	(0.035)	5.09	158
Frontier-stack strength (placebo)	0.333*	(0.178)	1.88	158
Panel B. Demand-dose decomposition				
Legacy match × sector demand, z_i	52.785***	(11.230)	4.70	161
Sector demand alone, B_i	23.115	(20.145)	1.15	161
Sector demand among legacy-matched firms	44.348*	(25.983)	1.71	101

Notes. Each row is a separate cross-sectional linear probability model with HC1 standard errors. Engagement equals one if either the dedicated field or profile narrative reports Y2K work in 1998 or 1999; an explicit no and absence from both waves are coded zero. Its mean is 0.47 in the 161-firm client-composition sample. Sector demand is $B_i = \sum_s \omega_{is} g_s$ over nine sectors, and the dose is $z_i = \text{Cap}_i B_i$. Scales differ across rows, so the robust t-statistics are the comparable evidence measures. Each row's heteroskedasticity-robust F equals its squared t-statistic. Conditional on the legacy match, the frontier-placebo t-statistic is 0.76 (legacy $t = 3.83$). These are construct-validation and supporting-dose diagnostics, not an independent IV design. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

and document the limited timing support for STPI mediation. Local formal education supply is a boundary on contemporaneous confounding rather than a test of the firm model. Because it asks a distinct question using a district school-entry comparison, its measurement and specification are introduced in Appendix B.10.

Observation, selection, and inference. The estimation universe contains the 235 firms listed in the 1995 directory, tracked across the seven directory waves from 1992 to 2003. The panel is unbalanced, and directory reappearance, outcome disclosure, and legal survival are distinct processes. Absence from a later directory is not legal exit. The principal coefficients are therefore incumbent- and observation-conditional, with the post-deadline estimates identified by firms observed on both sides of the deadline.

Balanced-panel estimates, inverse-probability weights, and Lee-type bounds assess sensitivity to observed composition under their stated assumptions; none eliminates selection on unobservables. Section 8 reports the full selection battery and keeps baseline-comparability reweighting separate from reappearance weighting.

For the core NASSCOM design, firm-clustered standard errors are the primary inference procedure. Permutation tests that reassign the predetermined legacy match are reported as sensitivity exercises under explicit exchangeability restrictions, including cohort- and size-stratified variants. The match was not randomly assigned, so these tests are not design-based randomization inference and do not replace the clustered standard errors. The regional-Prowess and education sections specify their own inference procedures.

6 Growth, Persistence, and the Upgrading Margin

This section tests whether the temporary foreign-demand shock left exposed firms persistently larger after demand receded and, if so, through which margins. I first use the NASSCOM self-reported directory panel to estimate the scale response and its persistence, and corroborate the input and foreign-revenue margins with matched Prowess audited statements.[4] I then examine certification, clients, foreign markets, and service scope in the directory. I close with two broader comparisons using the universe of firms reported in Prowess: IT versus other industries, and firms in regions with more versus less pre-shock IT activity.

I find that exposed firms became persistently larger after the deadline. The scale effect is largest and most consistent for employment; registered capital also rises, although its

[4]Prowess coverage is concentrated among listed and larger unlisted firms.

estimate is sensitive to functional form. The employment result is stable across alternative measures of firms' pre-shock technological match to Y2K demand. Certification, including CMM process maturity, rises and persists, whereas reported clients, foreign markets, and service scope do not show comparable robust differentials. Among firms matched across the two sources, audited compensation and fixed assets move in the same direction, and higher Y2K demand intensity predicts higher foreign revenue. In the full Prowess sample, IT export earnings rise around the deadline relative to other industries and remain higher afterward, although differential pre-trends preclude a causal interpretation. The regional comparison instead stays within IT, relating outcomes to district-level IT presence measured through fiscal 1996 rather than using non-IT industries as the comparison group; district-based exposure also covers later entrants. Across roughly sixteen effective district cells, post-deadline export earnings and the rough export-activity proxy are higher under a permutation reference test, while fixed assets accelerate; post-deadline sales and total-asset levels remain imprecise.

Overall, these findings are consistent with the model's central sequence. The model predicts that, holding organizational capability fixed, the scale response should end when remediation demand expires. If the shock induces upgrading, however, the resulting capability can support a larger ordinary business after the deadline. I find evidence consistent with this prediction: the legacy match predicts persistent employment and a lasting rise in certification, including CMM, after the deadline, when the match no longer has direct remediation value. While the evidence does not identify certification as the mediator of employment, the finding shows persistent scale alongside a persistent capability-related response—the joint pattern the model predicts when a temporary demand shock induces durable upgrading.

6.1 The scale response and its persistence

Table 5 reports estimates of the split-window difference-in-differences specification in equation (19) for employment, registered capital, software exports, and revenue. The coefficient β_1 compares changes during the 1998–99 remediation period between firms with and without the pre-shock legacy match, while β_2 makes the same comparison after the deadline. The contrast $\beta_2 - \beta_1$ asks whether this relative gain persists or widens after the deadline.

In the log-linear estimates (Panel A), both inputs exhibit persistence, although their timing differs. The employment differential is modest and statistically indistinguishable from zero during remediation, at 0.172 log points, but rises to 0.425 log points after the deadline. This corresponds to approximately 53%, or roughly 50 employees at the 1995

Table 5: The legacy match and the timing of firm scale-up

	Inputs		Outputs	
	(1) Employment	(2) Capital	(3) Software exports	(4) Revenue
Panel A. DiD Estimates: OLS on log outcomes				
Legacy match × Remediation period (1998–99, β_1)	0.172	0.560**	0.102	−0.075
	(0.133)	(0.258)	(0.287)	(0.213)
Legacy match × Post-deadline period (2001–03, β_2)	0.425**	0.914***	0.246	0.060
	(0.172)	(0.301)	(0.320)	(0.262)
Difference ($\beta_2 - \beta_1$): post-deadline minus remediation	0.252**	0.354*	0.145	0.136
	(0.116)	(0.211)	(0.220)	(0.184)
Firm fixed effects	✓	✓	✓	✓
Cohort × year fixed effects	✓	✓	✓	✓
M&A event controls	✓	✓	✓	✓
Comparison-group mean, 1995	95	$19.6m	$0.88m	$1.12m
Observations (firm-years; firms)	953 (199)	738 (161)	716 (153)	814 (173)
Panel B. DiD Estimates: PPML on outcomes in levels				
Legacy match × Post-deadline period (2001–03, β_2)	0.581***	0.362	0.079	0.128
	(0.218)	(0.246)	(0.252)	(0.274)
Observations (firm-years; firms)	953 (199)	738 (161)	716 (153)	814 (173)

Notes. The legacy match is L_i^{bin}, an indicator for whether firm i reported any qualifying legacy technology before the shock. Each Panel A column reports a firm-level difference-in-differences regression. The omitted pre-period is 1992–95. The first two coefficient rows are the legacy-match differentials during the remediation window (1998–99) and after the January 1, 2000 deadline (2001–03); the third row is the linear contrast between them from the same regression. Panel A estimates OLS on the log outcome. Panel B estimates the same model by PPML on the outcome in levels, retains both period interactions, and reports the post-deadline coefficient β_2. PPML uses a log link; its coefficients are semi-elasticities, so $\exp(\beta) - 1$ gives the proportional conditional-mean difference for a binary exposure. Outcome availability differs across columns. Comparison-group means are measured in employees or US$ millions. All specifications include firm fixed effects, cohort-by-year fixed effects, and M&A event controls. Standard errors clustered by firm are in parentheses. The PPML (Panel B) sample is constructed to match the OLS (Panel A) sample to isolate estimator and functional-form sensitivity. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

comparison-group mean of 95; the post-minus-remediation contrast is itself significant, and a permutation test that reassigns the legacy-match indicator gives $p = 0.012$ for the post-deadline estimate. In contrast, registered capital adjusts earlier. Its differential is already 0.560 log points during remediation and rises to 0.914 after the deadline, corresponding to approximately 149% higher registered capital. Thus, the rise in registered capital begins while remediation work is under way, whereas employment expansion is concentrated after the deadline. Neither differential reverses when the temporary demand ends.

Panel B re-estimates each outcome by PPML in levels, using the same firm-years as

Panel A for that outcome. Differences across panels therefore reflect estimator and functional form, not sample composition. Employment remains positive and statistically significant, with a post-deadline coefficient of 0.581, implying approximately a 79% difference in the conditional mean. The capital coefficient falls to 0.362 and becomes imprecise. Across estimators, employment therefore carries the persistent scale result. Capital indicates earlier balance-sheet expansion but does not have the same functional-form robustness. The employment response is also tilted toward reported software and engineering headcount: the post-deadline coefficient is 0.374, against 0.070 for other non-technical workers, with the same ordering under PPML (Appendix Table B.5).

The estimates for the directory's output measures, software exports and revenue, are directionally as expected but not statistically significant; because both are disclosed by selected and changing subsets of firms, they also mix output response with reporting behavior. The matched Prowess statements below provide a separate audited check on the foreign-revenue margin.

The employment pattern is robust across alternative measures of the pre-shock technology match to Y2K demand. Appendix Table B.2 reports five measures defined in Section 5.1: the baseline legacy indicator, the coherent-stack gate, the legacy-category count, the standardized stack-coherence dose, and the coherent-layer count. Employment is positive under all five measures and both estimators; the three continuous depth measures are statistically significant in both, while the binary measures trade precision across functional forms. Because the measures have different units and enter separate regressions, the exercise establishes robustness to alternative measurement rather than a ranking by stack depth. Capital remains positive but generally imprecise. Directory exports and revenue are insignificant under OLS and become positive for the depth measures only under PPML, reinforcing the conclusion that the output evidence is sensitive to functional form and remains suggestive.

Figure 6 plots event-study estimates of the employment differential across directory waves under the two exposures the design develops: the baseline legacy indicator and, as the intensity margin, the demand dose from Section 5.3. Under both, the differential becomes positive during remediation and remains positive after the deadline. The 1992 coefficients are small and imprecise and, because exposure is measured from the 1995 stack, provide only a conditional check on differential pre-trends rather than strict placebos. Because the exposures have different units, the common result is their timing and persistence rather than their relative magnitudes. Appendix Figure B.1 reports the corresponding paths under the depth measures, the stack-coherence dose and the coherent-stack gate. Appendix Figure B.2

reports the corresponding event studies for registered capital. There, the binary path rises around 1999 and remains positive afterward, while the stack-coherence estimates are less precise after the deadline, consistent with the earlier capital adjustment in Table 5.

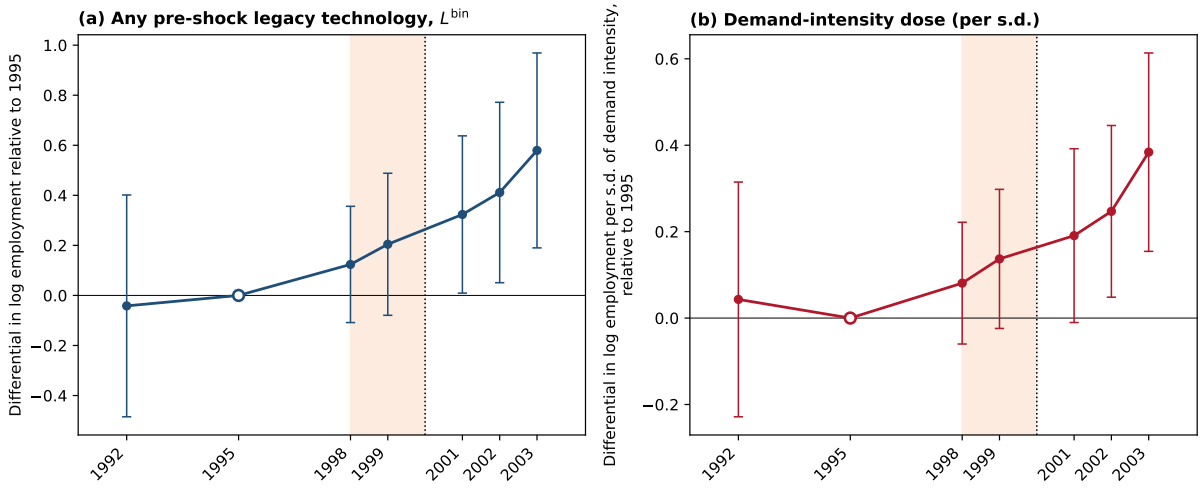


Figure 6: Employment dynamics under the baseline legacy indicator and the demand-intensity dose.

Notes: Each panel plots coefficients from a separate OLS event-study regression of log employment on the indicated exposure interacted with directory-wave indicators. Panel (a) uses L_i^{bin}, an indicator for whether firm i reported any pre-shock legacy technology ($N = 953$ firm-years). Panel (b) uses the demand-intensity dose of Section 5.3—the predetermined legacy match interacted with the firm's pre-shock client-sector Y2K-spending intensity—scaled so the plotted coefficient corresponds to a one-standard-deviation increase; it is estimated on the subsample with a defined dose ($N = 669$ firm-years). Both specifications include firm fixed effects, cohort-by-year fixed effects, and M&A event controls. The 1995 interaction is omitted (open markers). Whiskers are 95% confidence intervals based on standard errors clustered by firm. The shaded area denotes the 1998–99 Y2K-remediation period, and the dotted line marks the January 1, 2000 deadline.

Lastly, I ask whether exposed firms were also more likely to remain in business. Legal survival is measured from corporate-registry records, which yield a determinate status for 198 of the 235 cohort firms; the remaining 37 cannot be classified. Among the 198, the pre-shock legacy match is associated with an 11.7-percentage-point higher survival probability in the adjusted linear model; the corresponding probit estimate is 11.3 percentage points. The positive association is consistent with the persistent scale response but is sensitive to how the 37 unclassified firms are coded, so the survival evidence remains suggestive (Appendix Figure B.3).

Audited corroboration among matched firms. Matched audited statements provide annual corroboration for selected input and foreign-revenue patterns, although the samples are small and selected toward larger filers. The binary matched universe contains 32 linked

firms—23 with the pre-shock legacy match and nine without—and employee counts are too sparse to reproduce the directory employment result. In that comparison, audited compensation and fixed assets move in the same direction as the directory inputs. When the audited statements are organized into the same directory-wave windows, export earnings are higher during and after remediation (Appendix Table B.15; the annual and wave-grain specifications are in Online Appendix Section H). Figure 7 instead uses the continuous demand-intensity dose, defined for 19 firms and nonzero for 16. From fiscal 1993 through 1997, neither audited outcome shows a sustained differential trend. Fixed assets then rise during remediation and settle at roughly half a log point per standard deviation, while export earnings turn positive around the deadline and reach 0.82 log points in fiscal 2004. The pooled post-deadline export estimate is 0.49 log points ($p = .006$, clustered by firm). Given these small selected samples, the NASSCOM panel continues to carry the employment and upgrading results.

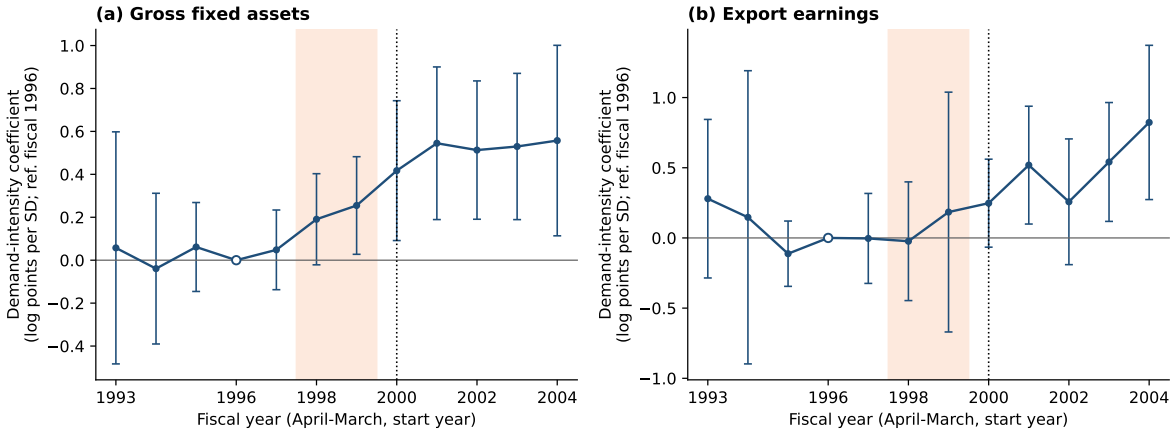


Figure 7: Audited corroboration among matched firms.

Notes: Each panel reports a separate annual event study of the log of the indicated positive audited outcome on the continuous demand-intensity exposure z, scaled per standard deviation and interacted with fiscal-year indicators, relative to omitted fiscal 1996 (open marker). Both specifications include firm and year fixed effects; whiskers are 95% confidence intervals from firm-clustered standard errors. The dose is defined for 19 matched firms and is nonzero for 16. After conditioning on a positive outcome report and removing fixed-effect singletons, panel (a) uses 19 firms, 16 with nonzero dose ($N = 190$), and panel (b) uses 17 firms, 15 with nonzero dose ($N = 165$). The display covers fiscal 1993–2004. Shading marks the fiscal 1998–99 remediation window; the dotted line marks fiscal 2000, the first fully post-deadline year because the January 2000 deadline falls inside fiscal 1999. Panel (a) reports gross fixed assets and panel (b) export earnings.

6.2 The Upgrading Fingerprint

The persistent employment differential does not reveal why exposed firms remained larger after the deadline. It could reflect a durable change in organizational capability, a buyer

relationship or another asset that helps the firm remain active abroad, or both (Melitz 2003; Roberts and Tybout 1997; Das et al. 2007; Verhoogen 2008; Atkin et al. 2017a). The directory permits a closer comparison because it records certification, clients, foreign markets, foreign presence, service scope, and technology stacks. I refer to the joint pattern across these outcomes as an upgrading fingerprint.

The certification systems. The directory records firms' reports of ISO-family certification and CMM maturity status. ISO 9000-series certification, including the software-oriented TickIT scheme, concerns a documented quality-management system assessed by an independent certification body; it does not certify product quality and is not specific to software (Arora and Asundi 1999; Arora and Athreye 2002). The Capability Maturity Model (CMM), developed by Carnegie Mellon's Software Engineering Institute, is a software-process maturity framework whose formal appraisals rate the software-process maturity of an organization or defined organizational unit on a five-level scale (Paulk et al. 1993; Masters and Bothwell 1995; Athreye 2005). Both may change internal processes or reassure foreign buyers about process quality, channels emphasized in the Indian software literature (Banerjee and Duflo 2000; Arora and Asundi 1999). Because the directory records the credential rather than how it creates value, the estimates cannot distinguish internal process change from buyer assurance. The 0–4 ladder used below is an ordinal index of reported depth: 0 for neither ISO-family certification nor CMM, 1 for ISO-family certification only, 2 for any reported CMM credential not identified as Level 4 or 5, 3 for CMM Level 4, and 4 for CMM Level 5. The ladder ranks reported credentials, but a one-rung increase need not represent the same change in cost or capability at every point. Any CMM report takes precedence over the ISO rung: a firm reporting CMM without ISO is assigned the rung corresponding to its reported CMM level, or rung 2 if the level is not identified; I do not impute an ISO credential.

Reported certification status rarely moves down the ladder. Across 780 successive observed within-firm wave transitions in the 235-firm legacy-match universe, 609 are unchanged, 162 are increases, and nine are decreases. Certification also became much more common in the full, changing directory cross-sections: the share of firms reporting an ISO-family or CMM credential rose from 12.3 percent in 1995 to 52.5 percent in 2003, while the CMM share rose from 1.3 to 19.3 percent.[5] The estimating question is whether firms with the pre-

[5]Contemporaneous policy also encouraged certification. Beginning in 1993–94, a Government of India scheme reimbursed 75 percent of admissible ISO-certification expenses for registered small-scale units (Office of the Development Commissioner (MSME), Government of India 2005). NASSCOM's 1999 directory reports

shock legacy match certified earlier or reached higher reported maturity than firms without it.

Which margin persists. Table 6 compares reported clients (column 1), the certification ladder (column 2), and the number of distinct foreign countries named in firms' reported markets or foreign-branch locations (column 3). Panel A estimates the split-window difference-in-differences specification in all three columns, so the cross-outcome comparison does not mix estimators. The 2003 directory-format change limits the two count outcomes to a 2001–02 post window, whereas certification uses 2001–03, so Panel C's cross-outcome tests do not use identical post windows. Panel B re-estimates the two count outcomes by PPML as a functional-form check; certification is omitted because its ordered ladder changes almost exclusively upward and does not provide an informative Poisson comparison. Because the outcomes use different units, their coefficients cannot be compared directly. Nor does finding a significant certification estimate and an insignificant estimate for another outcome establish that the two effects differ. Panel C therefore standardizes each outcome within its own estimation sample, stacks the three regressions, and directly tests the differences between the standardized differentials. Figure 8 plots the corresponding event studies, with columns 1–3 matching panels (a)–(c); the figure shows the annual paths, while the table summarizes the remediation and post-deadline windows.

The certification column carries the persistent differential. Certification rises by 0.319 more ladder points during remediation and by 1.066 more after the deadline for firms with the pre-shock legacy match than for firms without it. Firms without the match also certify over these years, so the estimates capture differential timing and reported depth rather than certification confined to the legacy-matched group. The gap widens from 0.58 ladder points in 1999 to 1.54 in 2003. The 1.066 coefficient is the total legacy-match differential used for the common cross-outcome comparison in this section. Section 7 conditions on the pre-shock firm bundle with which the match traveled and reports the residual beyond those observables; it does not replace the estimand used here.

In contrast, clients and foreign countries do not show a robust persistent differential. Their remediation-window estimates are positive but imprecise, and their post-deadline es-

that software exporters holding ISO 9000 or SEI CMM Level 2 or higher were eligible for Special Import Licences equal to 5 percent of f.o.b. export value (NASSCOM 1999, pp. 24–25, 40–41). Common policy timing is absorbed by the cohort-by-year effects, although differential take-up by legacy match or size could still amplify the estimated gap. The import-license rule used the same qualification for ISO 9000 and CMM Level 2 or higher, so it offered no additional policy return to the Level 4/5 deepening that contributes to the ladder response.

timates are not statistically distinguishable from zero. Because the confidence intervals include zero but also sizeable positive responses, the data do not detect a persistent client or foreign-country response; they do not rule one out.

On Panel C's standardized scale, the post-deadline certification differential is 0.81 standard deviations larger than the client differential and 0.82 standard deviations larger than the foreign-country differential. Two different size checks qualify these comparisons. The first reassigns the legacy-match indicator only among firms in the same 1995 employment tercile. Under this exercise, the certification–foreign-countries contrast remains statistically distinguishable from zero ($p = .006$), whereas the certification–clients contrast does not ($p = .150$).

The second check allows the three 1995 employment groups to follow separate annual paths. With those size-specific paths absorbed, the standardized post-deadline differentials are 0.46 for clients, 0.59 for certification, and 0.02 for foreign-country reach. The certification–clients difference is then 0.13 standard deviations, while the certification–foreign-countries difference remains 0.57 standard deviations. Size adjustment therefore makes the client comparison less robust, but it does not show that client counts failed to respond. Taken together, the results show a robust certification response without comparably robust responses in the two expansion margins.

One important caveat is that certification had already begun to diverge before remediation. Its 1992 lead is −0.338 ladder points (SE 0.122); among the 91 firms observed in both pre-period waves, 16 of the 56 firms with the pre-shock legacy match first certified by 1995, compared with none of the 35 firms without it. A linear continuation of this early divergence accounts for 67–94% of the post-deadline certification path. After subtracting that continuation, the residual annual certification coefficients remain positive and average about 0.2 within-sample standard deviations, so this exercise bears mainly on magnitude rather than sign. Figure 8 should therefore not be read as a clean break at the deadline.

The post-deadline differential is nevertheless not confined to those early adopters. Dropping them leaves a differential of +0.87 ladder points, significant at the one-percent level. With size-by-year controls, the annual post-deadline coefficients remain positive but attenuate by 30–46%, and the estimated deadline kink falls from +0.207 to +0.125 and becomes imprecise. The estimate also remains positive under alternative ladder codings and among firms uncertified in 1995 (Appendix Tables B.7 and B.10). These checks limit but do not eliminate the pre-trend concern. The evidence is therefore an earlier ISO-family divergence followed by further deadline-period widening; the decomposition below shows that CMM

Table 6: The upgrading fingerprint: which outcome of the demand response persists

	Buyer mass	Certification	Foreign-market reach
	(1) Clients	(2) Certification	(3) Foreign markets
Panel A. OLS (log clients and markets, conditional on a positive report; the 0–4 certification ladder)			
Legacy match × Remediation period (1998–99, β_1)	0.059	0.319***	0.145
	(0.186)	(0.087)	(0.143)
Legacy match × Post-deadline period (β_2)	0.080	1.066***	0.150
	(0.194)	(0.175)	(0.161)
Mean of Y, control pre-period	3.7	0.05	2.6
Observations (firm-years; firms)	559 (205)	984 (235)	769 (211)
Panel B. PPML (clients and markets in levels; certification omitted — see note)			
Legacy match × Remediation period (1998–99, β_1)	0.200	—	0.194
	(0.182)		(0.168)
Legacy match × Post-deadline period (β_2)	0.111	—	0.038
	(0.196)		(0.158)
Observations (firm-years; firms)	559 (205)	—	855 (235)
Firm fixed effects	✓	✓	✓
Cohort × year fixed effects	✓	✓	✓
M&A event controls	✓	✓	✓

Panel C. Standardized post-deadline contrasts

		One-sided p-value		
Contrast	Difference (SE)	Clustered	Permutation	Within-size
Certification − clients	+0.811 (0.314)	0.005	0.004	0.150
Certification − foreign markets	+0.818 (0.193)	< 0.001	0.001	0.006

Notes. Each column estimates equation (19) using the predetermined 1995 legacy-match indicator L_i^{bin}. All specifications include firm and entry-cohort-by-year fixed effects and M&A event controls; standard errors are clustered by firm. The omitted pre-period is 1992–95, the remediation window is 1998–99, and the post-deadline window is 2001–03 for certification and 2001–02 for the count outcomes.

Panels A–B. Panel A reports OLS estimates for log clients conditional on a positive report, the 0–4 certification ladder, and log foreign markets. The mean-of-Y row reports levels. Panel B reports PPML estimates for the two count outcomes; $\exp(\beta) - 1$ is the proportional differential. No PPML certification estimate is reported because the ordered ladder changes almost exclusively upward and provides no informative Poisson comparison.

Panel C. Each outcome is standardized within its own estimation sample before the three equations are stacked; the stacked system contains 235 firms. All p-values test whether the certification contrast is positive. "Clustered" uses the stacked system's firm-clustered covariance; the permutation columns reassign the legacy match unrestrictedly or within pre-period size groups (999 draws). These are sensitivity checks, not design-based randomization inference. The within-size reassignment weakens the certification–clients contrast but not the certification–markets contrast. Figure 8 and Appendix Table B.7 report the pre-trend and coding checks. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

carries a smaller lead.

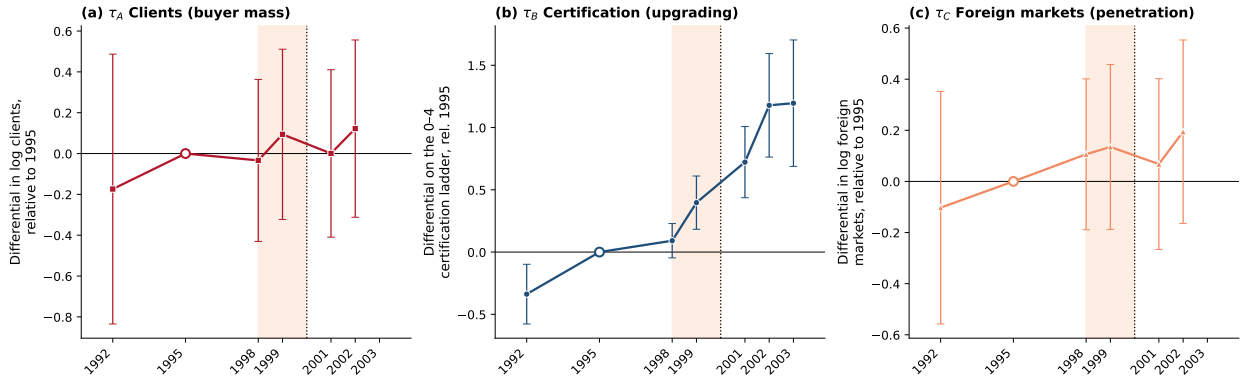


Figure 8: The upgrading fingerprint: certification ratchets up, while reported clients and foreign-market reach show no robust persistent differential.

Notes: Each panel plots an OLS event study of the predetermined 1995 legacy-match indicator L_i^{bin} interacted with directory-wave dummies, relative to the omitted 1995 wave; the open marker denotes this normalization. Panel (a) plots log reported clients, panel (b) the 0–4 certification ladder, and panel (c) the log markets-inclusive foreign-country count. The client outcome conditions on a positive report; comparison firms enter client reporting from a lower baseline, so this panel measures client-count intensity among reporters. Foreign-market reach uses the markets-inclusive count; the branch-country field has no usable pre-1998 baseline. Both count outcomes exclude the post-2002 directory-format change, so their post-deadline window is 2001–02. All specifications include firm and entry-cohort-by-year fixed effects and M&A event controls; whiskers are 95% confidence intervals from firm-clustered standard errors. The shaded band marks the 1998–99 remediation window and the dotted line the January 2000 deadline. The certification-minus-clients contrast (+0.81 SD, unrestricted permutation $p = .004$) does not survive a size-respecting permutation null ($p = .15$), while the certification-minus-markets contrast (+0.82 SD) does ($p = .006$); the two are always reported together. Table 6 reports the split-window coefficients, the standardized cross-outcome tests, PPML checks for the count margins, and both permutation variants.

The certification response and nearby outcomes. Table 7 first separates the certification ladder into its components. Column (1) reports the 0–4 ladder. Columns (2) and (3) use separate, nonexclusive indicators equal to one when the directory reports CMM status or an ISO-family credential. The post-deadline differentials are 32.6 percentage points for CMM and 24.2 percentage points for the ISO-family indicator. CMM is the closer observed measure of software-process maturity, and its pre-deadline divergence is smaller but still nonzero: the 1992 coefficient is −0.067 with a standard error of 0.036.

Beginning in 2001, free-text certification reports become the only source and can reveal CMM levels more readily than the earlier checkbox matrix, although earlier directories also contained some certification text. In 1999, levels were already known for 24 of 28 CMM-reporting firms, so greater level visibility in the later source could directly move at most four 1999 CMM reporters above rung 2. The post-deadline estimate remains positive when the outcome is any reported ISO-family or CMM credential and when the ladder is capped at

rung 2 (Appendix Table B.7). Newly revealed high CMM levels therefore do not account for the result by themselves, although other differences in what the two source formats record may still contribute.

Columns (4) and (5) turn to client composition and reported service and capability breadth. Column (4) is the share of classified named client relationships that are foreign, including Indian subsidiaries of multinationals; client names that cannot be classified are excluded. It is not a revenue share. This share is already about 0.80 among exposed firms before the shock and does not rise after the deadline. Column (5) counts the service and capability categories a firm checks in the directory's structured grid, excluding industry categories. Because the grid changes across editions, the count is standardized within each year. Its post-deadline differential is −0.055. A separate measure constructed from free-text service descriptions is weakly positive at +0.144 (SE 0.082), and an upgrading composite that excludes certification has no precisely estimated persistent differential. On a standardized scale, certification exceeds the foreign-client-share response by 1.00 standard deviation. Certification therefore moves without a comparably robust change in client mix or reported service and capability breadth. Figure 9 plots the corresponding event studies for the ladder, reported CMM status, and the foreign-client share.

Table 7: The certification response and nearby outcomes

	Certification components			Alternative margins	
	(1) Cert. ladder	(2) Any CMM	(3) Any ISO family	(4) Foreign-client share	(5) Service/capability breadth
Legacy match × Remediation period (1998–99, β_1)	0.319***	0.079***	0.174***	−0.027	0.043
	(0.087)	(0.028)	(0.060)	(0.077)	(0.146)
Legacy match × Post-deadline period (β_2)	1.066***	0.326***	0.242***	−0.028	−0.055
	(0.175)	(0.057)	(0.067)	(0.090)	(0.176)
Mean of Y, control pre-period	0.05	0.01	0.02	0.64	8.6
Observations (firm-years; firms)	984 (235)	984 (235)	984 (235)	609 (201)	781 (235)
Firm & cohort×year FE, M&A controls	✓	✓	✓	✓	✓

Notes. Each column reports a separate OLS difference-in-differences estimate using the predetermined legacy-match indicator L_i^{bin}, with firm and entry-cohort-by-year fixed effects and M&A event controls; standard errors are clustered by firm. Columns (1)–(3) report the 0–4 certification ladder and separate, nonexclusive indicators for reported CMM and ISO-family certification. Column (4) is the share of classified reported client names that are foreign, including Indian subsidiaries of foreign multinationals, not a revenue share. Column (5) is a within-year standardized count of service and capability categories in the directory grid, excluding industries; its mean row reports the raw count. Columns (1)–(4) use 1992–95 as the omitted pre-period and 2001–03 as the post period. Column (5) is observed in 1995, 1998–99, and 2001–02; it uses 1995 as the omitted pre-period and 2001–02 as the post period. For the certification–client-share contrast, both outcomes are standardized within their estimation samples; the post-deadline difference is +1.00 standard deviation (SE = 0.29, $p < 0.001$). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Other directory measures provide more limited evidence about the form of upgrading. The employment decomposition above is tilted toward reported software and engineering

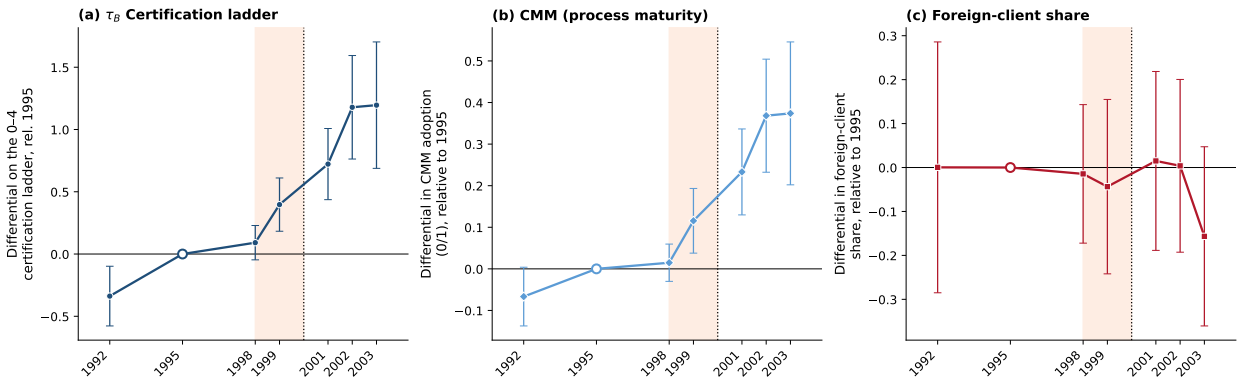


Figure 9: Certification rises persistently, while the foreign-client share does not.

Notes: Each panel is a separate OLS event study of the indicated outcome on the predetermined legacy-match indicator L_i^{bin} interacted with directory-wave dummies, relative to the omitted 1995 wave; the open marker denotes this normalization. All specifications include firm and entry-cohort-by-year fixed effects and M&A event controls; whiskers are 95% confidence intervals from firm-clustered standard errors. Panel (a) is the 0–4 certification ladder, panel (b) reported CMM status, and panel (c) the share of a firm's classified clients that are foreign. The foreign-client share is not a revenue share. CMM has a smaller pre-period lead than the overall ladder, but its 1992 coefficient remains marginally negative. The 2003 dip in panel (c) reflects the directory-format change that added newly classified clients to the denominator. Table 7 reports the component estimates and the certification-minus-foreign-client-share contrast.

roles, but these labels do not measure skills, wages, or productivity.

The visible technology stacks provide a separate check. Among the 160 firms that report comparable stacks in both 1995 and 1999, frontier breadth rises by 1.21 categories among exposed firms and by 0.91 among comparison firms. Both groups therefore add frontier technologies. The 0.30 differential has a standard error of 0.25, so the evidence does not show strong differential modernization but cannot rule out a modest difference. The sample covers only the shock window and cannot speak to stack changes after the deadline. The raw legacy-count contrast is affected by mean reversion and becomes uninformative after conditioning on its baseline (Appendix A.5).

Where the certification response is concentrated. The preceding results ask what changed after the shock. The next exercise asks a different question: whether pre-shock stack coherence predicts which firms later certify. The baseline indicator L_i^{bin} equals one if a firm reports any of 13 qualifying legacy technologies and classifies 135 of the 235 firms as exposed. Figure 10 instead uses a stricter partition on the common 225-firm sample: 71 firms report no legacy core layer, 99 report at least one core layer but not a coherent stack, and 55 report a coherent stack. Within this common sample, the baseline exposed group contains 54 coherent firms, 79 incomplete-stack firms, and two firms with no core layer. The

stricter partition is therefore not a literal decomposition of L_i^{bin}. Relative to firms with no legacy core layer, the post-deadline certification differential is +0.17 ladder points for an incomplete stack and +0.86 for a coherent stack.

Appendix Table B.3 enters core-layer presence and coherence jointly. The incomplete-minus-no-core estimate is +0.171 ladder points (SE 0.235, $p = .469$). The coherent-minus-incomplete increment is +0.685 ladder points (SE 0.249, $p = .0065$), making the coherent-minus-no-core total +0.856. This is heterogeneity in the depth of the predetermined technical match: it shows where the later certification response is concentrated, not that certification caused employment growth. Section 7 separately examines post-1998 foreign-presence formation among surviving firms and pre-shock connectivity as a possible gate. Those outcomes leave durable participation as a live complement or alternative to organizational capability.

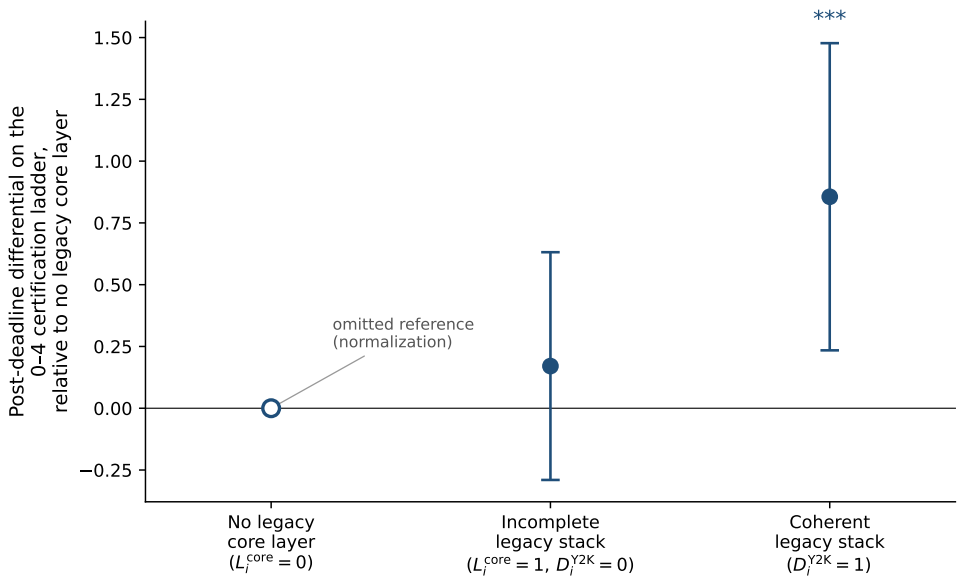


Figure 10: The certification response is concentrated among firms with coherent legacy stacks.

Notes: The figure plots post-deadline (2001–03) difference-in-differences coefficients on the 0–4 certification ladder by predetermined legacy-stack category, relative to firms with no legacy core layer. The open marker denotes the omitted no-core group; zero is a normalization, not an estimate. The categories are mutually exclusive: no legacy core layer ($L_i^{\text{core}} = 0$); at least one legacy core layer but no coherent stack ($L_i^{\text{core}} = 1$, $D_i^{\text{Y2K}} = 0$); and a coherent legacy stack ($D_i^{\text{Y2K}} = 1$). Coherence means passing the production-anchored IBM-mainframe or AS/400 gate: both require a legacy language, an operating-system or hardware anchor, and at least three coherent layers within one ecosystem, while the mainframe gate also requires a database or transaction layer. The regression is a split-window version of equation (19) on the common 225-firm technical-match sample, with firm fixed effects, entry-cohort-by-year fixed effects, M&A controls, and remediation- and post-window interactions for reporting verbosity, frontier breadth, and 1992 stack-field observability. Whiskers are 95% confidence intervals from firm-clustered standard errors. Stars test each displayed category against the omitted no-core group; $^{***}p < 0.01$. Appendix Table B.3 reports the direct coherent-minus-incomplete contrast.

6.3 Sector-level timing and broader coverage

The research design using the NASSCOM panel follows firms observed in the 1995 directory. This subsection instead uses the full Prowess universe to compare the export path of software and IT-enabled-services firms with firms in other industries. Prowess provides standardized annual financial statements across a wide range of industries and has been used to study firm adjustment in India (Alfaro and Chari 2010; Goldberg et al. 2010; De Loecker et al. 2016). The analysis universe contains 26,374 listed and larger unlisted audited filers, including 793 software and IT-enabled-services firms. Its annual coverage lets me trace the sectoral export differential before, during, and after Y2K.

Figure 11 focuses on export earnings, the audited Prowess outcome most directly linked to foreign Y2K demand. It reports event-study estimates, relative to fiscal 1996, of the gap between software/ITES firms and a broad non-IT comparison group of manufacturing and service firms. Computer hardware is excluded from both groups because the treatment is defined around software-services work; finance is excluded because its accounts and deregulation path are not comparable. Panel (a) reports log export earnings. The software/ITES differential was already rising before remediation, but it rises more sharply during remediation, flattens after fiscal 2000, and remains elevated through the 2001–02 technology bust. The corresponding split-window estimates are 1.06 log points during fiscal 1998–99, 1.66 in fiscal 2000, and 1.46 in fiscal 2001–04.

Panel (b) traces the share of firms with a nonzero export-earnings record, a rough measure of the extensive margin. The indicator equals one when Prowess records positive export earnings and zero when the field is explicitly zero or unreported in an otherwise observed annual statement. The pre-remediation path is difficult to assess because the early software/ITES sample is thin: the fiscal-1990 estimate rests on only seven observations, and the remaining estimates fluctuate and are imprecise. The differential rises to 13 percentage points in fiscal 2000 and remains 13–15 points higher through fiscal 2004. In raw levels, the proxy falls from .548 in fiscal 1990–97 to .487 in fiscal 2001–04 among software and IT-enabled-services firms, and from .475 to .323 among comparison firms. Both shares decline, but the decline is much smaller among software and IT-enabled-services firms.

A secular rise in software exports could explain part of the post-period gap, while changes in Prowess reporting could either widen or narrow the measured differential. A linear continuation of the pre-1997 export-earnings trend accounts for roughly 38% of the later gap, compared with roughly 99% of the sales gap, and the export path is similar when software/ITES firms are compared with trade-insulated non-tradable services and with other

services (Appendix Table B.16). The shape of the export path therefore corroborates the NASSCOM evidence that the effects of the Y2K shock persisted beyond the deadline.

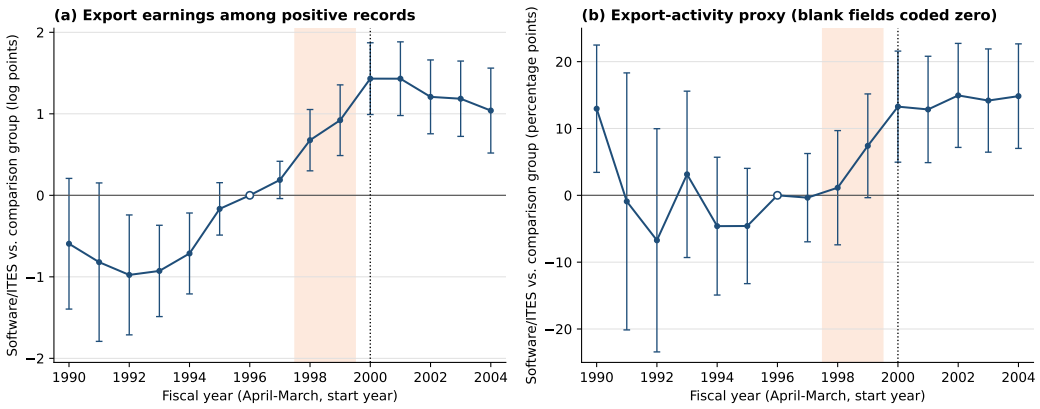


Figure 11: Audited export timing in the full Prowess universe.

Notes: The plots are firm- and year-fixed-effect event studies of predetermined software/ITES status, normalized to fiscal 1996 (open marker); whiskers are firm-clustered 95% confidence intervals. Panel (a) logs positive Prowess export earnings (N = 37,735; 292 contributing software/ITES firms). Panel (b) codes a positive record as one and an explicit zero or nonreporting in an observed annual statement as zero (N = 95,260; 579 contributing software/ITES firms). In the 97,676-record source frame, 39,072 records are positive, 62 are explicit zeros, and 58,542 are unreported; panel (b) is therefore a reporting-inclusive proxy, not verified export entry. Its fiscal-1990 estimate (0.130, SE 0.049) uses seven software/ITES observations. Shading marks fiscal 1998–99; the dotted line marks fiscal 2000, the first fully post-deadline year. See Appendix Table B.16 for the comparison groups and complete specification.

Online Appendix Section H extends the sector comparison to four audited investment and input margins. Among firms reporting positive values, the employee-count proxy widens after remediation, while the capex differential peaks in fiscal 2000 and remains positive afterward. These patterns align with the directory employment result and the matched-sample evidence on capital deepening. Booked R&D and royalties plus software charges show no detectable differential among positive reporters. Because employee and R&D coverage is thin and changes over time, these estimates describe reporting-subset patterns rather than sector-wide input changes.

Regional coverage with entrants. The regional design provides an entrants-inclusive within-IT check using local IT presence measured through fiscal 1996. Appendix Table B.17 and the annual paths in Online Appendix Section H relate audited outcomes to this exposure, defined as each district's share of IT firms among audited non-finance filers with sales observed through fiscal 1996, leaving the firm's own presence out. Because the exposure does not require a 1995 directory record, it is available for entrants. The post-deadline

export-earnings coefficient is +0.31 log points per standard deviation (permutation-reference $p = .041$). The rough export-activity proxy is 5.6 percentage points higher ($p = .034$; acceleration $p = .016$). Gross fixed assets, sales, and compensation also accelerate relative to the shock window ($p = .034$, .042, and .030), while their post-deadline levels are imprecise; capex and total assets remain null. The 1993–95 annual paths are generally flat, and the export and fixed-asset differentials persist through the technology bust.

These estimates come from roughly sixteen effective geographic cells, and all quoted significance assessments use the 999-draw cell-bundle permutation reference test described in Appendix Table B.17. High-exposure cells are concentrated in the southern technology hubs, so a hub-favoring shock can reproduce the gradient. Restricting to large hub cells yields an export post coefficient of +0.41 log points ($p = .023$) but does not resolve that confound. Locations use 2019 registered offices, which agree with historical directory locations for 71% of matched firms. The historical-location override attenuates the export post coefficient to +0.25 log points ($p = .245$), while the acceleration contrasts on sales, fixed assets, and compensation hold under it (permutation-reference $p = .006$–$.011$). The location disagreements do not look like random relocations, so the override is a bound rather than a correction. The regional comparison is therefore an entrants-inclusive within-IT consistency check on export earnings, the rough export-activity proxy, and audited scale and input margins; it does not measure firm capability, foreign-market access, or export entry.

The boundary of audited accounts. The final Prowess exercise asks whether audited accounts can distinguish organizational upgrading from durable participation. Online Appendix Section H compares the persistence of investment margins—fixed assets and capex—with scale margins—sales, export earnings, and compensation. It does not find a statistically separable difference: the investment-minus-scale contrast is imprecise in the matched binary, matched continuous-dose, and full-Prowess designs, while both classes persist in the full-universe comparison. Financial statements observe these accounting margins, but not certification, reported clients, foreign-market reach, service scope, or comparable organizational routines. They can therefore corroborate selected inputs, foreign revenue, timing, and coverage; the distinction between organizational capability and durable participation remains anchored in the directory outcomes.

7 Observable Firm Characteristics and Upgrading

The results so far compare firms with and without a predetermined match to Y2K remediation. The timing of the demand shock was external to Indian software firms, but the legacy stack was not randomly assigned across them. It may have traveled with broader firm strength, established foreign relationships, or communications capacity that would themselves predict growth and certification. This section uses those predetermined characteristics to stress-test the two main results. I ask two questions. First, how much of the legacy-match differential survives when the observable characteristics enter as controls? Second, is the differential confined to firms that had already cleared particular margins, such as broad foreign reach or large initial scale? The first asks whether prior advantage explains the response. The second asks whether prior advantage was a precondition for it. A response confined to firms that had already established themselves abroad or grown large would mean the shock rewarded incumbency. A response that appears across those margins would mean that firms which had not yet proven themselves could convert an ordinary commercial opportunity into durable upgrading.

Neither result is explained by these characteristics, and neither is confined to firms that had already cleared those margins. Measured firm quality absorbs a substantial part of the certification differential, but what remains is still precisely estimated, and the employment differential is essentially unchanged. Legacy-matched firms with only one foreign market before the shock climbed the certification ladder as far as those already selling into several, and they grew in every tercile of initial size. Their own reported data links predict later certification without absorbing the legacy-match coefficient. What the panel cannot evaluate is the district infrastructure behind those links, since almost every firm was already in a served district before remediation began.

7.1 Conditioning on observable prior advantage

Table 8 asks how much of the legacy-match differential, the post-deadline gap between firms with and without the pre-shock match, survives when the predetermined firm characteristics enter as controls. Firm fixed effects already absorb the level differences between matched and unmatched firms, so the concern is not that legacy-matched firms were larger or older to begin with. It is that firms with those advantages may have been on a different path around the deadline for reasons unrelated to Y2K. Each characteristic therefore enters interacted with the remediation-window and post-deadline indicators, which lets it predict its

own change over those windows rather than a difference in levels alone. Columns (2)–(4) add one block of characteristics at a time, and column (5) adds all three. First, the firm-quality block contains 1995 employment, client-industry breadth, services breadth, age, and registered capital. Second, the market-position block is an index of five forms of prior foreign engagement: foreign-market reach, alliances, foreign affiliation, Western-market presence, and an international branch. Third, the communications block contains firm-reported data-link and email indicators and whether the firm's district had operational STPI service by 1995.[6]

Firm quality accounts for a substantial share of the certification differential but does not absorb the post-deadline result. Without controls, legacy-matched firms sit 1.066 rungs higher on the 0–4 ladder than comparison firms after the deadline, measured against their own pre-shock positions. Adding the five quality measures leaves 0.609 rungs, a 43% reduction, still precisely estimated. Attenuation of this kind is expected, since the characteristics are correlated with the legacy match and move over the same windows. The attenuation falls almost entirely on the remediation window. That coefficient drops from 0.319 to 0.083 with the quality block and to 0.072 with all observables, while the post-deadline coefficient survives. What the observable characteristics leave unexplained is therefore concentrated after the deadline, when the certification gap continued to widen.

Prior market position and communications capacity move the post-deadline coefficient less, to 0.928 and 0.891 when each block enters on its own. With all three blocks, it is 0.571 (column 5), about 46% below the baseline. The within-size permutation companion gives $p = .001$ under its stated exchangeability restriction. Replacing the broad data-link indicator with the stricter international-link flag moves it to 0.577, which is no change at all. While the gap between the unconditioned and conditioned estimates is not a causal decomposition, since the controls are correlated and imperfect proxies for what they measure, each number answers a question worth asking. The 1.066 estimate is the full legacy-match differential, including the pre-shock firm bundle with which the match traveled. The 0.571 estimate is the residual beyond that bundle.

Employment is even less sensitive to the same attributes. Its post-deadline coefficient is 0.425 at baseline, rises to 0.515 with the quality block, and is 0.476 with all observables, or roughly 61% more employees than comparison firms. The increase under the quality

[6]Client-industry breadth counts the Y2K-relevant client industries a firm listed in 1995 and services breadth the distinct services recorded in its directory entry; employment and registered capital enter in logs; age is years since incorporation. The five quality measures are standardized, so each coefficient is the change associated with a one-standard-deviation move, and the market-position index standardizes and averages its five components.

Table 8: The legacy-match differential survives conditioning on predetermined firm characteristics

	(1)	(2)	(3)	(4)	(5)
	Baseline	Firm quality	Market position	Communications	All observables
Panel A: certification ladder (0–4)					
Legacy match × Remediation	0.319***	0.083	0.250***	0.245***	0.072
period (1998–99, β_1)	(0.087)	(0.088)	(0.085)	(0.085)	(0.089)
Legacy match × Post-deadline	1.066***	0.609***	0.928***	0.891***	0.571***
period (2001–03, β_2)	(0.175)	(0.181)	(0.176)	(0.174)	(0.176)
Panel B: log employment					
Legacy match × Remediation	0.172	0.185	0.130	0.137	0.171
period (1998–99, β_1)	(0.133)	(0.146)	(0.135)	(0.141)	(0.146)
Legacy match × Post-deadline	0.425**	0.515***	0.356**	0.357**	0.476**
period (2001–03, β_2)	(0.172)	(0.194)	(0.172)	(0.170)	(0.185)
Firm quality × period		✓			✓
Market position × period			✓		✓
Communications × period				✓	✓
Missing indicators × period		✓			✓
Firm and cohort×year FE; M&A	✓	✓	✓	✓	✓
Within-size permutation p (cert.)					0.001
Pre-period mean: certification	0.107	0.107	0.107	0.107	0.107
Pre-period mean: log employment	4.343	4.343	4.343	4.343	4.343
Observations: certification	984	984	984	984	984
Observations: employment	953	953	953	953	953

Notes: Split-window difference-in-differences estimates corresponding to equation (19). The firm-quality block contains standardized 1995 employment, Y2K-client-industry breadth, services breadth, age, and registered capital. The market-position block contains a standardized index of five pre-shock foreign-engagement measures: foreign-market reach, alliances, foreign affiliation, Western-market presence, and an international branch. The communications block contains indicators for a firm-reported data link and email in 1995 and for location in a district with operational STPI service by 1995. Every characteristic is interacted with the remediation and post-deadline windows. The quality variables are median-imputed using the median across firms to preserve the estimation universe, and the corresponding missing-value indicators are included in columns (2) and (5). Missing firm counts are 9 for employment, 74 for client-industry breadth, 1 for services breadth, 11 for age, and 86 for capital. On the 108 firms with all five quality measures observed, the all-observables post coefficients are 0.634 (0.283) for certification and 0.127 (0.265) for employment; the latter subsample is uninformative. Replacing any reported data link with the stricter international-link flag gives post coefficients of 0.577 for certification and 0.465 for employment in column (5). Standard errors clustered by firm in parentheses. The permutation row reports the add-one two-sided p-value from 999 draws reassigning the legacy-match indicator within 1995-employment terciles, with missing employment as its own stratum. Legacy stacks were not randomly assigned, so this is a sensitivity under the stated exchangeability restriction rather than design-based randomization inference. $^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$.

block indicates mild suppression rather than positive confounding on this outcome. Market position and communications reduce the coefficient separately to about 0.36, but their joint inclusion with firm quality leaves an estimate close to the baseline.

Two of the quality measures are incomplete. Client-industry breadth is missing for 74 of the 235 firms and registered capital for 86, so the specifications replace a missing value with the median across firms and add an indicator marking the affected firms, which keeps them in the estimation universe without treating the median as if it had been observed. Because the imputation applies to controls rather than to the outcome or the exposure, it governs how much the controls can absorb, not whether the differential is there. Among the 108 firms with all five measures observed, and so with nothing imputed, the certification differential is 0.634 ladder points (SE 0.283), slightly larger than the conditioned full-sample estimate. That subsample is too small to test employment.

Firms' own reported connectivity predicts later certification but does not absorb the legacy-match coefficient. In Appendix Table B.14, predetermined own links strongly predict subsequent certification, but their 1992 coefficients show that connected and unconnected firms were already on different paths before the remediation market opened. When own connectivity and the legacy match enter together, the legacy coefficient declines from 0.955 to between 0.796 and 0.817 and remains precise. Online Appendix Section G shows the same pattern in reported link take-up. Exposed firms were already pulling ahead of comparison firms on reported links before remediation began, and the deadline produces no visible break. Connectivity therefore marks firms that were already on a different path, so its coefficient does not measure what infrastructure contributed.

7.2 Prior foreign reach and initial scale

Prior foreign reach could have been a precondition for the response rather than a correlate of it. Relational accounts of buyer–supplier matching motivate the question. Under incomplete contracts, incumbent relationships and reputation can determine which supplier receives difficult-to-contract work (Banerjee and Duflo 2000; Macchiavello 2022). A deadline pushes the other way. A client facing a fixed date and a shortage of programmers who can read its own code has reason to look past its established vendors. Theory does not settle which force dominates, and neither does the model of Section 4, since prior reach shifts ordinary demand and the timing of upgrading as well as access to Y2K orders. The directory records which firms were serving which foreign markets years before the shock, so it can settle where the response appears. It does not record buyer search or order receipt, so it cannot show how

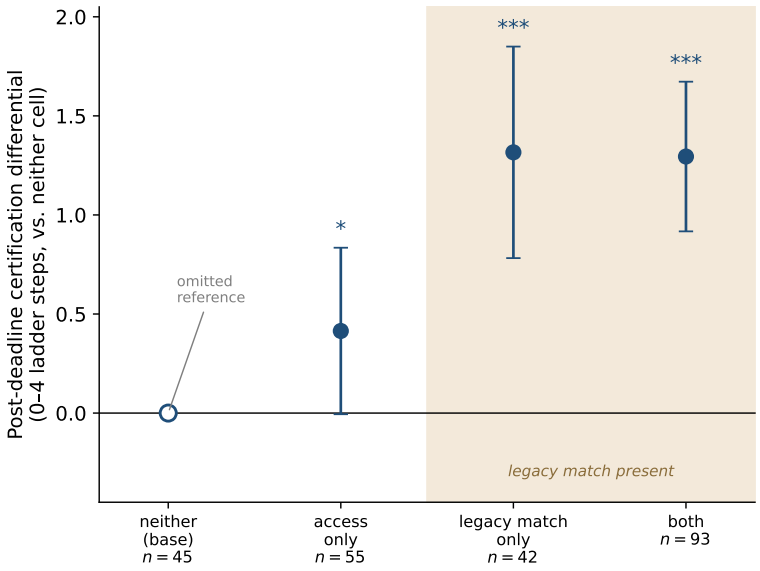


Figure 12: The certification response is not concentrated among firms with broad prior reach.

Notes: Post-deadline ({2001–03}) difference-in-differences coefficients on the 0–4 certification ladder by predetermined 1995 legacy-match×prior-reach cell, relative to the "neither" cell. The legacy match is L_i^{bin}; broad prior reach means more than one reported foreign market, so the limited-reach group includes firms reporting zero or one market. Firm counts appear beneath the cell labels. Whiskers are 95% confidence intervals from firm-clustered standard errors. Coefficients are within-firm and survivor-conditional. Only 19 of 42 limited-reach legacy-matched firms and 69 of 93 broad-reach legacy-matched firms contribute identifying pre- and post-deadline observations. An illustrative imputation assigning no certification gain to every non-returning limited-reach legacy-matched firm leaves that cell's differential relative to the neither cell at +0.60; this is not a bound against every selection process. Appendix Table B.12 reports the full quadrant estimates, an index-based classification, continuous specifications, pairwise tests, and permutation inference.

buyers and suppliers found each other.

Figure 12 divides firms by two 1995 characteristics: whether they reported a legacy technology and whether they served more than one foreign market. Limited reach therefore means zero or one reported market, not the absence of foreign engagement. Relative to firms with neither characteristic, the post-deadline certification coefficients are 1.316 for legacy-matched firms with limited reach and 1.295 for those with broad reach. The difference is 0.021 ladder points (SE 0.286; permutation $p = .948$).

The two legacy-matched cells have almost identical post-deadline coefficients. Both stand more than 1.3 ladder points above firms with neither characteristic, and the gap between them is under two percent of that. Firms serving at most one foreign market before the shock upgraded as far as those already selling into several. The estimates do not establish that the two are equal. The two-sided 95% interval for the limited-minus-broad difference runs from −0.54 to 0.58, so moderate differences in either direction remain consistent with

the data, and a one-sided interval rules out a limited-reach premium above roughly half a ladder point. The same one-sided bound is 0.21 ladder points under the index-based classification in Appendix Table B.12, which replaces the two binary cuts with median splits of a technical-match index and a prior-foreign-engagement index.

Differential attrition is the main threat to that comparison. Limited-reach legacy-matched firms are less likely to reappear in the directory after the deadline, at 45% against 74% for broad-reach legacy-matched firms. The imputation reported in the figure notes assigns no certification gain to every limited-reach firm that fails to return and still leaves that cell 0.60 ladder points above the neither cell, so successful returners alone do not generate the result. It cannot recover outcomes for firms that disappeared. The comparison also cannot be repeated at the more demanding coherent-stack threshold, where 45 of the 55 qualifying firms already had broad reach. The firms that do remain also move together on a margin the certification ladder does not capture. Legacy-matched firms were more likely to establish a physical presence abroad over the same window. Among firms observed in both 1998 and 2001 with no foreign presence in 1998, 60% of legacy-matched firms reported one by 2001 against 33% of other firms ($p = .012$), and within the legacy-matched group the rate is the same for limited- and broad-reach firms, at 62% and 59%. These are small, survivor-conditional cells, so the transition corroborates the reach pattern rather than identifying deadline-induced search (Appendix Figure B.6).

This design answers a different question from an engineered-access design. Atkin et al. (2017a) randomly assign access to a demanding foreign buyer and identify what that access does to upgrading. Prior reach is not assigned here. The question this setting can answer is whether the firms that converted an unassigned shock into durable upgrading had to be the ones already selling abroad. They did not.

Initial scale. Initial scale is the second margin that could have gated the response, and it did not. Appendix Table B.13 re-estimates the legacy-match differential within terciles of 1995 employment, cut at 35 and 125 employees. Post-deadline employment differentials are 0.50, 0.42, and 0.50 log points in the small, middle, and large terciles, or roughly 65%, 52%, and 65% more employees. Small legacy-matched firms grew as much as large ones. Large firms were not the only beneficiaries of the remediation boom, though the tercile intervals are wide enough that this is evidence against a large-firm-only pattern rather than a precise ranking across sizes.

Certification shows a size gradient that employment does not. The legacy-match estimates rise across the terciles, at 0.364, 0.416, and 1.005 ladder points, and standardized

1995 employment independently predicts 0.405 additional post-deadline rungs. The legacy match still adds 0.680 rungs once size is absorbed flexibly, so the certification response is not attributable to initial size. The lower-tercile estimates are too imprecise to establish that the legacy-match differential itself rises with size.

7.3 Communications as an enabling condition

Communications capacity is the third margin that could have gated the response, and it is the one margin the panel can date but not price. Offshore remediation was delivered over data links, and the panel measures that channel in two ways: firms' own reported links and email, which enter the controls above, and records from the Software Technology Parks of India (STPI), which date when the scheme's package of dedicated links, export facilitation, and tax benefits reached each district. Neither is a clean firm-level treatment. Own links are a firm's choice rather than an assignment. District service bundles connectivity with the rest of the policy package and does not establish that a given firm was connected, and some Y2K work was performed at client sites rather than delivered remotely. Instead, the two together establish when the infrastructure arrived relative to the demand.

The infrastructure arrived first, and by a wide margin. STPI service covered 64% of the 1995 firm universe and 92% by 1998. Of the 235 firms, 131 were in pilot districts served by 1991 and another 94 in districts first served by 1997, leaving three served only after the deadline and seven never recorded as served. The capacity to deliver offshore work was already in place across almost the whole panel before the demand for it appeared, so the timing of infrastructure cannot explain why these firms grew when they did. The same chronology leaves nothing to compare. Only seven post-deadline firm-years, contributed by four firms, occur in a not-yet-served district, and exactly one of those firms carries the 1995 legacy exposure, and center placement was deliberate rather than quasi-random. The panel therefore cannot price what the infrastructure contributed. That is a support failure rather than evidence that infrastructure was irrelevant. The institutional chronology and specification diagnostics are in Online Appendix Section G.

Connectivity does show up on the destination margin. In the firm-level access-margin estimates of Online Appendix Section G, connected legacy-matched firms reach more foreign markets, though the within-legacy-match gradient of roughly 0.45 log points is about half the reported triple difference, because the comparison-side gradient is negative. Client counts and reported export values show no corresponding pattern. Connectivity therefore marks which firms sold where, and, read with the pre-existing certification paths in Table B.14,

which firms were already ahead. It does not explain the durable upgrading.

Taken together, the technology a firm happened to run in 1995 predicts durable upgrading beyond everything observed here about its initial strength, foreign engagement, and communications capacity. Neither the certification nor the employment response belongs to the firms that were already large or already selling widely abroad. These exercises establish a robust separating role for the legacy match, not a causal ranking of technical fit over access, scale, or infrastructure. The next section addresses a different threat, whether the residual legacy-match differential reflects contemporaneous aggregate forces, especially the dot-com boom, rather than Y2K.

8 Alternative Explanations

This section evaluates the leading alternative explanations for the paper's two main findings—the persistent employment expansion among legacy-matched firms after the deadline, and their durable rise in certification. Because the specification absorbs any change common to all firms in a cohort-year cell, an alternative explanation must be correlated with the 1995 legacy stack and must act on employment or certification independently of the remediation shock. First, the Y2K remediation window sits inside the largest expansion in global IT demand in modern history, the dot-com bubble and subsequent burst. That boom may have rewarded legacy-matched firms independently of the deadline, in which case what mattered was general demand for software labor, not Y2K remediation. Second, the certification gap between the two groups may already have been widening before 1995 and simply kept widening, in which case the post-deadline difference is a trend continuing rather than a response to the shock. Third, two domestic schemes rewarded ISO and CMM credentials over the same years, in which case the upgrading was paid for by subsidy rather than demanded by foreign buyers. Fourth, the differentials may be artifacts of how firms reported, which firms the directory observed, or how the specification is drawn, in which case there is no response to explain. Fifth, formal education supply may have expanded where legacy-matched firms were concentrated, in which case part of the differentials could reflect a contemporaneous change in the local labor pool rather than the remediation shock.

I find that none of these alternatives accounts for the two differentials. I test each against evidence it should have produced if it were the explanation. First, general IT demand does not explain the legacy-match response specific to Y2K. A frontier technology stack predicts no additional Y2K engagement once the legacy stack is held fixed, and measured exposure

to general United States IT spending does not subsume the certification differential. The certification path also bends at the January 2000 deadline and at no other point, while audited software export earnings stay elevated through the 2001–02 dot-com bust rather than falling back with it. Second, the later certification differential is not confined to the firms already certified in 1995. Firms reporting no certification at all in 1995 also upgraded differentially after the deadline, and I bound how much of the magnitude a continuing pre-1995 divergence could still account for. Third, the response extends to reported process maturity, a margin no Indian scheme in this window subsidized. Fourth, the certification result survives alternative codings of the outcome, and both differentials survive the sample, measurement, specification, and attrition checks. Fifth, recognized private-secondary and English-medium private-secondary entry shows no robust, deadline-timed acceleration in IT-exposed districts, and newly founded formal institutions are unlikely to have expanded the supply of suitably trained workers at scale before the deadline.

8.1 The broader IT and offshoring boom

The Y2K window sits inside the dot-com boom, so legacy-matched firms may simply have been the firms best placed to capture a general rise in offshore demand. Three comparisons say otherwise.

The legacy match picks out the requirement Y2K created, not technical sophistication in general. Firms running the obsolete stack in 1995 were 0.327 more likely to report a Y2K engagement afterwards, against a sample mean of 0.47 (SE 0.076), and the category count and coherence measures rank firms the same way. An index built the same way from modern technologies predicts engagement weakly on its own and not at all once the legacy stack is included (Table 4 in Section 5). Remediation demand selected on the obsolete stack, not the frontier one.

General IT demand does not absorb the certification response. Table 9 builds two exposures from the same pre-shock client-industry shares, one weighted by United States IT spending overall and one by measured Y2K-remediation intensity. On its own, general spending predicts post-deadline certification weakly, at 0.197 ladder points per standard deviation (SE 0.109). Adding the firm-level legacy match cuts that to 0.127 and leaves the legacy coefficient at 0.903 (column 3). Client-sector composition by itself supplies little identifying variation, and the comparison covers certification only, for the 161 firms that report client shares.

The Y2K shock and the dot-com boom fall on different timelines, and the certification

Table 9: The legacy match versus general IT spending after the deadline

	Dependent variable: the 0–4 certification ladder			
	(1)	(2)	(3)	(4)
Panel A. Both exposures entered together				
General IT spending $x^{IT}\times$ Post (per SD)	0.197*	0.228**	0.127	−0.045
	(0.109)	(0.114)	(0.101)	(0.147)
Y2K demand intensity $x^{Y2K}\times$ Post (per SD)		−0.101	−0.040	−0.035
		(0.106)	(0.096)	(0.136)
Legacy match $L_i^{\text{bin}}\times$ Post			0.903***	
			(0.235)	
Firm fixed effects	✓	✓	✓	✓
Cohort × year fixed effects	✓	✓	✓	✓
M&A event controls	✓	✓	✓	✓
Fine size × year fixed effects				✓
Mean of Y (pre-period)	0.106	0.106	0.106	0.106
Observations	701	701	701	701

Panel B. Each exposure alone, under stricter fixed effects	Legacy match L_i^{bin}	IT spending x^{IT} (per SD)
Cohort × year	1.066*** (0.175)	0.197* (0.109)
Size-tercile × year	0.695*** (0.201)	0.069 (0.111)
Fine size × year	0.597** (0.268)	−0.054 (0.138)
Observations	984	701

Notes. The dependent variable is the 0–4 certification ladder; every column is the split-window form of equation (19) with the listed exposures each interacted with the shock and post windows (shock-window coefficients included but not reported; none is significant where the post term dies). Standard errors clustered on the firm in parentheses. x^{IT} weights the firm's pre-shock client-industry shares by the BEA general-IT-spending shifter, and x^{Y2K} weights the same shares by the measured Y2K-remediation shifter; both are scaled per cross-sectional standard deviation, and the sample is the 161 firms reporting client-industry shares (701 firm-years, of which 22 are firm-singletons that do not contribute to the within-firm estimates). Column (1) enters general IT spending alone; column (2) races the two shifters; column (3) adds the predetermined legacy match L_i^{bin}; column (4) replaces the legacy-match control with fine (0.1 log-point) 1995-size-by-year fixed effects, within which an aggregate-shifter exposure has little variation left (297 size-by-year cells for the 701 observations). Panel B re-estimates each exposure alone under successively stricter fixed effects; the legacy-match rows use the full 235-firm sample (984 contributing firm-years). The free-text composite index (a visible-stack measure) survives the legacy-match control and dies only under size-by-year effects, consistent with certification being the durable margin; that disclosure, the deadline-timing tests, the Euro comparison, and the placebo and permutation tests are reported in the text. $^{*}\,p<0.10$, $^{**}\,p<0.05$, $^{***}\,p<0.01$.

response follows the Y2K deadline. The dot-com boom ran through 2000 and gave way to a bust in 2001–02. Y2K remediation had a single date: it had to be finished by January 1, 2000. The certification path of legacy-matched firms bends above its own pre-shock trend at that date, by 0.207 ladder points (SE 0.082), as shown in Figure 8. The audited Prowess accounts then show the export margin outlasting the boom. Software and IT-enabled-services export earnings stay elevated relative to other industries through the 2001–02 bust instead of falling back with it, and the same persistence appears among the matched firms and in the within-IT

regional comparison (Section 6.3).

A second deadline in the same window would be a natural test, and the Euro conversion is the closest candidate. The Euro was anticipated too, but it arrived through a scheduled multi-stage changeover rather than on one date, so it cannot serve as an equally clean timing test. Euro-exposed firms were also larger, and once exposure is reassigned within size strata the post-2001 certification rise is no longer distinguishable, at a permutation p of .277 against .032 without the strata (Appendix Figure B.5). The Euro pattern is not separable from baseline size composition.

Three limits apply to these comparisons. A general boom can leave durable differences of its own when firms sink investment or form lasting buyer relationships, so persistence through the bust corroborates the account rather than separating it from a boom. The cross-industry comparison carries differential pre-trends, and the regional comparison rests on a small number of hubs. Engagement is reported after the shock, so an unmeasured firm advantage could still be correlated with both the legacy stack and subsequent upgrading.

8.2 Exposure timing and pre-shock trends

Y2K was known before remediation spending accelerated, so firms could in principle have acquired the legacy stack in anticipation of the work. Rebuilding the exposure from the 1992 directory is the strictest test that they did not. On that smaller sample of 124 firms the post-deadline certification differential is 0.67 ladder points (SE 0.31) under the binary legacy measure and 0.61 (SE 0.16) under the category count; employment is −0.02 (SE 0.26) and 0.25 (SE 0.11) (Appendix Table B.9). Certification survives the stricter base. The binary employment estimate is uninformative, and the category-count estimate is positive and precise.

Employment shows no pre-shock divergence but certification does. The 1992 employment interaction is −0.043 (SE 0.259), so nothing distinguishes the two groups' earlier paths, though a single pre-wave cannot rule out much (Appendix Table B.8). Certification is different. The gap between the two groups widened between 1992 and 1995, with a 1992 interaction of −0.338 ladder points (SE 0.122), and every one of the 16 firms certified by 1995 is legacy-matched. Carried forward in a straight line, that early divergence would account for 67–94% of the later gap.

How much of the certification result survives that continuation can be bounded (Rambachan and Roth 2023). The post-deadline difference survives trend violations up to a quarter of the observed lead per directory-wave step and is not distinguishable from zero at half (Appendix

Figure B.4). The wave-step scaling is demanding, because the single observed pre-step spans three calendar years while the late post-steps span one. Subtracting a straight-line extrapolation of the 1992–95 divergence outright leaves 0.24 ladder points, with an interval of $[-0.39, 0.88]$. The exercise bounds the size of the response. It does not pick a counterfactual path.

The later differential is not an arithmetic artifact of those early adopters. Among the 206 firms reporting no certification at all in 1995, legacy-matched firms still rise 0.783 more ladder points (SE 0.191). The same ordering holds under four further cuts, including size-quintile-by-year effects and the CMM margin among firms holding no CMM in 1995 (Appendix Table B.10). The response is made of first-time adopters, not of firms extending a lead they already held. Conditioning on 1995 certification removes the earlier difference by construction, so the flat pre-period in that subset follows from the design, and a persistent difference in latent adoption propensity among firms uncertified at that date remains admissible.

8.3 Certification policy

Neither of the two domestic schemes that rewarded certification over these years can account for the response, because each rewarded the different firms or margin. From 1993–94 a Government of India scheme reimbursed 75 percent of admissible ISO-certification expenses for registered small-scale units (Office of the Development Commissioner (MSME), Government of India 2005). From 1998 software exporters holding ISO 9000 or SEI CMM Level 2 or higher qualified for Special Import Licences worth 5 percent of f.o.b. export value (NASSCOM 1999, pp. 24–25, 40–41). Since cohort-by-year effects absorb common policy timing, the alternative needs differential eligibility or take-up among legacy-matched firms, not merely the existence of the schemes.

The reimbursement scheme excluded the firms that responded. The reimbursement scheme was restricted to registered small-scale units, and legacy-matched firms are the larger and older ones, at 301 employees against 95 at the 1995 baseline and ten years old against six (Appendix Table B.1). The data record neither program registration nor firm-level reimbursement, so size alone cannot settle who actually claimed it.

The import-licence rule set one qualification for ISO 9000 and for CMM Level 2 or higher, offering no further statutory reward above that threshold. The upgrading response from the Y2K shock runs partly above it. Censoring the ladder at rung 2 lowers the differential from 1.066 to 0.583 on the same 984 firm-years, leaving 0.483 ladder points of response above

the level the licence rewarded (Appendix Table B.7). No Indian scheme in this window subsidized CMM appraisal at all, and reported CMM status rises 0.326 (SE 0.057; Table 7).

What survives is a narrower version of the alternative. Firm-level participation in either scheme is unobserved, so policy could have amplified a response the demand shock started. It could not have produced the part of the response the schemes did not reward.

8.4 Measurement, observation, and estimation

Reporting format. The certification differential does not come from a change in how firms reported. Both the outcome and the legacy stack come from directory reports, and the format changed in 2001 from a checkbox matrix to free text, which can make CMM levels more visible. Earlier records already identify the level for 24 of the 28 firms reporting CMM in 1999, so at most four reports could have been reclassified upward for that reason. The differential stays positive under four codings of the outcome and the sample, including dropping the 2002 and 2003 waves (Appendix Table B.7).

Attrition. Differential attrition does not generate the results either. Directory reappearance is not legal survival, so the estimates are conditional on the 1995 incumbents observed on both sides of the deadline. Balanced-panel estimates, inverse-probability reweighting, and restriction to firms with a surviving legal entity all leave the post-deadline differentials positive. Lee bounds, which assume monotone selection, are uninformative for employment. The two groups reappear at almost identical rates after 2001, 63.6 against 63.8 percent, so the trim is under one observation and both bounds sit at the 0.500 point estimate. For certification a 4.9 percent trim gives bounds of 0.937 to 1.115 around 1.045.

The pooled employment rates mask selection that runs in opposite directions among initially small and initially large firms (Appendix Table B.13). The two Heckman corrections reproduce the uncorrected coefficient almost exactly, 0.579 against 0.580. Neither is used as the preferred estimate, because one rests on an exclusion restriction the data cannot defend and the other is identified only by functional form (Appendix Table B.11).

Estimators. The estimator does not carry the result. Exposure begins at the same date for all treated firms, so the forbidden comparisons associated with staggered adoption do not arise (de Chaisemartin and D'Haultfœuille 2020). Re-estimating with the imputation estimator of Borusyak et al. (2024) gives 0.60 for employment and 1.12 for certification, against 0.42 and 1.07 in the baseline, and the group-time estimator of Callaway and Sant'Anna (2021)

gives a post-deadline employment effect of 0.51 (Appendix Table B.6). Both alternatives are larger than the headline, not smaller.

Specification. The specification is not doing the work either. Alternative legacy-match constructions, PPML for employment, and nineteen sample and measurement toggles all preserve the pattern (Appendix Tables B.2 and B.7). Dropping the merger and acquisition controls, which could themselves respond to the shock, moves employment from 0.425 to 0.435 and certification from 1.066 to 1.099, so those controls are not carrying the results (Appendix B.4).

Inference. Firm-clustered standard errors are the primary inference procedure. Reassigning the legacy match at random reproduces the headline in about two percent of draws, whether the reassignment is unrestricted ($p = .018$) or restricted to firms with comparable client portfolios ($p = .020$).

These checks carry their own limits. None of the coding variants turns a self-reported credential into an audited measure, and Prowess corroborates selected financial outcomes without observing certification at all. The attrition exercises weaken a simple differential-attrition account and leave selection on unobservables unresolved. The alternative constructions are correlated reconstructions of the same directories, and cells selected for a strong first stage are not independent replications. Because the legacy match was not randomly assigned, the permutation results are sensitivity exercises under an exchangeability restriction rather than design-based randomization inference. Two further checks add no affirmative evidence. The foreign-affiliate placebo is imprecise, and the pseudo-deadline exercise recovers the 1992 certification lead that Section 8.2 already discloses, so both are reported in the appendix rather than counted as successful placebos.

8.5 Local formal education supply

The last alternative locates part of the response outside the firm. A contemporaneous expansion in local formal education where legacy-matched firms were concentrated could make the employment and certification differentials partly a labor-supply response rather than a demand-induced change in firms. The firm specification does not absorb district-specific changes in the local labor pool. I therefore examine one observable supply margin: entry by recognized private-secondary and English-medium private-secondary schools. Private-secondary founding rose sharply at the national level, but entry on these two margins did

not show a robust, deadline-timed acceleration in districts where IT employment was concentrated. Appendix B.10 reports the design, estimates, and limitations.

Timing provides the firmer bound. Y2K remediation relied on specialized technical skills that had to be available by January 2000. Because acquiring the relevant formal skills generally took longer than the remaining window, institutions founded after remediation demand accelerated in 1998 could not plausibly have expanded the supply of suitably trained workers at scale before the deadline. New formal-school entry is therefore unlikely to have been an important source of the immediate labor response. The evidence also weighs against a large, geographically concentrated entry response on the measured margins. It does not rule out smaller effects, new seats or quality changes within existing institutions, short-cycle training, migration by already-skilled workers, or education responses unfolding over a longer horizon.

Taken together, the five alternatives leave the two main differentials standing. A common late-1990s IT shock cannot by itself produce a differential graded by the 1995 legacy stack, and the certification response is not composed of the firms that had already certified. The response extends to first-time adopters and to reported process maturity, the audited export margin outlives the boom it would otherwise be attributed to, and newly founded formal institutions are unlikely to have supplied an important share of the immediate labor response.

The design cannot rule out every unobserved advantage correlated with the legacy stack, or differential take-up of certification policy, and what remains differs by outcome. For employment, the legacy match predicts a persistent relative expansion on a sparse but comparatively flat observed pre-path. For certification, it predicts broad subsequent upgrading whose exact magnitude stays sensitive to the 1992–95 divergence and to unobserved policy participation. That joint pattern supports a temporary demand shock followed by durable upgrading.

9 Conclusion

This paper uses Y2K as an expiration test for demand-driven upgrading. A large export opportunity opened for Indian software firms with expertise in legacy technologies and then closed at a fixed, externally determined date. If that opportunity affected only current demand, the exposed–comparison differential should have receded when remediation ended. Persistence instead requires a firm state that outlasted the market that created it. Among

firms already present in the 1995 industry directory, the predetermined legacy match strongly predicts subsequent Y2K engagement and a durable post-deadline divergence. The employment differential is modest and statistically indistinguishable from zero during remediation, but reaches 0.425 log points—approximately 53%—in 2001–03. Registered capital moves earlier, while matched audited accounts provide corroborating evidence on compensation, fixed assets, and export earnings. The clearest scale result is therefore not a contemporaneous employment expansion during remediation, but a larger operating scale after direct remediation demand had disappeared.

The post-deadline divergence also appears on a capability-related margin. Legacy-matched firms move 1.066 points farther up the 0–4 ladder of software-process and quality-management certification, with substantial differentials in both CMM status and ISO-family credentials. The response is concentrated among firms with coherent legacy production stacks, remains present among firms reporting no certification in 1995, and is not confined to the initially largest firms or those with broad prior foreign reach. The magnitude requires qualification. Certification had already begun to diverge between 1992 and 1995, and allowing predetermined firm quality, market position, and communications capacity to generate their own remediation and post-deadline paths reduces the differential to 0.571 ladder points. Reported client counts, foreign-country reach, client composition, and service breadth show no comparably robust persistent differentials, although the estimates do not rule out economically meaningful participation responses. The evidence thus documents persistent scale alongside a durable organizational-capability or quality-assurance differential. It does not establish certification as the mediator of employment or exclude an unobserved participation asset.

The model interprets this joint pattern as an investment-timing problem. A process and quality-management system is a sunk investment whose return rises with the scale and complexity of the firm's project load. A firm may know how to adopt such a system and still postpone doing so because ordinary demand does not yet justify paying the fixed cost. Temporary demand can change that calculation. Because the capability remains useful in ordinary projects after remediation ends, the finite Y2K premium need not cover the entire sunk cost; it need only make adoption now more valuable than waiting. The induced response is therefore concentrated near the no-Y2K adoption margin. Once the deadline passes, the legacy match no longer raises demand directly, but an organizational capability acquired while that match was valuable can continue to support ordinary business. This is the obsolete-match reversal: expertise in technologies abandoned by the frontier can

predict modern organizational capability after the commercial premium on that expertise has expired. Comparison firms may later reach their own adoption dates, so persistence through 2003 is an empirical result rather than a claim of permanent divergence.

The interpretation remains subject to important limits. The legacy match was not randomly assigned, and the causal reading depends on conditional parallel trends among firms with and without that match. The directory provides only one pre-1995 wave, certification is self-reported and has an early differential, and the main estimates are conditional on incumbent firms observed around the deadline. Exposure reconstructed solely from the 1992 technology stack continues to support the certification result, but the corresponding employment evidence is mixed. The design also does not separate the experiences bundled inside Y2K work. Greater project load is the model's baseline mechanism, but buyer standards, learning, financing, relationships, and access to global delivery networks may also have changed the return or cost of upgrading. The robustness exercises make several simple alternatives less persuasive without eliminating all unobserved advantage. Nor does the paper estimate Y2K's contribution to India's aggregate software-export takeoff. Its claim is narrower: firms positioned to exploit the episode underwent a differential process of organizational maturation and emerged operating at greater scale.

The broader implication is that foreign demand can transform firms not only through what buyers teach, what quality they require, or which markets they open, but also through when demand makes a fixed organizational investment worth undertaking. This is not a claim that temporary demand generally produces persistent growth. Persistence requires access to the opportunity, an induced investment that remains productive in ordinary business, and firms close enough to an adoption threshold for the temporary return to change their timing. When those conditions hold, however, a temporary market can be catalytic rather than merely transitory. India's legacy software expertise created access to a market scheduled to disappear; the firms positioned to serve it emerged with modern process credentials and operating capacity whose value extended beyond Y2K. A capability inherited from the past can therefore shape future advantage when a transient opportunity makes a more general capability worth building.

References

Alfaro, Laura and Anusha Chari, "India transformed? Insights from the firm level 1988–2005," *India Policy Forum*, 2010, *6*, 153–224. Volume 2009–10, National Council of Applied Economic Research and Brookings Institution; pages 212–224 are the discussants' comments and general discussion. Earlier circulated as NBER Working Paper 15448 (2009), DOI 10.3386/w15448.

Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P Vasquez, "The effects of joining multinational supply chains: New evidence from firm-to-firm linkages," *The Quarterly Journal of Economics*, 2022, *137* (3), 1495–1552.

Ali, Nadia, Giacomo De Giorgi, Aminur Rahman, and Eric Verhoogen, "What do market-access subsidies do? Experimental evidence from Tunisia," NBER Working Paper 33985, National Bureau of Economic Research July 2025.

Arora, Ashish and Jai Asundi, "Quality certification and the economics of contract software development: A study of the Indian software industry," NBER Working Paper 7260, National Bureau of Economic Research July 1999.

___ **and Suma Athreye**, "The software industry and India's economic development," *Information Economics and Policy*, 2002, *14* (2), 253–273.

Asher, Sam, Tobias Lunt, Ryu Matsuura, and Paul Novosad, "Development research at high geographic resolution: An analysis of night-lights, firms, and poverty in India using the SHRUG open data platform," *The World Bank Economic Review*, 2021, *35* (4), 845–871.

Athreye, Suma S., "The Indian software industry and its evolving service capability," *Industrial and Corporate Change*, 2005, *14* (3), 393–418.

Atkin, David, "Endogenous skill acquisition and export manufacturing in Mexico," *American Economic Review*, 2016, *106* (8), 2046–2085.

___ **, Amit K. Khandelwal, and Adam Osman**, "Exporting and firm performance: Evidence from a randomized experiment," *The Quarterly Journal of Economics*, May 2017, *132* (2), 551–615.

___ **, Azam Chaudhry, Shamyla Chaudry, Amit K Khandelwal, and Eric Verhoogen**, "Organizational barriers to technology adoption: Evidence from soccer-ball producers in Pakistan," *The Quarterly Journal of Economics*, 2017, *132* (3), 1101–1164.

Baldwin, Richard and Paul Krugman, "Persistent trade effects of large exchange rate shocks," *The Quarterly Journal of Economics*, November 1989, *104* (4), 635–654.

Banerjee, Abhijit V. and Esther Duflo, "Reputation effects and the limits of contracting: A study of the Indian software industry," *The Quarterly Journal of Economics*, 2000, *115* (3), 989–1017.

Bloom, Nicholas, Aprajit Mahajan, David McKenzie, and John Roberts, "Do management interventions last? Evidence from India," *American Economic Journal: Applied Economics*, April 2020, *12* (2), 198–219.

___ , Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does management matter? Evidence from India," *The Quarterly Journal of Economics*, February 2013, *128* (1), 1–51.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event-study designs: Robust and efficient estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Bustos, Paula, "Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms," *American Economic Review*, February 2011, *101* (1), 304–340.

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

Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal, "Misallocation in firm production: A nonparametric analysis using procurement lotteries," NBER Working Paper 31311, National Bureau of Economic Research 2023. Latest version CEPR Discussion Paper 21106, January 2026.

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout, "Market entry costs, producer heterogeneity, and export dynamics," *Econometrica*, May 2007, *75* (3), 837–873.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Fan, Tianyu, Michael Peters, and Fabrizio Zilibotti, "Growing like India—the unequal effects of service-led growth," *Econometrica*, 2023, *91* (4), 1457–1494.

Ferraz, Claudio, Frederico Finan, and Dimitri Szerman, "Procuring firm growth: The effects of government purchases on firm dynamics," NBER Working Paper 21219, National Bureau of Economic Research 2015. Latest revision 2021.

Garicano, Luis, "Hierarchies and the organization of knowledge in production," *Journal of Political Economy*, October 2000, *108* (5), 874–904.

Geanakoplos, John and Paul Milgrom, "A theory of hierarchies based on limited managerial attention," *Journal of the Japanese and International Economies*, September 1991, *5* (3), 205–225.

Gerschenkron, Alexander, *Economic Backwardness in Historical Perspective: A Book of Essays*, Cambridge, MA: Belknap Press of Harvard University Press, 1962.

Ghose, Devaki, "Trade, internal migration, and human capital: Who gains from India's IT boom?," *Review of Economics and Statistics*, 2026. Advance online publication, 24 March 2026; volume, issue, and pages not yet assigned. Earlier circulated as World Bank Policy Research Working Paper 9738 (2021).

Goldberg, Pinelopi K, Amit K Khandelwal, Nina Pavcnik, and Petia Topalova, "Multiproduct firms and product turnover in the developing world: Evidence from India," *The Review of Economics and Statistics*, 2010, *92* (4), 1042–1049.

Hanlon, W Walker, "Temporary shocks and persistent effects in urban economies: Evidence from British cities after the U.S. Civil War," *The Review of Economics and Statistics*, 2017, *99* (1), 67–79.

___ , "The persistent effect of temporary input cost advantages in shipbuilding, 1850 to 1911," *Journal of the European Economic Association*, 2020, *18* (6), 3173–3209.

Heath, Rachel and A. Mushfiq Mobarak, "Manufacturing growth and the lives of Bangladeshi women," *Journal of Development Economics*, 2015, *115*, 1–15.

Ilzetzki, Ethan, "Learning by necessity: Government demand, capacity constraints, and productivity growth," *American Economic Review*, August 2024, *114* (8), 2436–2471.

Jaravel, Xavier, "The unequal gains from product innovations: Evidence from the us retail sector," *The Quarterly Journal of Economics*, 2019, *134* (2), 715–783.

Jensen, Robert, "Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India," *The Quarterly Journal of Economics*, 2012, *127* (2), 753–792.

Juhász, Réka, "Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade," *American Economic Review*, 2018, *108* (11), 3339–3376.

Khanna, Gaurav and Nicolas Morales, "The IT boom and other unintended consequences of chasing the American dream," Working Paper 25-01, Federal Reserve Bank of Richmond 2025. DOI 10.21144/wp25-01. Earlier circulated as Center for Global Development Working Paper 460 (2017).

Lane, Nathan, "Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea," *Quarterly Journal of Economics*, 2025, *140* (3), 1683–1741.

Lileeva, Alla and Daniel Trefler, "Improved access to foreign markets raises plant-level productivity... for some plants," *The Quarterly Journal of Economics*, August 2010, *125* (3), 1051–1099.

Loecker, Jan De, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik, "Prices, markups, and trade reform," *Econometrica*, 2016, *84* (2), 445–510.

Loughran, Tim and Bill McDonald, "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks," *The Journal of Finance*, 2011, *66* (1), 35–65.

Macchiavello, Rocco, "Relational contracts and development," *Annual Review of Economics*, 2022, *14* (1), 337–362.

Masters, Steve and Carol Bothwell, "CMM Appraisal Framework, Version 1.0," Technical Report CMU/SEI-95-TR-001, Software Engineering Institute, Carnegie Mellon University February 1995.

McDonald, Robert and Daniel Siegel, "The Value of Waiting to Invest," *Quarterly Journal of Economics*, 1986, *101* (4), 707–727.

Melitz, Marc J., "The impact of trade on intra-industry reallocations and aggregate industry productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Mitrunen, Matti, "War Reparations, Structural Change, and Intergenerational Mobility," *Quarterly Journal of Economics*, 2025, *140* (1), 521–584.

NASSCOM, "Indian I.T. Software and Services Directory 1999–2000," New Delhi: National Association of Software and Service Companies 1999.

Office of the Development Commissioner (MSME), Government of India, "Achievements of the Small-Scale Industries Sector, May 2004–April 2005," Ministry of Micro, Small and Medium Enterprises, New Delhi 2005. https://dcmsme.gov.in/achvmnts.htm, accessed 29 July 2026.

Oster, Emily and Bryce Millett Steinberg, "Do IT service centers promote school enrollment? Evidence from India," *Journal of Development Economics*, 2013, *104*, 123–135.

Paulk, Mark C., Bill Curtis, Mary Beth Chrissis, and Charles V. Weber, "Capability Maturity Model for Software, Version 1.1," Technical Report CMU/SEI-93-TR-024, Software Engineering Institute, Carnegie Mellon University February 1993.

Rajan, Raghuram G. and Luigi Zingales, "Financial dependence and growth," *American Economic Review*, 1998, *88* (3), 559–586.

Rambachan, Ashesh and Jonathan Roth, "A more credible approach to parallel trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Roberts, Mark J. and James R. Tybout, "The decision to export in Colombia: An empirical model of entry with sunk costs," *American Economic Review*, September 1997, *87* (4), 545–564.

Shastry, Gauri Kartini, "Human capital response to globalization: Education and information technology in India," *Journal of Human Resources*, 2012, *47* (2), 287–330.

United States Senate, Special Committee on the Year 2000 Technology Problem, "Y2K aftermath—crisis averted: Final committee report," Technical Report, 106th Congress, 2nd Session, U.S. Government Printing Office, Washington, DC February 2000. Issued February 29, 2000. Robert F. Bennett, Chairman; Christopher J. Dodd, Vice Chairman.

U.S. Securities and Exchange Commission, "Statement of the Commission Regarding Disclosure of Year 2000 Issues and Consequences by Public Companies, Investment Advisers, Investment Companies, and Municipal Securities Issuers," Release Nos. 33-7558, 34-40277, IC-23366, IA-1738; 63 Fed. Reg. 41394 (Aug. 4, 1998) July 1998.

Verhoogen, Eric, "Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector," *The Quarterly Journal of Economics*, May 2008, *123* (2), 489–530.

——— , "Firm-level upgrading in developing countries," *Journal of Economic Literature*, 2023, *61* (4), 1410–1464.

A Data: Sources, Construction, and Validation

This appendix documents how I built the firm-year panel, what I did to check it, and what it cannot measure. Section A.1 describes the directory source and its digitization. Section A.3 describes the procedure that assigns stable firm identifiers, and Section A.4 reports the audits that test it, including a cross-validation against the corporate registry. Section A.5 states the limitations that bind the analysis. Section A.6 summarizes the constructed variables and Section A.7 gives the construction of the coherent legacy match used in Section 5. The Online Appendix carries the rest of the construction record: the page-level discovery sweep and the directory's non-profile content, the entity-resolution scoring and adjudication detail, panel summary statistics, and the integration of the external datasets.

A.1 The NASSCOM directories

The panel is built from the *Directory of Indian Software & Service Companies*, which the National Association of Software and Service Companies (NASSCOM) published annually through the 1990s and early 2000s. Each edition is a printed volume of several hundred pages carrying a structured profile page per member firm: name, address, chief executive, founding year, ownership and share capital, software and other headcounts, revenue and software-export figures, service offerings, named clients, the programming-language and platform stack, certifications, Indian and international branches, foreign alliances, and target markets and industries. From 1998 the directory adds dedicated *Year 2000 Solutions* and *Euro Currency Solutions* columns, in which firms report contemporaneously whether they take on Y2K- and EMU-related remediation work.

NASSCOM was the trade body for export-oriented Indian IT in this window, so the directory is a voluntary member roster rather than a population register. Its widely cited coverage figure of roughly 95% refers to industry *revenue*, which a handful of large exporters dominate, and inclusion required returning a detailed survey form each cycle (Arora and Athreye 2002). The directory therefore captures the economically significant universe densely and the small-firm tail sparsely and intermittently. Two external checks bear this out: a 144-firm hand audit against corporate-history sources finds that the panel contains every firm appearing in standard accounts of the industry through 2003, and registry checks adjudicate the status of 74% of listed firms (Section A.4).

Seven editions survive and have been digitized: 1992, 1995, 1998, 1999, 2001, 2002, and 2003. No separate edition exists for 1996, 1997, or 2000, so those years are missing by

publication schedule rather than by attrition, and *directory year* throughout the paper refers to the year on the cover. I audited all 4,116 raw scan pages cover to cover and extracted 3,194 per-firm profile pages, which deduplicate to 3,232 firm-records. Firm-records can exceed profile pages because some editions print two firms on one page; those pages are detected at extraction and split. Scan quality is uniformly good from 1998; the 1992 and 1995 editions use older paper stock whose typos surface in the platform and language fields and are handled in the tokenization cleanup below.

A.2 From printed pages to structured records

The profile pages were converted into structured firm-year records by a dual-pass large-language-model extractor (Anthropic Claude Sonnet 4.6, temperature zero). The choice followed a ten-firm pilot benchmarked against Google Document AI, the prior baseline: on every benchmarked field the language model matched or exceeded the optical-character-recognition system, it returned structured values for fields that system does not extract at all (city, phone, named clients), it recovered canonical service names where that system pulled mid-sentence fragments, and it recovered all five pilot firms on which that system had failed outright.

The two passes run against the same page under the same schema but different instructions. The first emphasizes precision and the second recall, and outputs merge with provenance: each list element records which pass produced it, and scalar fields follow an agree-or-prefer-precision rule with every disagreement logged. Roughly 28% of firms carry at least one scalar disagreement between passes, concentrated in list-field coverage, which is the redundancy the second pass exists to supply. Before the full run I validated the design on a 35-firm test sample spanning all seven editions against a gate fixed in advance, requiring at least 90% coverage on named clients, services offered, and Year 2000 solutions. All seven editions cleared the gate, and the full extraction processed all 3,194 pages with no parse failures.

The schema is a typed backbone plus an open-ended remainder, and three of its design choices matter for the analysis. First, revenue is stored raw as four parallel series — software exports, software revenue, domestic revenue, and total revenue — each entry keeping the printed year, raw string, value, currency, unit, and a projection flag, so currency and reconciliation choices can be revisited without re-extraction. Directory pages often print a current-year actual alongside a one-year projection, and the flag separates them. Second, services and clients each split into an *explicit* field printed under the corresponding heading

and an *inferred* field derived from surrounding text, with each value entering exactly one of the two. The headline measures use the explicit lists and inclusive variants use the union. The 1992 edition prints no clients heading at all, so its client information is legitimately inferred-only, and across the panel inferred client coverage runs 34–72% by year against 1.5–3.1% for explicit named clients — an asymmetry that reflects the directory's layout rather than the extraction. Third, address is stored as a raw string with city captured separately; the printed state field is sparse and geography is instead reconstructed downstream from address and city, reaching complete district coverage.

Two mechanical issues were identified and handled. Directory typos in the noisy list fields — split tokens and spacing variants of the same product name — were repaired by a targeted cleanup pass over the 241 firms with the worst surface-form drift, which deduplicates variants within a firm and never normalizes across firms. And some editions print two firms on a single profile page, in 25 instances in 1992, 26 in 1995, 49 in 2002, and one in 2003; the extractor decides the number of firms from page content and the validation test covered all known multi-firm pages.

Table A.1 reports non-missing coverage by edition for the analytically important fields. The identity backbone of firm name, address, and phone is extractor-gated at 100%, and city, chief executive, and founding year sit at 95–100% throughout. Two patterns in the table carry analytical weight. Contact-technology fields rise from zero in 1992 to near-universal by 2001, tracking the diffusion of the technologies themselves. And the technology-stack fields collapse from roughly 90% coverage in 1995–1999 to 7–13% by 2003, a directory layout change whose consequences for measurement are in Section A.5.

A.3 Firm identifiers

Tracking firms across editions requires resolving name changes, mergers and acquisitions, spin-offs, exits, and re-entries into a stable identifier. Without it the panel is a sequence of cross-sections; with it the panel supports the within-firm designs the paper uses. One principle governs the construction: an identifier captures the same legal entity continuing across years, and nothing else. A parent and its Indian subsidiary, an acquirer and its target, and a division and its separately incorporated successor all receive different identifiers, with the corporate relationships among them recorded in a separate events layer. This is what lets the regressions net merger-driven growth out of organic growth by controlling for event flags, rather than silently consolidating balance sheets across ownership changes.

Resolution runs in five stages. Within-edition deduplication first collapses same-year

Table A.1: Field coverage and self-report prevalence by directory edition (%)

Field	1992	1995	1998	1999	2001	2002	2003
Founding year	98.4	95.4	97.3	96.7	97.6	99.9	97.7
Email	0.0	13.9	90.5	96.7	99.5	99.0	98.9
Website	0.0	0.0	49.4	77.1	91.5	96.7	97.5
Software employees	45.6	94.5	93.2	91.0	87.8	92.4	89.7
Services (explicit)	91.2	45.1	89.5	89.0	85.5	90.8	86.1
Services (inferred)	92.0	99.2	98.5	94.5	97.7	94.0	95.8
Named clients (explicit)	2.4	3.0	2.7	3.1	1.5	1.7	2.0
Clients (inferred)	33.6	55.3	68.6	65.1	70.1	67.3	72.1
Markets	98.4	95.8	89.8	85.7	87.3	87.8	90.8
Industries served	65.6	68.4	58.9	54.9	57.6	82.4	95.9
Foreign alliances	34.4	29.1	29.7	22.2	19.8	56.1	54.4
Programming languages	88.0	91.6	83.0	73.9	26.9	12.1	6.9
Operating systems	95.2	93.7	92.0	86.7	19.5	3.9	2.1
Platforms	96.8	94.9	93.2	91.2	65.3	12.9	10.1
Year 2000 self-report (% *True*)	0.0	0.0	36.5	23.3	0.0	0.0	0.0
Euro self-report (% *True*)	0.0	0.0	17.5	18.6	0.0	0.0	0.0

Notes: Cells are the share of firm-records in each edition for which the field is non-missing; empty lists, empty strings, and missing values all count as missing. The two rows below the rule report self-report prevalence, the share of firm-records answering *True*, rather than coverage; they are zero outside 1998–1999 because only those editions carry the columns. The identity backbone of firm name, address, and phone, together with city, chief executive, and founding year, is omitted for space and is discussed in the text. The revenue series are structured fields whose usable value coverage is reported in Section A.6. Certification is measured from the directory's dedicated certification tables in 1995, 1998, and 1999 and from the profile text thereafter, not from a single per-edition field. The printed state field is not used; geography is reconstructed from address and city.

rows that refer to one firm through typos or suffix variants. Cross-edition candidate generation then proposes matches between each within-year unit and units up to three editions later, blocking on name similarity, executive name, founding year, normalized address, and branch overlap, and scoring each surviving candidate by a fixed weighted sum over thirteen signals. Two thresholds calibrated against hand-judged samples sort candidates into automatic links, an ambiguous band, and non-matches. The ambiguous pairs are then adjudicated by an escalating sequence of language-model review passes, followed by firm-level corporate-history research with web search that yields 1,224 dated corporate events carrying 3,408 source citations. Finally, I reviewed the machine output by hand across 115 adjudications covering every candidate rename link, every conflict, and every remaining uncertain pair. The guiding rule was to trust the researched corporate history but audit its join back to the panel, since the failure mode is a real rename attached to the wrong panel firm by substring matching. The review accepted 87 rename links, rejected 21 as wrong joins of

exactly that kind, confirmed three uncertain pairs as different firms, and added two links the pairwise machinery had missed.

Transitive closure over the accepted links assigns the identifiers: 1,248 firms spanning 3,232 firm-records, of which 450 appear in a single edition and 798 are tracked across two or more. The largest connected components correspond to well-documented corporate lineages, among them Citicorp Overseas Software's rename to i-flex Solutions, identified through a shared address and chief executive. Of the 124 firms in the 1992 cohort, 43 still appear in the 2003 edition, a directory-presence retention rate of 34.7%; Section A.4 shows this is a member-roster retention rate far below true legal-entity survival, and the analysis treats it accordingly. The Online Appendix reports the candidate-pair counts, the scoring signals, the adjudication verdicts at each escalation, and the provenance of the 3,684 links in the final identifier graph.

A.4 Validation

Four audits and two external cross-validations test the construction, each with a decision rule fixed before it ran. All four audits passed.

Extraction accuracy. A stratified random sample of 50 firms across the seven editions was checked field by field against the source PDFs: the identity backbone, the address-derived city, the three fields central to the analysis, and the largest non-empty revenue stream, for 16 fields per firm and 800 firm-field cells in all. I reviewed all 800 cells against the source pages and observed no defects, against a threshold set in advance that would have triggered rework at a 5% error rate. One non-finding is itself a check: the Year 2000 Solutions field is blank for every sampled 1992 and 1995 firm, as it must be, since the directory introduces that column only in 1998.

Citation accuracy of the corporate-history research. The 1,224 researched corporate events carry 3,408 source citations spanning 2,609 unique URLs on 612 domains. I classified domains into seven credibility tiers — regulatory primary sources, business and technology press, encyclopedic sources, registry-derived aggregators of official filings, business databases, retail-brokerage profile pages, and a questionable long tail — and fetched a stratified sample of 70 URLs, adjudicating each against the event it supports. Twenty-four were directly supported, 28 presumed supported through search where the page sits behind a bot-block, 14 partial in that the page names the right firm without surfacing the specific year or successor,

two unsupported, and two dead links. The strict hallucination rate is 2.9%, within the 5% pass threshold; counting dead links it is 5.7%, and both are reported here. Citation practice in the paper follows the same tiering, using regulatory and press sources as primary evidence, registry-derived aggregators for legal-entity continuity, and everything below only as corroboration.

Join integrity and rename recall. A programmatic audit re-derived each of the 87 accepted rename links in both directions and classified every non-exact name match against the review notes, finding no further wrong joins: 61 exact matches, 13 legal-suffix variants, seven cases where the researched name contains the panel name in full, and six members of chains already accepted at review. Precision audits cannot show how many real renames were missed, so I also assembled a 29-case ground-truth list of well-documented Indian IT renames in 1992–2003 from press archives, regulatory filings, and company histories. Recall on the 25 verifiable cases is 25 of 25, and the catch includes multi-step chains — one firm traverses four names across the panel window and is correctly held together — and renames that fell between editions. Independently, a research assistant's 144-firm investigation log, compiled by hand against external sources, groups all 11 of its rename and merger cases correctly in the panel.

Government-registry cross-validation. Because NASSCOM is a private member roster, I searched every listed firm in India's Ministry of Corporate Affairs registry and its public mirrors for its corporate identification number, incorporation date, registration status, registered office, capital, industry code, and last filing. The search serves four purposes. First, directory credibility: of the 1,248 firms, 839 (67.2%) are verified as registered companies and a further 87 (7.0%) are correctly non-company entities such as partnerships, proprietorships, foreign liaison offices, and government societies that sit legitimately outside the Companies Act, so registry checks adjudicate the status of 74.2% of listed firms. The remaining 322 are untraceable and concentrated in the small-firm tail. Second, identifiers: a corporate identification number is recovered for 862 firms (69.1%), which is the join key to the audited financial data. Third, founding-year validation: the registry incorporation date matches the self-reported founding year within one year for 82% of the 859 comparable firms, with most of the remainder explained by registry entries that predate the self-report, typically older parent or holding entities.

Fourth, the registry brackets survival. The 34.7% directory-presence rate for the 1992 cohort is a lower bound; registry status implies an upper bound near 88%, with a floor

around 70% if every untraceable micro-firm died. Three measurement facts support reading directory attrition as roster lapse rather than firm death. The firms the directory stops tracking are the small single-listing tail: untraceable 1992-cohort firms have a median of 21 employees against 100–145 for those that stayed or were acquired. Registry strike-offs are dated, and every dated disagreement between the registry and the researched histories, 13 of 13, is a post-2003 strike-off, which corroborates rather than conflicts, since India's mass strike-off campaigns came much later. And each registry match was scored on four independent identity signals, with 676 strong, 152 medium, and only 11 weak matches, of which just one is a 1992-cohort survivor. The registry has limits of its own. Active status over-counts operating firms, and its electronic records begin in 2006, so it confirms survivors but cannot date a pre-2006 death. The paper therefore never headlines a survival rate, treats directory presence as sample selection, and uses the registry for what it measures well: the identifier crosswalk and the founding-year validation.

A.5 Known limitations

Directory presence is sample selection, not survival. The most consequential limitation for the analysis is that appearance and disappearance are endogenous to a firm's choice to renew membership and return the survey. Directory exit conflates true exit, acquisition, rename, and roster lapse, and is concentrated in the small-firm tail; the registry validation above bounds the wedge, at 34.7% directory presence against roughly 70–88% legal-entity survival for the 1992 cohort. The paper accordingly models directory presence as a selection mechanism rather than a survival outcome, through the inverse-probability weights and selection corrections of Appendix B.7, and quotes no single survival rate. Three facts limit the damage: differential attrition by pre-shock exposure is about seven percentage points, the exposure measure is predetermined and so uncontaminated by post-shock reporting choices, and the registry supplies the tools to bound the selection.

The technology-stack fields hollow out after 2000. A directory layout change collapses coverage of the programming-language, operating-system, and platform fields from roughly 90% in 1995–1999 to 7–13% by 2003 (Table A.1). The exposure measures are unaffected, since they are built from the 1992 and 1995 editions, but any post-2000 technology-adoption measure must lean on certifications and service breadth rather than the stack fields, and the paper's upgrading measures do so. The raw legacy-count contrast across waves is affected by mean reversion for the same reason and is uninformative once conditioned on its

baseline.

Five entity boundaries admit more than one reading. Two lineages were re-grouped after audit: the Pentafour lineage, where two separately incorporated firms under shared chairmanship had been grouped as one and were split, and the Silverline–Hexaware component, where a link rejected at manual review had remained active in one linkage layer because the rejection filter was applied asymmetrically. Both fixes are in the panel used throughout, and together they reassign nine firm-years, or 0.28% of the panel. Two further cases were audited and retained — the Wipro Systems and Wipro Technologies boundary, which tracks the operational hardware-versus-IT distinction, and the Tata Unisys and Tata Information Systems pair, verified as distinct joint ventures with different partners, cities, and founding years. One northern lineage is retained conservatively and flagged, because one edition's entry plausibly belongs to a different firm of the same name and the registry cannot confirm the distinction.

Smaller items. The 1999 edition lacks an alphabetical index in the surviving scan, so the recall safety net for that year is one layer thinner than for surrounding editions and the 1999 panel was cross-validated against the 1998 and 2001 indexes instead. In the 2002 edition, 49 joint-venture co-brand pages are mis-tagged as two-firm layouts in an audit column, which leaves the firm count correct. The directory's dedicated certification tables exist only for 1995, 1998, and 1999, so the post-1999 certification series comes from the profile text and the construction bridges the format split explicitly.

A.6 Constructed variables

The analytical variables are built from the panel in deterministic documented steps, and the replication package's codebook records each variable's exact construction rule, sources, coverage, and alternatives. This subsection summarizes what a reader needs in order to interpret the results, and closes with the caveats that bind how the variables can be read.

Foundations. Revenue series are reconciled actual-over-projection and converted at Reserve Bank fiscal-year rates; software exports and revenue are usable at 75–87% of firm-years, and projections overshoot realized values by a median of 44%, so projections never substitute for actuals. Employment coalesces the component headcounts with the printed total, reaching 94–98% coverage per year against 22–88% for the raw printed total, and capital

coalesces subscribed into authorized capital at 55–90%. Headquarters geography reaches complete district and state coverage on 2001 Census boundaries, reconstructed from address and city.

Exposure. The paper's exposure measures are predetermined, built from the 1992 and 1995 editions alone. The lead measure is the legacy-stack indicator L_i^{bin}, a thirteen-technology screen of the COBOL and mainframe families over the pre-shock stack fields, which marks 135 of the 235 firms in the 1995 estimation universe. Its refinement, the coherence gate D_i^{Y2K} and its dose, is in Section A.7. Around these sit the companion measures used for validation and robustness: the realized Y2K self-report, taken as the union of the 1998–99 checkbox and the profile narrative, which yields 405 firm-years against 264 for the checkbox alone because text-only reporters carry real engagements; the Euro self-report; a Y2K-heavy client-industries measure; pre-shock foreign reach, where the foreign-country count discriminates because the any-foreign indicator saturates near 90%; a foreign-affiliation placebo; and internet readiness from parsed connectivity fields. Realized self-reports are contemporaneous with treatment and are used only to validate the predetermined measures, never to define them.

Outcomes. The upgrading analysis distinguishes the client margin, built from the explicit and inferred client lists; the certification margin, a ladder in the spirit of Verhoogen (2008) spanning no certification through ISO to CMM levels; and the market-reach margin, built from foreign country counts using the variant that includes stated target markets, since branch counts have no pre-1998 base. Certification is complemented by four sibling upgrading measures — service breadth, a foreign-client share built by classifying all 7,583 distinct client names, the direction of technology-stack adjustment, and product breadth — plus a standardized composite. Baseline cross-correlations among these components run 0.06–0.28, so comparisons of their persistence are informative rather than mechanical. Pre-shock technical-match and foreign-access indices standardize over available components without imputing zeros for non-reporters, and the match-by-access classification used in Section 7 is reported in two co-equal versions, a composite-index split and a transparent raw split, whose cell agreement of 69–76% makes the dual reporting a robustness check.

Measurement caveats. Four caveats bind. First, directory coverage rises across editions and a 2002–03 format change inflates all count-based measures, so count trajectories are read within-firm and conditional on reporting guards, and 2002–03 levels are treated as

suspect. Second, missing is not zero: every count variable carries a reporting-guard companion, and genuinely missing values — for example the roughly 43% of pre-shock export volumes, missing selectively toward the small-firm tail — are never imputed and never enter composites. Third, predetermined and realized exposure are never pooled, since the predetermined measures come from 1992–1995 and the realized self-reports from 1998–99 and they play different roles by design. Fourth, the standardized composites are relative measures standardized within edition, so their levels are not comparable across years and persistence claims built on them are within-firm and relative. Every aggregate count in this appendix was re-derived independently from the underlying data, and a 470-cell spot check re-derived ten constructed variables per sampled firm-year from the raw panel, with no discrepancies.

A.7 The coherent legacy match

The empirical object behind the depth exposure of Section 5.1 is not how many legacy technologies a firm listed, but whether it had assembled a coherent legacy business-systems production environment before the shock. This subsection gives the classification behind the coherence gate D_i^{Y2K} and its dose. The measure is rule-based and uses no estimated weights; the weighted and latent constructions mentioned at the end are robustness only.

Classification. Each raw stack token is canonicalized and tagged on three axes: a stack *layer* — language, operating system, database, transaction or batch processing, hardware, plus optional fourth-generation tooling — an *era*, and a binary *core-legacy* flag. The decisive choice is that era is not core. Several legacy-era technologies are deliberately excluded from the core count because Y2K was a defect in two-digit *business-date* logic rather than in legacy computing as such. Scientific and teaching languages such as FORTRAN, Pascal, BASIC, and Ada, generic assembler, and the open-systems Unix variants — AIX, Solaris, HP-UX, which are the destination of modernization rather than its legacy source — are legacy-era but score zero on core legacy. The frontier client-server and enterprise-web technologies, including Visual Basic, the COM, CORBA, and EJB middleware families, and Java, are frontier rather than legacy. Genuine legacy-core technologies were rescued from the ambiguous tail, among them VSAM, MCP, GUARDIAN, VME, GCOS, NATURAL, FOCUS, CSP, TELON, and DMS II. The crosswalk classifies 470 tokens, and the eight ambiguous vendor and assembler tokens are held non-core, since they are the most likely false-positive source.

Ecosystems and layers. Complementarity is computed within an ecosystem, because legacy systems were proprietary. Four ecosystems appear in the data. The IBM mainframe contributes five production layers and the IBM AS/400 four; the DEC VAX and a proprietary group spanning Unisys, Tandem, ICL, Wang, HP 3000, and Prime contribute four each and enter as a heterogeneity arm rather than the headline exposure. The midrange and minicomputer architectures integrate transaction processing into the operating system, so the transaction layer drops from their denominator. Table A.2 details the two headline ecosystems layer by layer. The IBM-mainframe rule requires at least three of five layers, including a legacy language, an operating-system or hardware anchor, and a database or transaction layer; the AS/400 rule requires at least three of four, including a language and an operating-system or hardware anchor. The dose $\max\{C_{i,\mathrm{MF}}/5,\, C_{i,\mathrm{AS400}}/4\}$ records the completeness of the firm's best-matched ecosystem, where each C counts the production layers filled.

Table A.2: The two headline legacy ecosystems, by stack layer

Layer	IBM mainframe	IBM AS/400
Language	COBOL, PL/I *(support: Assembler, REXX)*	RPG, RPG/400, COBOL/400, COBOL
Operating system	MVS, OS/390, VSE *(support: VM/CMS)*	OS/400
Database / files	DB2, IMS, IDMS, VSAM	DB2/400
Transaction / batch	CICS, JCL, IMS DB/DC	— *(integrated in OS/400)*
Hardware	S/390, ES/9000, 30xx, 43xx	AS/400
Depth (4GL / tools)	NATURAL, FOCUS, CSP, TELON	SYNON, CL/400, SQL/400

Notes: Core fillers, in roman, count toward the firm's coherent-layer total and the gate. Support technologies, in italic, are legacy signal but not one of the counted production layers and enter only the item-response robustness vocabulary. Depth tooling populates a separate depth flag. The AS/400 has no separately counted transaction layer because transaction processing is integrated into OS/400. COBOL is the one shared language and counts toward whichever ecosystem the firm otherwise anchors. The two minicomputer ecosystems, DEC VAX and a proprietary group, follow the same four-layer logic and enter as a heterogeneity arm rather than the headline exposure.

Timing and missing values. The legacy match is measured on the pre-shock union of the 1992 and 1995 directories. Both years are predetermined relative to the shock, and pooling reduces false negatives from intermittent directory printing, since a deep legacy stack is sticky. A firm whose stack field was printed but empty has match zero; a firm whose field was never printed in either pre-year is coded exposure-missing rather than zero and drops from the estimation sample. The 1998–99 directories are reserved for first-stage validation

and never used to define treatment.

Alternative constructions. Beyond the headline gate and dose I compute a loose gate without the database-or-transaction clause, an ecosystem-agnostic count of coherent core layers, a broad-legacy gate spanning the minicomputer ecosystems, a weighted-complementarity score with within-ecosystem pairwise weights, and a two-parameter item-response latent score over the core-legacy vocabulary residualized on reporting verbosity. The headline results are invariant to these alternatives. Reporting verbosity itself is handled in estimation rather than by penalizing the index: every depth specification includes post-shock interactions with the log count of non-legacy tokens and with frontier breadth, and a separately constructed frontier-stack index serves as a placebo. That placebo is informative but not clean in both directions. The frontier index does not reproduce the certification pattern, since its event-study path is flat, yet it weakly predicts realized Y2K uptake ($t \approx 2.2$, against $t \approx 4.7$ for the gate). This is consistent with frontier strength and the legacy match being correlated among larger firms rather than with a frontier channel for the upgrading.

A.8 External sources and replication

Eight external sources join the directory panel, and the Online Appendix gives their vintages, crosswalks, and integration limitations in full. Three of the joins bear directly on how the results should be read. First, all geography runs through a canonical 2001-Census district spine; firm geocoding is city-level, and because firm clusters in three metropolitan areas straddle 2001 district boundaries, district-level analysis treats Delhi, Mumbai, and Kolkata as single metropolitan units. Second, the crosswalk to Prowess audited accounts matches 457 of the 1,248 firms, of which 454 are usable, and match probability is strongly size-selected — 91.5% for listed firms, 28% for unlisted, and 2.6% for firms with no registry record. This is why Prowess serves as a large-firm corroboration layer and never the headline sample. The audited panel used in the paper covers 3,319 firm-years for 230 matched firms, and its median $\log_{10}$ ratio of audited to self-reported revenue is -0.008 across overlapping firm-years. Third, the demand shifter is built from 110,082 10-K and 10-Q filings, extracted under a nulls-only-outside-window discipline with all sampled cost verbatims verified in the source text, amendment-deduplicated, with vendors and special-purpose entities excluded and conglomerates consolidated. The result is a \$52.0 billion disclosed client-demand universe, inside the contemporaneous \$100–115 billion estimates of total US remediation spending, aggregated from filer industry codes to the eight directory client-industry groups plus a residual,

revenue-weighted, and supplemented by $8.34 billion of federal Y2K outlays.

A replication package accompanies the paper. It contains the extraction, entity-resolution, variable-construction, and estimation code, together with the prompts and frozen model outputs of every model-assisted step; the audit artifacts behind every validation result reported here; a codebook documenting each constructed variable's construction rule, source fields, coverage, and alternatives; and a per-exhibit provenance table mapping every number printed in the paper to the dataset columns and code that produce it. The directory scans themselves are available subject to the source archives' terms.

B Additional Results

This appendix reports the exhibits behind the body's results, organized by the question each answers. Section B.1 describes the estimation sample. Sections B.2 and B.3 ask whether the results depend on how the legacy match is measured and which other firm margins moved. Sections B.4–B.7 take up the three standard challenges to a difference-in-differences design of this shape: the estimator, the pre-period, and sample selection. Section B.8 reports the full observable-characteristics exercises of Section 7, Section B.9 the audited accounts, and Section B.10 the local formal-education-supply margin. The Online Appendix carries the remaining material: the full exposure grid, the demand shifter's internal diagnostics, the infrastructure and connectivity exercises, the additional audited-accounts exhibits, and the coefficient table underlying the education figure.

B.1 Baseline composition of the estimation sample

Firms with the pre-shock legacy match differ from comparison firms at baseline, and Table B.1 reports by how much. The legacy stack was typically assembled by established companies, so matched firms are larger, older, more often public limited, and more oriented toward banking and finance. The normalized differences are between 0.4 and 0.8 on size, age, and client-industry mix. Identification does not require these firms to be comparable in levels, since firm fixed effects absorb level differences and the design reads upgrading off within-firm changes timed to an externally dated deadline. The table matters for two other reasons. Reporting itself differs by exposure — exports are reported by 70% of matched firms against 48% of comparison firms — so the revenue and export rows are conditional on reporting and are read jointly with the reporting-share rows. And the communications row shows the infrastructure margin was close to balanced before the shock, at 67% against 60%

of firms in a district with an operating STPI center.

Table B.1: Baseline covariates by exposure, 1995

Covariate	Legacy match	Comparison	Diff.	Norm. diff.	N_e/N_c
Total employees	301.36	95.27	206.10	0.49	132/94
log(1+employees)	4.74	3.76	0.98	0.75	132/94
Firm age (years)	10.19	5.92	4.26	0.51	134/90
Capital (reconciled, USD m)	60.00	19.57	40.43	0.46	90/59
Software revenue (USD m)	6.39	1.12	5.26	0.53	111/75
Software exports (USD m)	5.37	0.88	4.49	0.48	94/48
Reports revenue	0.82	0.75	0.07	0.18	135/100
Reports exports	0.70	0.48	0.22	0.45	135/100
Reports capital	0.67	0.59	0.08	0.16	135/100
Lists any client (guard)	0.65	0.46	0.19	0.39	135/100
No. industries served	2.08	1.31	0.77	0.50	135/100
Public limited	0.47	0.21	0.26	0.56	135/100
Private limited	0.35	0.53	-0.18	-0.37	135/100
No. foreign countries	0.35	0.00	0.35	0.32	135/100
Any foreign country	0.08	0.00	0.08	0.42	135/100
Serves banking/finance	0.54	0.20	0.34	0.75	135/100
Serves telecom	0.16	0.07	0.09	0.27	135/100
Serves manufacturing	0.48	0.24	0.24	0.52	135/100
Any certification	0.19	0.04	0.15	0.47	135/100
STPI center in district	0.67	0.60	0.07	0.14	135/100
Region: Maharashtra	0.29	0.34	-0.05	-0.11	135/100
Region: Delhi	0.23	0.25	-0.02	-0.05	135/100
Region: Karnataka	0.20	0.14	0.06	0.16	135/100

Notes: The 1995 cross-section of the estimation universe (235 firms; 135 legacy-matched, 100 comparison). *Diff.* is the legacy-matched mean minus the comparison mean. *Norm. diff.* is the Imbens–Rubin normalized difference, the mean difference divided by the pooled standard deviation $\sqrt{(s_1^2 + s_0^2)/2}$; unlike a t-test it does not shrink mechanically with sample size, and a value above roughly 0.25 in absolute value signals substantial imbalance. N_e/N_c is the number of legacy-matched and comparison firms with the covariate observed. The dollar-denominated and capital rows are conditional on reporting, so they are read jointly with the reporting-share rows.

B.2 Does the result depend on how the legacy match is measured?

No: the scale response holds under five codings of the pre-shock match, and the certification response concentrates in the coherence margin. Table B.2 re-estimates the split-window specification across the measures defined in Section 5.1 — the binary indicator, the coherence gate, the legacy-category count, the standardized coherence dose, and the coherent-layer count — for four scale outcomes under both estimators. Employment is positive under all five measures and both estimators. Because the measures carry different units and enter

separate regressions, the exercise establishes robustness to alternative measurement rather than a ranking by stack depth.

Table B.3 reports the complementarity test behind Figure 10, reparametrized so that the coherence margin carries its own clustered standard error. For certification the response is concentrated there: the incomplete-minus-no-core estimate is small and imprecise, while the coherent-minus-incomplete increment is large and distinguishable from zero. The employment and capital gradients are ordered the same way but imprecise.

Figure B.1 traces the employment paths under the two depth measures on their common sample. The dose path is positive after the deadline with a negative, imprecise 1992 coefficient; the gate path is imprecise throughout, and its 1992 coefficient is negative and marginally significant.

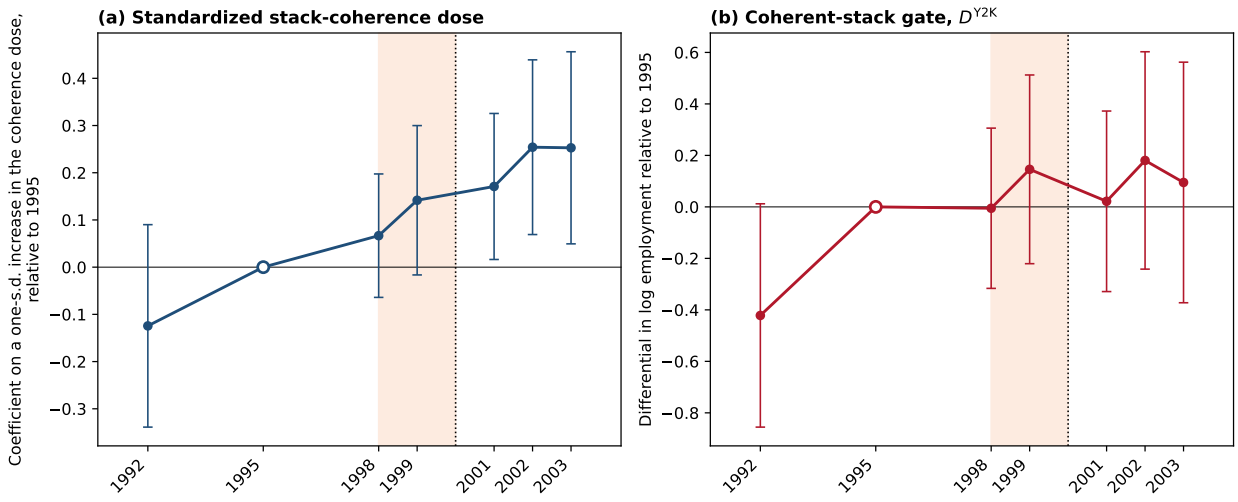


Figure B.1: Employment event studies under depth measures of the pre-shock legacy match.

Notes: Each panel plots coefficients from a separate OLS event-study regression of log employment on the indicated depth measure interacted with directory-wave indicators, relative to the omitted 1995 wave (open markers), on the common depth-measurable sample ($N = 934$ firm-years). Panel (a) uses the standardized stack-coherence dose, so the plotted coefficient corresponds to a one-standard-deviation increase. Panel (b) uses the coherent-stack gate D_i^{Y2K} and additionally includes the reporting-verbosity and frontier-breadth controls of Table B.2. All specifications include firm fixed effects, cohort-by-year fixed effects, and M&A event controls. Whiskers are 95% confidence intervals from firm-clustered standard errors. The shaded band marks the 1998–99 remediation window and the dotted line the January 2000 deadline.

Table B.2: Post-deadline scale responses across alternative measures of the legacy match

	Inputs		Outputs	
	(1) Employment	(2) Capital	(3) Software exports	(4) Revenue
Panel A. DiD estimates: OLS on log outcomes				
Binary technical-match measures				
Any legacy technology, L_i^{bin}	0.341*	0.760**	0.058	−0.198
	(0.189)	(0.358)	(0.309)	(0.280)
Coherent legacy stack, D_i^{Y2K}	0.222	0.350	0.112	−0.207
	(0.203)	(0.369)	(0.327)	(0.331)
Depth measures				
Legacy-category count	0.096**	0.151*	0.050	0.024
	(0.048)	(0.078)	(0.066)	(0.064)
Coherence dose (standardized)	0.213**	0.173	0.064	−0.085
	(0.098)	(0.165)	(0.163)	(0.157)
Coherent-layer count	0.149**	0.178*	−0.009	−0.042
	(0.066)	(0.106)	(0.102)	(0.102)
Firm-years	934	726	714	803
Firms (clusters)	195	158	152	170
Panel B. DiD estimates: PPML on outcomes in levels				
Binary technical-match measures				
Any legacy technology, L_i^{bin}	0.407	0.453	−0.093	0.197
	(0.255)	(0.287)	(0.369)	(0.332)
Coherent legacy stack, D_i^{Y2K}	0.706***	0.369	1.002***	0.660
	(0.234)	(0.413)	(0.324)	(0.407)
Depth measures				
Legacy-category count	0.137**	0.119	0.233***	0.258***
	(0.058)	(0.075)	(0.058)	(0.065)
Coherence dose (standardized)	0.367***	0.161	0.485***	0.453***
	(0.115)	(0.175)	(0.148)	(0.168)
Coherent-layer count	0.226***	0.175*	0.285***	0.329***
	(0.078)	(0.104)	(0.090)	(0.111)
Firm-years	934	726	714	803
Firms (clusters)	195	158	152	170

Notes: Each cell reports the post-deadline coefficient β_2, the 2001–03 legacy-match differential relative to the 1992–95 base, from a separate version of equation (19); the 1998–99 remediation interaction (β_1) is also estimated but not reported. The binary measures indicate any pre-shock legacy technology, L_i^{bin}, and a coherent legacy production stack, D_i^{Y2K}; the depth measures are the legacy-category count, the standardized coherence dose, and the coherent-layer count. Panel A estimates OLS on log outcomes and Panel B PPML with a log link on outcomes in levels, on identical outcome-specific observations, so differences across panels reflect estimator and functional form rather than sample composition. All specifications draw from the 225-firm universe for which stack depth is observed and include firm fixed effects, cohort-by-year fixed effects, M&A event controls, and reporting-verbosity and frontier-breadth controls; standard errors clustered by firm are in parentheses. The L_i^{bin} estimates differ from Table 5 because of the restricted sample and the additional controls. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.3: The formal complementarity test: presence versus coherence of the legacy match

	(1) Certification ladder	(2) Log employment	(3) Log capital
Presence (β: incomplete stack − no legacy core layer)			
× Remediation (1998–99, β_1)	0.012	−0.056	0.168
	(0.113)	(0.161)	(0.326)
× Post-deadline (2001–03, β_2)	0.171	0.222	0.193
	(0.235)	(0.214)	(0.372)
Coherence (δ: coherent stack − incomplete stack)			
× Remediation (1998–99, δ_1)	0.249**	0.211	0.592*
	(0.120)	(0.173)	(0.306)
× Post-deadline (2001–03, δ_2)	0.685***	0.161	0.299
	(0.249)	(0.216)	(0.381)
Coherent stack − no legacy core layer ($\beta_2 + \delta_2$)	0.856***	0.383	0.492
	(0.317)	(0.252)	(0.459)
Observations (firm-years)	961	934	726

Notes: The nested legacy-match decomposition behind Figure 10, reparametrized so the coherence margin carries its own standard error. Each column is a single regression that enters two predetermined indicators jointly: *presence*, whether a firm held any legacy core layer, and *coherence*, whether the firm passed the IBM-mainframe or AS/400 gate of Section A.7. Because a coherent stack is nested in the presence indicator, the presence coefficient β_s is the incomplete-minus-no-core-layer contrast and the coherence coefficient δ_s is the coherent-minus-incomplete contrast; their sum, in the bottom block, is the coherent-minus-no-core-layer total. The estimates come from the split-form version of equation (19) on the common 225-firm legacy-match sample, with firm and entry-cohort-by-year fixed effects, M&A event controls, and remediation- and post-window interactions for reporting verbosity, frontier breadth, and 1992 stack-field observability; standard errors clustered by firm in parentheses. The decomposition partitions firms on the predetermined technical match, so it identifies which firms later upgraded rather than certification as a causal mediator. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.4 asks what the reduced form implies per engaged firm. It reports the sample-specific first stages of reported Y2K engagement on the legacy match, the corresponding two-stage coefficients with their Anderson–Rubin sets, and the uninstrumented engagement association. The exclusion restriction required to read these as causal is implausible, since the predetermined legacy match may affect unreported projects, client acquisition, hiring, and credibility as well as the directory engagement field. The table is a scaling sensitivity, not an instrumental-variables estimate of an effect per engaged firm.

Table B.4: Reported Y2K engagement and employment: a scaling sensitivity

	Δ ln employment from the 1995 base	
	(1) To the remediation window, 1998–99	(2) To the post-deadline window, 2001–03
First stage: engagement on L_i^{bin}	0.274***	0.326***
	(0.078)	(0.086)
Robust first-stage Wald F	12.3	14.2
Reduced form: outcome on L_i^{bin}	0.214	0.575***
	(0.141)	(0.190)
Two-stage least squares: outcome on engagement	0.783	1.764**
	(0.550)	(0.763)
Anderson–Rubin 95% set	[−0.238, 2.446]	[0.574, 4.346]
OLS: outcome on engagement, uninstrumented	0.078	0.027
	(0.160)	(0.202)
Firms	173	144

Notes: Each column uses firms observed in 1995. The dependent variable is mean log employment over the firm's available observations in 1998–99 or 2001–03, minus log employment in 1995. Controls are log(1 + employment) and firm age in 1995, their missing-value indicators, private- and public-limited legal form, and indicators for Maharashtra, Delhi, and Karnataka. Standard errors in parentheses are heteroskedasticity-robust (HC1). Engagement equals one if the dedicated field or profile narrative reports Y2K work in 1998 or 1999; an explicit no and absence from both waves are coded zero. The excluded instrument is L_i^{bin}, the predetermined 1995 legacy indicator. These controlled, sample-specific first stages differ from the validation regressions in Table 4, and because the long differences use 1995 alone as the base they do not exactly rescale the pooled-pre coefficients in equation (19). In each column the two-stage coefficient equals the reduced form divided by the first stage, and the Anderson–Rubin sets are obtained by grid inversion and are robust to a weak instrument. The two-stage coefficients are scaling sensitivities rather than causal estimates per engaged firm. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.3 Which other firm margins moved?

Three margins beyond headline employment are reported here: the composition of the workforce, registered capital, and legal-entity survival. Table B.5 splits total headcount into

Table B.5: Employment growth by worker type

	Employment (log headcount)		
	(1) Total	(2) Software & eng.	(3) Other (non-tech.)
Panel A. DiD estimates: OLS on log outcomes			
Legacy match × Remediation period (1998–99, β_1)	0.172	0.130	−0.209
	(0.133)	(0.137)	(0.157)
Legacy match × Post-deadline period (2001–03, β_2)	0.425**	0.374**	0.070
	(0.172)	(0.186)	(0.193)
Difference ($\beta_2 - \beta_1$): post-deadline minus remediation	0.252**	0.244**	0.280**
	(0.116)	(0.116)	(0.141)
Firm fixed effects	✓	✓	✓
Cohort × year fixed effects	✓	✓	✓
M&A event controls	✓	✓	✓
Comparison-group mean, 1995	95	66	33
Observations (firm-years; firms)	953 (199)	875 (187)	851 (183)
Panel B. DiD estimates: PPML on outcomes in levels			
Legacy match × Post-deadline period (2001–03, β_2)	0.581***	0.654***	−0.134
	(0.218)	(0.241)	(0.338)
Observations (firm-years; firms)	953 (199)	875 (187)	851 (183)

Notes: Each column is the firm-level difference-in-differences of equation (19) on the 235-firm headline universe, with L_i^{bin} (an indicator for any pre-shock legacy technology) as the legacy-match measure. The omitted pre-period is 1992–95; β_1 and β_2 are the legacy-match differentials during the 1998–99 remediation window and after the January 1, 2000 deadline (2001–03), and the difference row is the linear contrast between them from the same regression. All specifications include firm fixed effects, cohort-by-year fixed effects, and M&A event controls; standard errors clustered by firm are in parentheses. Panel A is OLS on log headcount; Panel B is PPML on headcount in levels. The comparison-group mean is the 1995 headcount. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

software-engineering and other staff under both estimators. The post-deadline expansion is concentrated in the technical workforce, at 0.374 log points against 0.070 for other staff, with the same ordering under PPML. The Online Appendix reports the same event study run on software-only headcount, which tracks total employment, so the scale result is not an artifact of non-software staff.

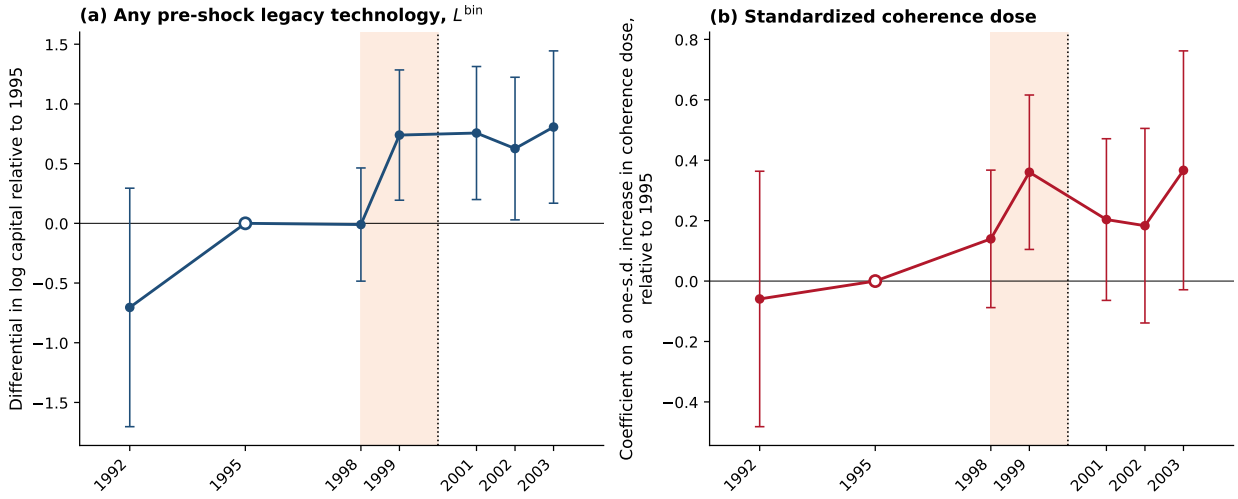


Figure B.2: Capital dynamics under two measures of the pre-shock legacy match.

Notes: Each panel plots coefficients from a separate OLS event-study regression of log capital on the indicated predetermined legacy-match measure interacted with directory-wave indicators, relative to the omitted 1995 wave. Capital is subscribed or authorized capital as reported in the directory. Panel (a) uses L_i^{bin}, an indicator for any reported pre-shock legacy technology; panel (b) uses the standardized coherence dose, so the plotted coefficient is the effect of a one-standard-deviation increase. Both panels use the same estimation sample of 726 firm-years from 158 firms. All specifications include firm fixed effects, cohort-by-year fixed effects, and M&A event controls. Whiskers are 95% confidence intervals from firm-clustered standard errors. The shaded band marks the 1998–99 remediation window and the dotted line the January 2000 deadline.

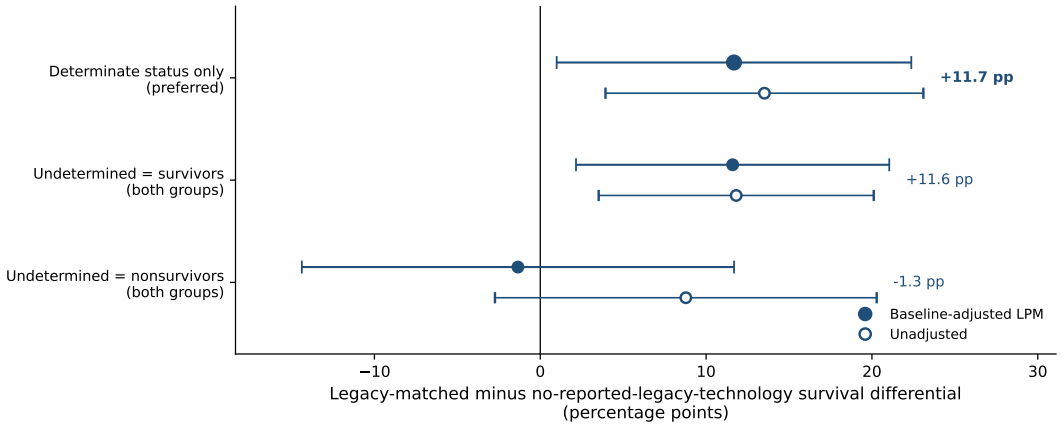


Figure B.3: Legal-entity survival of the 1995 cohort, by pre-shock legacy match.

Notes: Each marker reports the difference in the probability of legal-entity survival between firms with any reported pre-shock legacy technology, $L_i^{\text{bin}} = 1$, and firms without, in percentage points. Legal status is defined independently of continued appearance in the directory. The 1995 cohort contains 235 firms, 135 legacy-matched and 100 without, of which 198 have determinate subsequent legal status and 37 remain undetermined. Filled markers report linear probability models controlling for baseline firm size and open markers unadjusted differences in group means; the numerical labels are the baseline-adjusted estimates. The top row restricts to firms with determinate status and is the preferred specification; the remaining rows code every undetermined firm in both groups as a survivor or as a nonsurvivor. Whiskers are 95% confidence intervals from heteroskedasticity-robust standard errors. The common codings are sensitivity analyses rather than formal bounds.

B.4 Does the result depend on the estimator or the sample?

Neither. Two-way fixed-effects difference-in-differences can be biased when treatment effects are heterogeneous, because already-treated units enter the comparison group. Exposure here turns on at a single date, which rules out the staggered-adoption pathology, and Table B.6 confirms this directly by reporting the imputation, group-time, and two-way fixed-effects estimates side by side on both headline outcomes. The estimates bracket one another, and the imputation estimator returns somewhat larger post-deadline coefficients than the baseline on both outcomes. Three estimator-package conventions hold throughout and are stated once: the imputation estimator runs without singleton removal, the group-time estimator runs in panel mode, and its not-yet-defined pre-period cell is suppressed rather than reported.

Table B.6: Estimator robustness: the headline under alternative difference-in-differences estimators

Estimator	Exposure	Remediation (1998–99)	Post-deadline (2001–03)	Inference
Panel A. Log employment				
Two-way fixed effects	Legacy match	0.17 (0.13)	0.42** (0.17)	Firm-clustered; $N = 953$
Imputation	Legacy match	0.19 [−0.08, 0.45]	0.60 [0.25, 0.96]	Firm bootstrap; permutation $p = .008$
Group-time	Legacy match	0.14	0.51	134 vs. 96 firms
Two-way fixed effects	Demand intensity	0.18 (0.15)	0.43** (0.21)	Firm-clustered; $N = 669$
Imputation	Demand intensity	0.16 [−0.16, 0.46]	0.57 [0.16, 1.04]	Firm bootstrap
Group-time	Demand intensity	0.14	0.51	80 vs. 80 firms
Heterogeneity-robust	Legacy match	0.14	0.50	Placebo −0.21 (0.19)
Heterogeneity-robust	Demand intensity	0.14	0.50	Placebo −0.25 (0.19)
Heterogeneity-robust	Demand intensity, per SD	0.22	0.62	Placebo −0.14 (0.21)
Panel B. Certification ladder (0–4)				
Two-way fixed effects	Legacy match	0.32*** (0.09)	1.07*** (0.18)	Firm-clustered; $N = 984$
Imputation	Legacy match	0.29 [0.15, 0.48]	1.12 [0.78, 1.49]	Firm bootstrap; permutation $p = .001$
Group-time	Legacy match	0.26	1.11	135 vs. 100 firms
Two-way fixed effects	Demand intensity	0.17 (0.12)	0.83*** (0.22)	Firm-clustered; $N = 679$
Imputation	Demand intensity	0.16 [−0.06, 0.40]	0.88 [0.47, 1.32]	Firm bootstrap
Group-time	Demand intensity	0.12	0.76	80 vs. 81 firms
Heterogeneity-robust	Legacy match	0.26	1.13	Placebo −0.35 (0.08)
Heterogeneity-robust	Demand intensity	0.12	0.79	Placebo −0.36 (0.13)
Heterogeneity-robust	Demand intensity, per SD	*unestimable*		Collinear after listwise deletion
Firm FE, cohort × year FE, M&A controls	✓ (two-way fixed-effects rows)			

Notes: Each row re-estimates the split-window differential of equation (19) under a different estimator. Parentheses hold standard errors and brackets hold bootstrap confidence intervals; two-way fixed-effects rows cluster by firm. Two exposures are reported: the predetermined binary legacy match, and the demand-intensity measure z_i above its median among the 161 firms reporting client-industry shares. The imputation estimator is Borusyak et al. (2024), fit on untreated observations; singleton fixed effects are retained because a 1995-cohort matched firm's effect is identified from its single pre-wave. The group-time estimator is Callaway and Sant'Anna (2021) with never-treated comparison firms and wave-rank event time, since the waves are unevenly spaced. The heterogeneity-robust rows are de Chaisemartin and D'Haultfœuille (2020), with the remediation and post-deadline columns the simple means of effects one to two and three to five; their certification placebo reproduces the 1992 certification lead disclosed in Section 6.2 under an independent estimator rather than adding new information. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.7 re-runs the split-window specification changing one sample or measurement choice at a time. Panel A varies the sample, Panel B the coding of the certification outcome,

and Panel C the treatment of the reported dollar outcomes. No sample toggle flips a sign: the post-deadline employment coefficient stays between +0.28 and +0.66 and the certification coefficient between +0.78 and +1.36. The certification result also holds when the outcome is a binary on-ladder credential, when the ladder is capped at rung 2, and when the 45 firms ever observed at CMM Levels 4–5 are dropped. A twentieth toggle not in the table removes the M&A event controls, since acquisition activity may itself respond to the shock: employment is 0.435 (SE 0.174) against the headline 0.425, and certification 1.099 (SE 0.179) against 1.066, so the controls do no work the results depend on.

Table B.7: Sample and measurement toggles

	Log employment			Certification ladder		
Toggle	Remediation	Post	Obs.	Remediation	Post	Obs.
Panel A. Sample toggles						
Headline specification	0.17 (0.13)	0.42** (0.17)	953	0.32*** (0.09)	1.07*** (0.18)	984
Drop 2003 (format change)	0.18 (0.13)	0.38** (0.17)	863	0.32*** (0.09)	1.00*** (0.17)	892
Drop 2002 and 2003	0.18 (0.13)	0.32* (0.17)	751	0.31*** (0.09)	0.78*** (0.15)	777
Drop entity-resolution collision firm-years	0.18 (0.13)	0.38** (0.17)	928	0.30*** (0.09)	0.99*** (0.18)	959
Balanced panel (81 firms)	0.29 (0.18)	0.60** (0.24)	512	0.47*** (0.15)	1.36*** (0.26)	522
Drop Delhi, Mumbai, and Kolkata	0.22 (0.22)	0.66*** (0.25)	396	0.40*** (0.15)	1.26*** (0.28)	407
Drop Bangalore	0.17 (0.15)	0.39** (0.19)	787	0.34*** (0.09)	1.09*** (0.18)	814
Drop top-decile 1995 size (24 firms)	0.18 (0.14)	0.41** (0.18)	842	0.22*** (0.08)	0.90*** (0.18)	870
Drop re-entrant firms (16)	0.13 (0.14)	0.39** (0.19)	881	0.29*** (0.09)	1.05*** (0.19)	909
Plain year FE, no cohort interaction	0.17 (0.13)	0.42** (0.17)	953	0.32*** (0.09)	1.07*** (0.17)	984
Drop the 16 early certification adopters	0.13 (0.14)	0.36* (0.18)	853	0.20** (0.08)	0.87*** (0.18)	884
Exposure redefined: two or more legacy categories	0.26* (0.14)	0.43** (0.18)	953	0.39*** (0.10)	1.04*** (0.20)	984
Drop foreign-affiliated firms	0.18 (0.17)	0.28 (0.25)	588	0.34*** (0.12)	0.98*** (0.24)	603
Panel B. Coding of the certification outcome						
Binary on-ladder credential	—	—	—	0.17*** (0.06)	0.26*** (0.07)	984
Ladder capped at rung 2	—	—	—	0.25*** (0.07)	0.58*** (0.10)	984
Exclude 45 firms ever at CMM Levels 4–5	—	—	—	0.14** (0.07)	0.31*** (0.09)	701
Panel C. Measurement toggles (log software exports; log revenue)						
Headline (no winsorizing, freshness ≤ 2)	0.10 (0.29)	0.25 (0.32)	716	−0.08 (0.21)	0.06 (0.26)	814
Winsorized 1/99 within wave	0.08 (0.29)	0.20 (0.32)	716	−0.08 (0.21)	0.04 (0.26)	814
Winsorized 5/95 within wave	0.07 (0.27)	0.19 (0.30)	716	−0.12 (0.20)	0.04 (0.23)	814
Exclude projection-flagged rows	0.04 (0.28)	0.21 (0.37)	591	−0.10 (0.21)	−0.04 (0.27)	683
Freshness lag ≤ 1	−0.15 (0.31)	−0.16 (0.35)	591	−0.28 (0.19)	−0.11 (0.25)	670
No freshness mask	0.12 (0.28)	0.28 (0.32)	749			
Firm FE, cohort × year FE, M&A controls	✓ (except where the toggle states otherwise)					

Notes: Each row re-estimates the split-window specification of equation (19) on the predetermined binary legacy match, changing only the stated toggle; standard errors in parentheses are clustered by firm. Panels A and B report log employment and the 0–4 certification ladder; Panel C reports log software exports in the first three columns and log revenue in the last three, so its column headings refer to those outcomes rather than to employment and certification. Panel B changes only the certification coding, so its employment columns are empty by construction. The freshness mask in Panel C drops reported dollar figures whose printed year lags the directory year by more than the stated number of years; the sign change under the tighter lag reflects pre-period composition, since that cut removes 115 rows, 76 of them pre-period. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.5 What survives departures from parallel trends?

The certification result carries a disclosed 1992 lead, and this subsection quantifies what the lead can and cannot explain. Three exercises bear on it: rebuilding the exposure from the earlier directory alone, bounding the post-deadline effect under transparent violations of parallel trends, and isolating the margin on which the lead cannot operate at all.

Table B.8 first separates the two ways the pre-period can enter. The pooled specification combines 1992 and 1995 into one pre-period; the 1995-reference specification instead adds an explicit 1992 lead term. The employment estimate is unchanged across the two and its 1992 interaction is −0.043 (SE 0.259), while the certification estimate is modestly smaller against the disclosed 1992 certification differential.

Table B.8: Pooled pre-period versus 1995-reference specification

	ln Employment		Certification ladder	
	Pooled	1995-ref	Pooled	1995-ref
Legacy match × 1992 (lead)	—	−0.043	—	−0.341**
		(0.259)		(0.140)
Legacy match × Remediation (β_1)	0.172	0.162	0.319***	0.234***
	(0.152)	(0.142)	(0.100)	(0.088)
Legacy match × Post-deadline (β_2)	0.425**	0.414**	1.066***	0.977***
	(0.197)	(0.188)	(0.200)	(0.192)
Omitted pre-period	1992 & 1995	1995	1992 & 1995	1995
Observations	984	984	1,015	1,015

Notes: Each column estimates the two-window difference-in-differences of equation (19) on the headline L_i^{bin} exposure, with firm and cohort-by-year fixed effects, M&A controls, and firm-clustered standard errors. The *Pooled* columns omit both pre-shock waves, imposing a common 1992–1995 pre-shock legacy-match differential; the *1995-ref* columns add an explicit legacy-match × $\mathbf{1}[1992]$ lead term so the remediation and post-deadline coefficients are referenced to 1995. A negative 1992 coefficient means the differential was smaller in 1992 than in 1995. The pooled columns reproduce the two-window estimates in the body. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.9 rebuilds the legacy match from the 1992 edition alone, which is the strictest available timing check, at the cost of a 124-firm base and a first-stage t near two. Certification replicates on this base under the binary measure and under the category count. The binary employment cell is uninformative: its confidence interval covers both the 1995-base estimate and its negation, which is a power problem rather than a failed replication.

Table B.9: Replication on the 1992 exposure base

Exposure measure	Remediation window (1998–99)	Post-deadline (2001–03)	Observations	First-stage t
Panel A. Log employment				
Legacy match, binary (L_i^{bin})	0.06 (0.21)	−0.02 (0.26)	455	2.2
Legacy match, category count	0.20** (0.08)	0.25** (0.11)	455	1.6
Demand intensity (z_i)	−0.05 (0.14)	0.07 (0.18)	293	2.2
Client-industry shares only (B_i)	−0.05 (0.08)	−0.04 (0.15)	293	2.1
Triangulation composite	0.21** (0.11)	0.24* (0.14)	455	2.4
Panel B. Certification ladder (0–4)				
Legacy match, binary (L_i^{bin})	0.30** (0.15)	0.67** (0.31)	464	2.2
Legacy match, category count	0.25*** (0.07)	0.61*** (0.16)	464	1.6
Demand intensity (z_i)	0.18*** (0.07)	0.29* (0.17)	299	2.2
Client-industry shares only (B_i)	−0.07* (0.04)	−0.14 (0.13)	299	2.1
Triangulation composite	0.16* (0.08)	0.52*** (0.17)	464	2.4
Firm FE, cohort × year FE, M&A controls	✓	✓		

Notes: Each row is one estimate of the split-window specification in equation (19), with exposure rebuilt entirely from the 1992 directory rather than from 1995. All specifications include firm fixed effects, entry-cohort-by-year fixed effects, and M&A event controls; standard errors in parentheses are clustered by firm. The 1992 base contains 124 firms and is never pooled with the 1995 base. Continuous measures enter per standard deviation. The exercise is underpowered by construction: the 1992 cohort is 91 firms and the first-stage t on reported Y2K engagement is about two throughout. The binary employment confidence interval, $[-0.54, +0.50]$, covers both the 1995-base estimate of +0.43 and its negation. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure B.4 reports robust confidence sets for the post-deadline average certification effect under the relative-magnitudes restriction of Rambachan and Roth (2023), which bounds the post-period violation of parallel trends between consecutive directory waves by $\bar{M}$ times the observed pre-period deviation. The effect survives $\bar{M} = 0.25$ and is no longer distinguishable from zero at $\bar{M} = 0.5$. Two readings of that number should be kept apart. The wave-step scaling is harsh, because the one observed pre-period step spans three calendar years while the late post-period steps span one, so at $\bar{M} = 1$ the restriction grants a confound as much drift per single post year as it accumulated over the whole 1992–95 window. Rescaled per calendar year, the point estimate survives drift up to 1.3 times the observed pre-period slope. The exact linear extrapolation of the lead, the harshest smooth counterfactual, leaves +0.24 with an interval that includes zero.

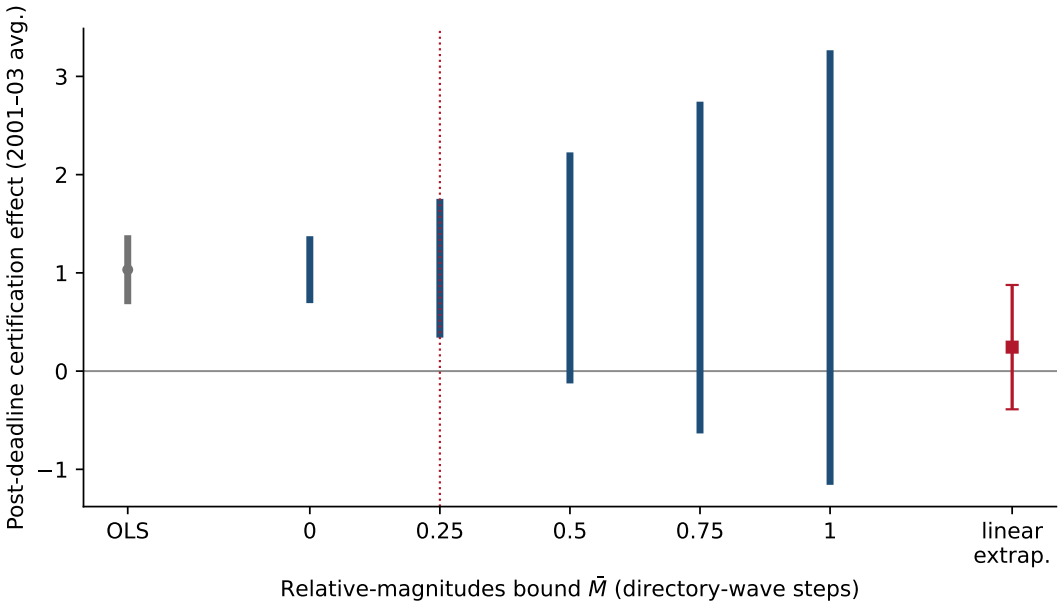


Figure B.4: Honest-trends bounds on the post-deadline certification effect.

Notes: Robust 95% confidence sets (Rambachan and Roth 2023) for the post-deadline average certification effect over 2001–03, allowing post-period violations of parallel trends up to $\bar{M}$ times the observed 1992 lead per directory-wave step under the relative-magnitudes restriction, with the conditional–least-favorable hybrid and one pre-period. The leftmost gray interval is the unadjusted estimate. The set excludes zero at $\bar{M} = 0.25$ and includes it from $\bar{M} = 0.5$, widening monotonically over the grid to $\bar{M} = 2$. The red marker is the delta-method interval after subtracting the exact linear calendar-time extrapolation of the lead. The wave-step scaling is conservative here, because the single pre-period step spans three calendar years against one for the late post-period steps.

The bounds discipline the level of the effect; Table B.10 shows the margin the lead cannot touch. Restricting to the 206 firms with no certification at all in 1995 removes the early adopters and, because the ladder rarely reverts, removes the pre-deadline differential by construction: the 1992 coefficient in that subset is −0.020. The post-deadline response there is 0.783 ladder points and survives size-quintile-by-year fixed effects at 0.532. The CMM margin among firms with no CMM in 1995, and the full sample under baseline-certification-by-year fixed effects, give the same answer. A continuation of the early adopters' trend cannot produce differential first-time adoption among firms that had never certified.

Table B.10: New adoption: the certification response among firms uncertified before the shock

	Legacy match × Shock	Legacy match × Post	Observations
Panel A. Firms with no certification in 1995 (certification ladder; 206 firms)			
Baseline specification	0.210**	0.783***	839
	(0.084)	(0.191)	
With size-quintile × year fixed effects	0.107	0.532***	838
	(0.090)	(0.202)	
Panel B. Firms with no CMM in 1995 (indicator for any CMM; 232 firms)			
Baseline specification	0.073**	0.326***	969
	(0.028)	(0.057)	
With size-quintile × year fixed effects	0.048*	0.175***	968
	(0.028)	(0.060)	
Panel C. Full sample with baseline-certification × year fixed effects (235 firms)			
Baseline-certification × year	0.223***	0.826***	973
	(0.084)	(0.190)	
Adding size-quintile × year	0.115	0.570***	972
	(0.090)	(0.198)	

Notes: Each row is the split-window form of equation (19) on the indicated sample and outcome, with firm and cohort-by-year fixed effects, M&A event controls, and standard errors clustered on the firm in parentheses. Panel A restricts to the 206 firms (110 legacy-matched, 96 comparison) with no certification at the 1995 wave; because the ladder rarely reverts, this subset has no pre-deadline certification by construction (pre-period mean 0.000), so the margin is free of the disclosed 1992 lead — the estimated 1992 event-study coefficient in the subset is −0.020 (0.106). Panel B repeats the exercise on the CMM margin among firms with no CMM in 1995, and Panel C keeps the full sample and absorbs baseline-certification-by-year fixed effects instead of restricting. The lead-free property of Panels A and B follows from the restriction and is not an independent pre-trend test. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.6 The Euro-conversion comparison

A second deadline-driven shock, the Euro conversion, is reported here as texture rather than as a second design. Figure B.5 plots the certification paths under the predetermined Y2K legacy match and under self-reported Euro-conversion work on the same axes. The Y2K path bends at its own deadline. The Euro path rises through the technology bust, which a boom-driven account does not predict, but Euro-exposed firms are larger and the rise is not separable from size composition, so Section 8 reads it as qualitative texture.

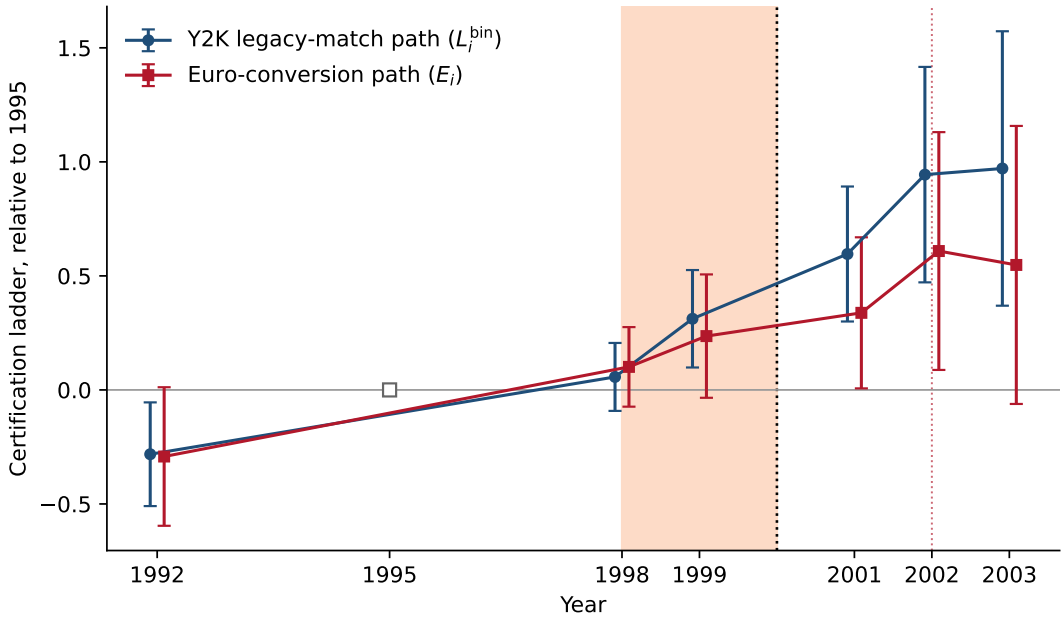


Figure B.5: Certification paths under the Y2K legacy match and Euro-conversion exposure.

Notes: Joint event-study paths of the 0–4 certification ladder on the predetermined legacy match L_i^{bin} and on self-reported Euro-conversion work, relative to the omitted 1995 wave, with firm and cohort-by-year fixed effects, M&A event controls, and firm-clustered 95% intervals. The shaded band marks the 1998–99 remediation window, the dotted black line the January 2000 Y2K deadline, and the dotted red line the January 2002 Euro cash changeover.

B.7 Does directory attrition drive the result?

It does not, under each of the four corrections reported in Table B.11, though none of them addresses selection on unobservables. Panel A reports Lee bounds. Differential attrition on employment is 0.3 percentage points, which implies a trim of less than one observation, so the bounds are degenerate at the point estimate; on certification the trim is 4.9 percent and the bounds are tight around the point estimate. The monotonicity assumption deserves one caveat: the near-equal pooled employment observation rates conceal tercile gaps of opposite sign, so the pooled bounds are the aggregate discipline and the size composition is handled by the scale gate in Table B.13. Panel B reports two Heckman variants, both disclosed as fragile — the exclusion candidate plausibly enters the outcome equation, and the no-exclusion variant is identified off functional-form nonlinearity alone — and neither is used to recover the main estimand. Their substantive content is that the corrected and uncorrected coefficients agree to two decimals while the selection term itself is significant: selection exists and is orthogonal to the exposure effect. Panel C reweights by the inverse probability of directory reappearance and Panel D restricts to firms with a surviving legal

entity, excludes confirmed exits, and excludes absorbed-or-renamed firms in turn. All four leave the post-deadline differentials positive.

Table B.11: Selection and attrition corrections

	Log employment	Certification ladder	Trim / diagnostic
Panel A. Lee bounds on the long difference			
Point estimate	0.500	1.045	
Bounds	[0.500, 0.500]	[0.937, 1.115]	
Trim share (observations)	0.3% (0)	4.9% (4)	
Any-post presence, matched vs. comparison	0.64 vs. 0.64	0.65 vs. 0.62	
Panel B. Heckman two-step, both variants disclosed as fragile			
Corrected, with exclusion restriction	0.579 (0.209)	0.643 (0.193)	λ: $t = -2.9$ / -1.7; $\rho = -0.67$ / -0.39
Corrected, no exclusion restriction	0.556 (0.535)	0.720 (0.433)	ρ pins at 1.00
Uncorrected long difference	0.580	0.652	
Panel C. Inverse-probability weighting for directory reappearance			
Unweighted (headline)	0.42** (0.17)	1.07*** (0.18)	$N = 953$ / 984
Reappearance-weighted	0.38** (0.18)	1.05*** (0.18)	$N = 953$ / 984
Panel D. Survival-typology restrictions (post-deadline)			
All firms (headline)	0.42** (0.17)	1.07*** (0.18)	230 / 235 firms
Surviving legal entities only	0.39** (0.18)	1.13*** (0.19)	173 / 174 firms
Excluding confirmed exits	0.33* (0.18)	1.07*** (0.18)	208 / 211 firms
Excluding absorbed or renamed firms	0.55*** (0.18)	1.12*** (0.19)	193 / 198 firms

Notes: Each cell corrects the post-deadline differential of equation (19) for a different treatment of directory attrition, with standard errors in parentheses clustered by firm. Selection is having any post-2001 directory observation with a non-missing outcome. Panel A reports Lee bounds under the assumption that exposure moves selection in one direction only; the trim column gives the implied trim share and the number of observations it removes. Panel B reports the Heckman two-step in two variants, one using legal-entity survival as the exclusion restriction and one identified off functional-form nonlinearity alone; the diagnostic column carries the inverse Mills t-statistics and the estimated ρ for the employment and certification equations respectively. Panel C reweights by the inverse probability of directory reappearance, and Panel D restricts the sample by researched survival status, with firm counts reported for the employment and certification samples respectively. Panels A and B are cross-sectional long differences with 1995 controls rather than the panel post-deadline coefficient, so within those panels the informative comparison is corrected against uncorrected rather than against the headline. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.8 Observable prior advantage, reach, and scale

This subsection reports the full versions of the three exercises in Section 7. Table B.12 gives the complete legacy-match-by-reach quadrant behind Figure 12: the post-deadline and remediation-window coefficients for every cell under both match measures, the reparametrized differences between limited- and broad-reach legacy-matched firms with their permutation p-values, the continuous match-by-access specification, and the within-matched-firm limited-reach split.

Table B.12: Certification and employment by legacy match×pre-shock access cell

	Certification ladder	ln employment
Panel A: legacy-match indicator L_i^{bin} × foreign access (base: neither)		
Legacy match only (L_i^{bin}, low access) × Remediation	0.273** (0.125)	0.207 (0.246)
Legacy match only (L_i^{bin}, low access) × Post	1.316*** (0.272)	0.453 (0.305)
Both (L_i^{bin}, high access) × Remediation	0.446*** (0.094)	0.150 (0.168)
Both (L_i^{bin}, high access) × Post	1.295*** (0.193)	0.313 (0.210)
Access only (no L_i^{bin}, high access) × Remediation	0.149 (0.109)	−0.018 (0.185)
Access only (no L_i^{bin}, high access) × Post	0.414* (0.214)	−0.155 (0.260)
Legacy match only − both [Post]	0.021 (0.286)	0.141 (0.284)
Legacy match only − access only [Post]	0.902*** (0.301)	0.608* (0.323)
RI p: legacy match only (1-sided); vs. both (2-sided)	0.001; 0.948	
Cell firms (pre+post identifying)	legacy match only 42 (19); both 93 (69); access only 55 (34)	
Panel B: broad technical-match index × access index (base: neither)		
Technical match only (high index, low access) × Remediation	0.333*** (0.108)	0.262 (0.181)
Technical match only (high index, low access) × Post	1.158*** (0.224)	0.356 (0.249)
Both (high index, high access) × Remediation	0.512*** (0.121)	0.378* (0.194)
Both (high index, high access) × Post	1.417*** (0.218)	0.769*** (0.242)
Access only (low index, high access) × Remediation	0.150* (0.090)	0.292 (0.214)
Access only (low index, high access) × Post	0.211 (0.182)	0.408 (0.290)
Technical match only − both [Post]	−0.259 (0.287)	−0.414** (0.201)
Technical match only − access only [Post]	0.947*** (0.252)	−0.052 (0.258)
RI p: technical match only (1-sided); vs. both (2-sided)	0.001; 0.375	
Cell firms (pre+post identifying)	technical match only 55 (37); both 61 (43); access only 50 (33)	
Panel C: continuous technical match, access, and their interaction (per SD, × Post)		
L_i^{bin} system: legacy match × Post	0.546*** (0.093)	0.232** (0.094)
L_i^{bin} system: access × Post	0.169* (0.096)	−0.013 (0.071)
L_i^{bin} system: legacy match × access × Post	0.003 (0.085)	0.035 (0.072)
Index system: technical match × Post	0.622*** (0.083)	0.124 (0.094)
Index system: access × Post	0.191** (0.085)	0.158* (0.086)
Index system: technical match × access × Post	0.055 (0.046)	−0.017 (0.071)
Panel D: within legacy-matched firms, low-access split × Post		
L_i^{bin} system (42 of 135 low-access)	0.080 (0.289)	0.137 (0.280)
Index system (64 of 135 low-access)	−0.302 (0.261)	−0.522** (0.222)

Notes: Cell-saturated split-window difference-in-differences (eq. (19)) on the 235-firm universe; firm and entry-cohort-by-year fixed effects, M&A controls, standard errors clustered by firm. Panel A uses L_i^{bin} and the union count of foreign markets (> 1 defines high access, not the panel's branches-only field); Panel B uses seven-component technical-match and five-component access indexes split at their 1995 medians. In both systems "access" is realized pre-shock foreign engagement, a proxy for prior market access rather than a constraint measure, and the two systems' cell assignments agree for 68.9% of firms. The match-only-minus-both row is the reparametrized post-deadline difference with its own clustered t; permutation inference permutes the firm-level quadrant-label tuple (999 draws, add-one p). Post coefficients are within-firm and survivor-conditional, and directory-presence attrition is higher in the legacy-matched limited-reach cell (45% against 74%); crediting non-returning firms with no gain leaves the cell effect at +0.60 and +0.78. The employment column is reported for completeness: its apparent access gradient appears in the index system only and does not survive the drop-early-adopter and pre-trend checks. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure B.6 reports the foreign-presence transition summarized in Section 7.2. Among firms observed in both 1998 and 2001 that reported no physical foreign presence in 1998, legacy-matched firms were more likely to report one by 2001, and the rate is the same for limited- and broad-reach legacy-matched firms. Foreign presence is a durable participation state, the cells are small, and the sample conditions on reappearing in both waves, so the figure corroborates the reach pattern rather than identifying deadline-induced buyer search.

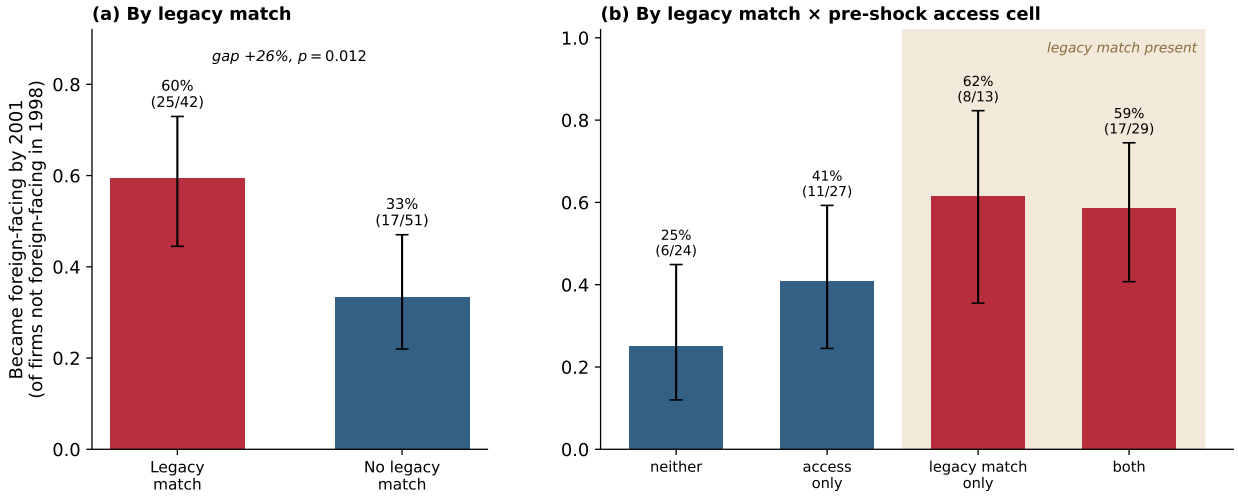


Figure B.6: Legacy-matched firms acquired foreign-facing presence during the boom across both prior-reach groups.

Notes: Among firms present in both the 1998 and 2001 directory waves and *not* foreign-facing in 1998, the share that had become foreign-facing by 2001. Foreign-facing means listing at least one foreign country in the directory's physical-presence fields, a stricter object than the 1995 market-reach measure used in Figure 12. Panel (a) splits by the legacy-match indicator L_i^{bin} and panel (b) by the legacy-match-by-prior-reach cell of Figure 12. Whiskers are Wilson 95% intervals and the legacy-match gap test is a firm-level two-proportion z-test. Cell counts appear in the figure. The sample is small and survivor-conditional, so the figure documents the transition pattern rather than a buyer-search mechanism or a separate treatment effect.

Table B.13 reports the scale gate. Panel A re-estimates the legacy-match differential within each 1995-size tercile, Panel B enters the match jointly with standardized 1995 employment and then absorbs size flexibly, and Panel C reports within-size permutation inference, the certification deadline kink, and directory-presence attrition by tercile. The attrition rows locate the one place where directory presence differs materially by exposure: the small tercile, where 48% of matched firms and 65% of comparison firms appear in any post-2001 wave.

Table B.13: The scale gate: the legacy-match effect within 1995-size terciles

	Remediation window	Post-deadline	N	Firms
Panel A: legacy-match (L_i^{bin}) effect within size terciles				
Employment, small (< 35)	0.061 (0.311)	0.498 (0.482)	282	75
Employment, mid (35–125)	0.238 (0.240)	0.422 (0.261)	299	75
Employment, large (> 125)	0.327 (0.250)	0.500* (0.285)	363	76
Certification, small (< 35)	−0.026 (0.138)	0.364 (0.387)	288	75
Certification, mid (35–125)	0.095 (0.147)	0.416* (0.243)	306	75
Certification, large (> 125)	0.289 (0.219)	1.005** (0.412)	369	76
Panel B: joint legacy-match-and-size specification (post-deadline)				
Employment: $L_i^{\text{bin}}\times$ Post		0.473** (0.194)	953	
Employment: firm size × Post		−0.045 (0.076)	953	
Employment: $L_i^{\text{bin}}\times$ Post, size×year FE		0.482** (0.186)	952	
Certification: $L_i^{\text{bin}}\times$ Post		0.624*** (0.192)	984	
Certification: firm size × Post		0.405*** (0.069)	984	
Certification: $L_i^{\text{bin}}\times$ Post, size×year FE		0.680*** (0.206)	983	
Panel C: within-size inference, deadline kink, and attrition				
Within-size-tercile RI: employment post	obs. 0.425	$p = 0.020$	999	
Within-size-tercile RI: certification post	obs. 1.066	$p = 0.001$	999	
Certification deadline kink (headline)		0.156*** (0.055)	984	
Certification deadline kink (size×year FE)		0.106* (0.062)	983	
Any-post presence, small tercile	comparison 0.65	exposed 0.48		
Any-post presence, mid tercile	comparison 0.60	exposed 0.62		
Any-post presence, large tercile	comparison 0.69	exposed 0.77		

Notes: Size is predetermined 1995 employment; terciles cut at 35 and 125 employees on the 235-firm universe (the 9 missing-size firms form their own stratum, excluded from the per-tercile splits in Panel A). Panel A re-estimates the split-window difference-in-differences (eq. (19)) on the L_i^{bin} indicator within each size tercile. Panel B enters L_i^{bin} jointly with standardized 1995 log employment, each interacted with both windows (firm-size×year fixed-effects rows absorb size flexibly). Panel C reports within-size-tercile permutation inference (999 draws, exposure permuted within size strata), the certification deadline-kink test (exposure×linear trend plus exposure×post-2000 kink, at tercile granularity), and the share of firms in each tercile×exposure group present in any post-2001 wave. All regressions include firm and entry-cohort-by-year fixed effects and M&A controls; standard errors clustered by firm in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.14 enters firms' own reported connectivity and the legacy match in the same regression. Predetermined own links strongly predict subsequent certification, but their 1992 coefficients show that connected and unconnected firms were already on different paths before the remediation market opened. When the two enter together the legacy coefficient declines from 0.955 to between 0.796 and 0.817 and remains precise. Connectivity therefore marks firms that were already ahead rather than measuring what infrastructure contributed. The Online Appendix reports the firm-level access-margin levels, the take-up paths behind them, and the STPI rollout chronology and timing diagnostics.

Table B.14: Own connectivity predicts certification, but does not absorb the predetermined legacy match

	Legacy match	Any reported link		International link	
	Only	Link only	Link + match	Link only	Link + match
	(1)	(2)	(3)	(4)	(5)
Panel A. Coefficients on the 1995 own-link flag					
Own link × 1992 lead		−0.395**	−0.344**	−0.392	−0.306
		(0.167)	(0.133)	(0.229)	(0.199)
Own link × Shock, 1998–99		0.337***	0.309***	0.382***	0.347***
		(0.072)	(0.056)	(0.110)	(0.097)
Own link × Post, ≥2001		0.856***	0.653***	0.786***	0.567***
		(0.096)	(0.099)	(0.198)	(0.173)
Panel B. Coefficients on the predetermined legacy match					
L_i^{bin}× 1992 lead	−0.338***		−0.263***		−0.244***
	(0.105)		(0.086)		(0.066)
L_i^{bin}× Shock, 1998–99	0.219**		0.162***		0.144***
	(0.077)		(0.055)		(0.049)
L_i^{bin}× Post, ≥2001	0.955***		0.796***		0.817***
	(0.120)		(0.099)		(0.098)
Firm FE, cohort × year FE, M&A	✓	✓	✓	✓	✓
Observations	971	971	971	971	971

Notes: The outcome is the 0–4 certification ladder. Each column is one regression estimated by OLS with firm fixed effects, entry-cohort-by-year fixed effects, and M&A controls; standard errors in parentheses are clustered by metro unit. The own-link variables are predetermined 1995 indicators for whether the directory reports any data link (106 of the 235 exposure-universe firms) or the stricter international link (65); they measure reported connectivity, not the pre-shock foreign-market access of Section 7. Columns (3) and (5) enter own connectivity and the legacy match together, each with the full set of own-link-by-time and L_i^{bin}-by-time interactions, and report both coefficient blocks from the same regression. The 1992 row is a lead relative to the omitted 1995 wave: a negative coefficient means the differential had already widened between 1992 and 1995, before the remediation market existed. The own-link flags are not randomly assigned and bundle connectivity with pre-existing firm scale and internationalization, so the table is descriptive. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.15: Audited replication on the matched Prowess sample (reduced-form split)

	(1) ln sales	(2) ln export earnings	(3) ln gross fixed assets	(4) PAT (USD m, level)
Panel A. 1995 legacy match (L_i^{bin})				
Legacy match × Shock	−0.92	1.52***	−0.05	6.88
	(0.93)	(0.54)	(0.54)	(7.42)
Legacy match × Post	−0.77	1.28*	0.63	22.15*
	(1.21)	(0.71)	(0.69)	(12.03)
Observations	135	119	135	135
Panel B. Demand intensity (z_i, per standard deviation)				
Exposure × Shock	−0.41	0.51	0.08	2.33
	(0.46)	(0.43)	(0.26)	(3.25)
Exposure × Post	−0.04	0.54	0.80***	13.98
	(0.62)	(0.38)	(0.20)	(9.30)
Observations	93	84	93	93
Firm and cohort × year FE	✓	✓	✓	✓

Notes: The reduced-form shock-and-post split of Table 5, re-estimated on audited Prowess financial statements for the 32 NASSCOM firms matched to Prowess (23 legacy-matched, 9 comparison). Sales, export earnings, and gross fixed assets are in logs; profit after tax is in levels because losses are real. Panel A uses the 1995 legacy-match binary and Panel B the demand-intensity measure z_i, restated per firm-level standard deviation (multiplier 0.0032) so that its coefficients are comparable to the per-standard-deviation conventions used elsewhere. All specifications include firm and entry-cohort-by-year fixed effects; standard errors clustered by firm are in parentheses. The comparison group is nine firms, so the table confirms direction rather than supplying an independent magnitude. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.9 Audited accounts

Three audited-accounts exhibits corroborate the directory results under assumptions and coverage that differ from the directory design's. Table B.15 re-estimates the reduced-form split on audited Prowess accounts for the matched firms. Audited export earnings reproduce the foreign-revenue growth the directory under-measures, and gross fixed assets deepen after the deadline. The comparison group is nine firms, so the annual paths behind this table are noisy and the pooled split is the cleaner statement; the Online Appendix reports those annual paths, the wave-grain matched specifications, and the additional investment and input margins. Tata Consultancy Services, the industry's largest exporter, operated as a division of unlisted Tata Sons until 2004 and has no separately audited statements in the sample window, so it is absent from every matched frame.

Table B.16 reports the cross-industry comparison behind Figure 11, including the al-

ternative comparison groups and the pre-trend accounting. A linear continuation of the pre-1997 export-earnings trend accounts for roughly 38% of the later gap, against roughly 99% of the sales gap, and the export path is similar when software and IT-enabled-services firms are compared with trade-insulated non-tradable services and with other services. The broad comparison is also exposed to the 1991-onward trade-liberalization reforms, which is why the trade-insulated comparison is reported beside it.

Table B.16: The IT boom in the full Prowess universe: cross-industry difference-in-differences

Panel A: software/ITES vs. firms outside IT, hardware, and finance

	Sales (1)	Export earnings (2)	Gross fixed assets (3)	Compen-sation (4)	Total assets (5)	Export activity proxy (6)
IT × Shock (1998–99)	0.801***	1.059***	0.470***	0.609***	0.662***	0.062**
	(0.129)	(0.172)	(0.127)	(0.138)	(0.110)	(0.031)
IT × 2000	1.513***	1.658***	1.098***	1.362***	1.312***	0.146***
	(0.147)	(0.203)	(0.148)	(0.154)	(0.125)	(0.036)
IT × Post (2001–04)	1.391***	1.457***	1.318***	1.302***	1.483***	0.155***
	(0.149)	(0.209)	(0.147)	(0.165)	(0.127)	(0.032)
Acceleration (Post − Shock)	0.590***	0.398***	0.848***	0.693***	0.820***	0.093***
	(0.109)	(0.152)	(0.078)	(0.104)	(0.081)	(0.025)
Mean of Y (pre-period)	1.66	−0.44	0.96	−1.11	1.87	0.48
Observations	84,435	37,735	90,548	84,680	94,874	95,260

Panel B: the comparison-group consistency test (sales and export earnings)

	Sales			Export earnings		
	Outside IT, hardware, finance	Non-tradable services	Other services	Outside IT, hardware, finance	Non-tradable services	Other services
IT × Shock (1998–99)	0.801***	0.851***	0.750***	1.059***	1.224***	1.060***
	(0.129)	(0.136)	(0.137)	(0.172)	(0.187)	(0.197)
IT × Post (2001–04)	1.391***	1.456***	1.260***	1.457***	1.787***	1.574***
	(0.149)	(0.160)	(0.165)	(0.209)	(0.227)	(0.236)
Observations	84,435	19,214	10,898	37,735	6,428	4,873
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓

Notes: Each column is a separate regression on the full Prowess universe of Indian listed firms and larger unlisted firms filing audited accounts (fiscal years 1990–2004), with firm and year fixed effects and standard errors clustered by firm in parentheses. Treatment is the predetermined software and IT-enabled-services classification, with incumbents' labels frozen at March 1998 and entrants' at Prowess entry, and hardware excluded from both sides. Outcomes (1)–(5) are logs of USD millions among positive values. Column (6) equals one when Prowess records positive export earnings and zero when the field is zero or unreported in an otherwise observed statement; it is a rough extensive-margin proxy, and it records a smaller decline among software/ITES firms (raw shares .548 → .487, against .475 → .323). IT × 2000 is the first fully post-deadline fiscal year, which the NASSCOM waves skip, and the acceleration row reparametrizes Post − Shock. Because treatment varies at the industry level with only two treated industries, industry-clustered standard errors are unreliable and firm clustering is reported throughout. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.17 reports the entrants-inclusive within-IT regional design and the permutation reference test from which every significance assessment quoted in Section 6.3 is drawn. The design runs on roughly sixteen effective geographic cells, which is why the permutation

reference governs rather than the district-clustered standard errors, and high-exposure cells concentrate in the southern technology hubs, so a hub-favoring shock can reproduce the gradient.

B.10 Local formal education supply

This appendix investigates whether the capability-gated Y2K demand shock produced a downstream response in local formal education. I find no robust, deadline-specific acceleration in recognized school entry in the districts where IT activity was concentrated. The timing makes newly founded formal institutions an implausible major source of the immediate labor response, and the estimates weigh against a large, geographically concentrated entry response on the measured margin.

Returns, opportunity costs, and access. Labor-demand shocks create two potentially opposing incentives for education. Higher skill premia increase the return to acquiring education, while better current employment opportunities increase the opportunity cost of remaining in school. Which force dominates depends partly on whether students can access the new jobs and how much additional education those jobs require. Export-manufacturing jobs in Mexico increased dropout because adolescents could take them before completing school (Atkin 2016). In Bangladesh, proximity to garment jobs increased schooling among younger girls even as older girls entered employment (Heath and Mobarak 2015). Randomized recruitment for business-process jobs and the arrival of IT-enabled service centers also increased schooling or training in India (Jensen 2012; Oster and Steinberg 2013). A job that students can take immediately can raise the opportunity cost of attendance; a job that rewards an attainable increment of schooling can instead raise educational investment. Both responses are constrained when the necessary qualification cannot be obtained within the opportunity's horizon.

Y2K combined strong demand for specialized technical skills with a short, fixed deadline. Because acquiring the relevant formal skills generally took longer than the remaining window, institutions founded after remediation demand accelerated in 1998 could not plausibly have expanded the supply of suitably trained workers at scale before January 2000. The immediate labor response therefore had to draw primarily on skills already in place. School entry nevertheless provides a downstream test: providers could have expanded if they expected the boom to raise returns to technical education beyond the deadline.

Table B.17: Local pre-period IT presence and the Y2K window: the within-IT regional design

Panel A: headline exposure — local IT share measured through fiscal 1996 (count-weighted, per SD)

	Sales (1)	Export earnings (2)	Gross fixed assets (3)	Capex (4)	Compen-sation (5)	Export-activity proxy (6)
Share × Shock (1998–99)	−0.060	0.139	0.050	0.132	0.133	0.016
	(0.107)	(0.119)	(0.097)	(0.116)	(0.167)	(0.024)
	[.623]	[.351]	[.709]	[.246]	[.435]	[.429]
Share × 2000	0.212	0.348	0.284	0.269	0.329	0.024
	(0.163)	(0.195)	(0.109)	(0.232)	(0.225)	(0.030)
Share × Post (2001–04)	0.126	0.314**	0.294	−0.035	0.331	0.056**
	(0.134)	(0.142)	(0.112)	(0.142)	(0.191)	(0.025)
	[.445]	[.041]	[.111]	[.749]	[.122]	[.034]
Acceleration (Post − Shock)	0.187**	0.175*	0.243**	−0.167	0.198**	0.040**
	(0.070)	(0.097)	(0.063)	(0.118)	(0.076)	(0.011)
	[.042]	[.089]	[.034]	[.270]	[.030]	[.016]
Mean of Y (pre-period, levels)	9.74	7.85	4.08	1.79	1.99	0.55
Observations	2,322	1,309	2,468	793	2,344	2,659
Firms	510	292	533	224	508	573
District cells	29	21	32	22	30	34

Panel B: the export-earnings post coefficient across design variants

	≤fy1995 shares	Directory-location override	All-firms denomi-nator	Hub cells (≥50 firms)	Compensation-weighted	Incumbents only
Share × Post (2001–04)	0.317**	0.252	0.314**	0.406**	0.238	0.256
	(0.152)	(0.136)	(0.141)	(0.110)	(0.111)	(0.147)
	[.013]	[.245]	[.040]	[.023]	[.564]	[.139]
Observations	1,309	1,309	1,309	1,281	1,309	630
Firms	292	292	292	284	292	86
District cells	21	22	21	14	21	15
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓

Notes: Each column is a separate regression on the predetermined-IT Prowess universe (software and IT-enabled services, entrants included; fiscal years 1990–2004), with firm and year fixed effects. The exposure is the district cell's pre-period IT presence — IT firms with audited sales through fiscal 1996 over non-finance firms, leaving the firm's own presence out — built from Prowess alone and scaled per standard deviation of the firm-level share. Districts are 2019 registered-office locations harmonized to undivided-2001 states with metro collapses, which agrees 71 percent with the 1990s directory locations of the matched firms. Standard errors in parentheses are clustered at the district cell. Brackets report two-sided p-values from 999 joint permutations of the pre-period cell bundles, each draw rebuilding the leave-own-out exposure; stars follow that permutation reference test under cell-bundle exchangeability rather than design-based randomization inference. The fiscal-2000 row carries no bracket because the test covers the shock, post, and acceleration parameters jointly. Outcomes (1)–(5) are logs of USD millions among positive values, and the pre-period mean row reports fiscal 1990–97 levels on each column's fitted sample. Column (6) equals one when export earnings are positive and zero when explicitly zero or unreported, a rough extensive-margin proxy that does not identify export entry. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Aggregate entry and the geographic test. The national pattern makes such a response plausible. Outside Uttar Pradesh, private-unaided secondary founding rose from an annual average of 2,653 schools in 1995–97 to 4,261 in 2000–02, an increase of 61 percent, while public and aided entry shows no comparable step. The geographic design below asks whether that expansion was disproportionately concentrated in IT-exposed districts.

Let F_{dt} denote the number of secondary schools reporting founding in district d and year t, and let E_d be the district's standardized IT-employment share in the 1998 Economic Census. I estimate

$$\mathbb{E}[F_{dt}] = \exp\left(\sum_{k\neq 1998} \theta_k E_d \mathbf{1}[t=k] + \alpha_d + \delta_{s(d),t}\right), \tag{21}$$

where α_d are district effects and $\delta_{s(d),t}$ are state-by-year effects. The latter absorb the national expansion and state-specific changes in school policy. This is not the paper's legacy-match design: E_d records where IT employment was located as the remediation window opened. The estimand is a relative tilt in recognized school entry toward those districts.

The comparison is thinner than the district count suggests. Five districts account for 73 percent of the variance in E_d. Therefore, I supplement district-clustered inference with a sensitivity exercise that globally reassigns exposure across districts. Because exposure was not assigned and exchangeability is not guaranteed, this exercise diagnoses how much of an estimate the design's geographic leverage can generate; it is not exact randomization inference. The preferred sample excludes Uttar Pradesh. An administrative recognition wave makes the state 53 percent of all-India private-secondary foundings in the 1998 reference year, compared with 11–15 percent in adjacent years. Including Uttar Pradesh produces the same qualitative pattern.

Figure B.7 presents the evidence in two steps. Panel (a) shows the aggregate pattern that motivated the investigation: private-secondary founding rose sharply while public and aided founding did not. Panels (b) and (c) then plot the adjusted district-exposure coefficients in their original units. Recognized private-secondary and English-medium private-secondary entry show no stable post-deadline acceleration. For private-secondary entry, the 1999 and 2000 coefficients are -0.012 and -0.007 log points per standard deviation of exposure; for English-medium entry, they are -0.008 and -0.003. Later estimates change sign. Every displayed coefficient lies within its corresponding exposure-reassignment reference range, and all post-1998 reassignment p-values exceed 0.25.

The clustered and reassignment-based joint pre-period diagnostics disagree, so these

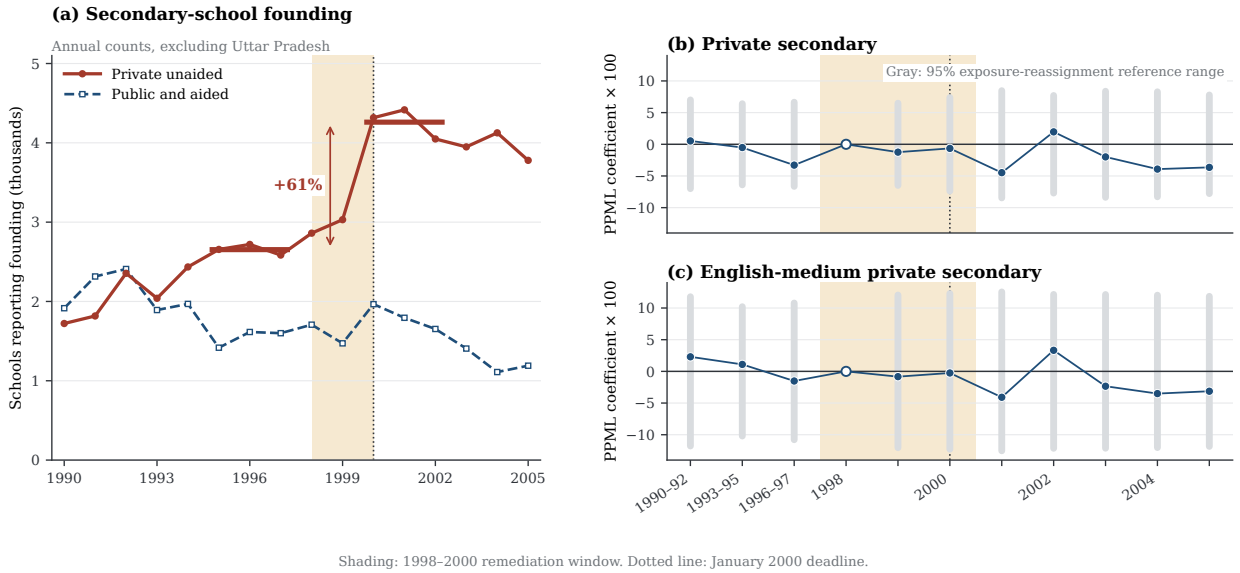


Figure B.7: Secondary-school founding around the Y2K window.

Notes: Panel (a) reports annual counts of schools recording each founding year outside Uttar Pradesh; 2000 is also a common heap in retrospectively reported founding years. Panels (b) and (c) report PPML event-study coefficients on interactions between period and the district's standardized IT-employment share in the 1998 Economic Census, multiplied by 100, with district and state-by-year fixed effects; 1998 is omitted and the pre-period years are binned. Gray vertical segments are symmetric exposure-reassignment reference ranges centered at zero, with half-width the 95th percentile of the absolute coefficient over 999 global reassignments of exposure across districts. They are exchangeability and leverage benchmarks, not confidence intervals, minimum detectable effects, or exact randomization inference. Founding dates are retrospective and school management and medium are observed in a later census, so the exercise conditions on survival and later classification. Shading marks the 1998–2000 remediation window and the dotted line the January 2000 deadline. Online Appendix Section I reports the underlying coefficients.

paths should not be interpreted as precise causal zeros. In the best-supported post-period years, the exposure-reassignment reference half-widths are approximately 6–8 log points per standard deviation for private entry and 12–13 log points for English-medium entry. The evidence therefore weighs against a large and geographically concentrated response on these measured margins while leaving smaller responses unresolved.

Interpretation and limits. Y2K may have increased the expected return to technical education, but the short deadline favored workers who already possessed the relevant skills. Students who had not acquired those skills had little time to complete formal training before the opportunity expired. Any formal-education response therefore had to operate on a longer horizon. Evidence that India's broader IT boom subsequently increased schooling, engineering enrollment, and mobility is consistent with such delayed adjustment (Shastry 2012; Ghose 2026).

This timing difference raises a distributional possibility rather than establishing a distributional result. The first-round gains from a short-lived opportunity with high skill requirements may accrue disproportionately to firms, workers, and places already above the relevant threshold instead of diffusing immediately through local human-capital formation. Skilled workers may also move toward the opportunity, but the evidence here does not identify a Y2K-induced migration response.

The evidence supports a narrower conclusion. Newly founded recognized schools are unlikely to have supplied a quantitatively important share of the immediate response, and their entry did not show a robust, deadline-timed tilt toward exposed districts. The data do not observe new seats or quality within existing institutions, short-cycle private IT training, worker origins, wages, consumption, rents, or remittances. They therefore neither estimate how the gains were distributed nor show that migration widened inequality. The education evidence motivates a hypothesis about the local reach of capability-gated booms; testing its distributional and structural-transformation consequences requires evidence beyond this paper.

C Model Derivations and Extensions

This appendix supports Section 4. Section C.1 derives the organizational-capacity microfoundation behind the quadratic normal form, and Section C.2 states the corresponding conditions for general convex operating costs. Section C.3 proves the phase-level adoption cutoff and its complementarity thresholds, and treats unrestricted upgrading dates, implementation lags, and observed certification. The remaining sections develop four extensions: multi-tier capability, prior access and deadline search, a persistent participation state, and a buyer-assurance procurement gate. None of them changes the paper's baseline exposure primitive, the predetermined technical match L_i.

C.1 Organizational capacity and scale–capability complementarity

This subsection derives the complementarity between operating scale and organizational capability. The allocation problem is a reduced-form model motivated by work on limited managerial attention (Geanakoplos and Milgrom 1991) and knowledge-based hierarchies (Garicano 2000); it abstracts from the explicit hierarchical information and problem-solving structures in those papers. Set $C_{iq}(0) = 0$. For $x > 0$, suppose an operation of scale x

consists of a continuum of project or delivery units $j \in [0, x]$. The firm assigns effective organizational attention $z_j > 0$ across these units, subject to total coordination capacity A_{iq}. The loss at a unit receiving attention z is $\ell_i(z)$, where $\ell_i'(z) < 0$ and $\ell_i''(z) > 0$. The loss function is the same in both organizational states; upgrading increases effective capacity from A_{i0} to A_{i1}, so $A_{i1} > A_{i0} > 0$. Organizational cost is

$$C_{iq}(x) = \min_{\{z_j > 0\}_{j \in [0,x]}} \int_0^x \ell_i(z_j)\, dj \qquad \text{subject to} \qquad \int_0^x z_j\, dj \leq A_{iq}. \tag{22}$$

Because ℓ_i is decreasing, the capacity constraint binds. Convexity and Jensen's inequality then imply equal allocation. Since average attention is A_{iq}/x, the optimal allocation is $z_j^* = A_{iq}/x$ for every unit, and hence

$$C_{iq}(x) = x\ell_i\left(\frac{A_{iq}}{x}\right). \tag{23}$$

Let $h_i(z) = \ell_i(z) - z\ell_i'(z)$. Differentiating equation (23) gives

$$C_{iq}'(x) = h_i\left(\frac{A_{iq}}{x}\right), \qquad h_i'(z) = -z\ell_i''(z) < 0, \qquad C_{iq}''(x) = \frac{A_{iq}^2}{x^3}\ell_i''\left(\frac{A_{iq}}{x}\right) > 0. \tag{24}$$

Since $A_{i1} > A_{i0}$ and h_i is decreasing, the derivative of the fixed-scale operating-cost saving is

$$\frac{d}{dx}\left[C_{i0}(x) - C_{i1}(x)\right] = h_i\left(\frac{A_{i0}}{x}\right) - h_i\left(\frac{A_{i1}}{x}\right) > 0. \tag{25}$$

Thus the return to organizational capability rises with scale. Expansion spreads a fixed organizational capacity across more operating units, while the upgraded system increases effective capacity. The same calculation gives $C_{i1}'(x) < C_{i0}'(x)$, so upgrading lowers marginal operating cost and supports greater optimal scale.

Optimized return under general attention losses. The attention structure also disciplines the return after scale is reoptimized. Substituting $z = A_{iq}/x$ for $x > 0$ gives

$$v_{iq}(m) \equiv \sup_{x \geq 0}\{mx - C_{iq}(x)\} = A_{iq}G_i(m), \qquad G_i(m) \equiv \max\left\{0, \sup_{z>0}\frac{m - \ell_i(z)}{z}\right\}. \tag{26}$$

Whenever the value is finite, $G_i(m)$ is weakly increasing and convex because it is the supremum of affine functions of m with nonnegative slopes. The optimized return to capability is therefore

$$\Delta_i(m) = v_{i1}(m) - v_{i0}(m) = (A_{i1} - A_{i0})G_i(m), \tag{27}$$

which is also weakly increasing and convex.

Strict results require an interior optimum. Let $h_i(z) = \ell_i(z) - z\ell_i'(z)$ as above, and suppose each relevant demand index lies in the interior range

$$m \in \left(\lim_{z \to \infty} h_i(z), \lim_{z \downarrow 0} h_i(z) \right). \tag{28}$$

Because $h_i'(z) < 0$, there is a unique attention level $z_i^*(m)$ satisfying $h_i(z_i^*(m)) = m$. It is common across organizational states, while optimal scale is proportional to capacity:

$$x_{iq}^*(m) = \frac{A_{iq}}{z_i^*(m)}, \qquad G_i(m) = \frac{m - \ell_i(z_i^*(m))}{z_i^*(m)} = -\ell_i'(z_i^*(m)). \tag{29}$$

The envelope and inverse-function theorems then give

$$G_i'(m) = \frac{1}{z_i^*(m)} > 0, \qquad G_i''(m) = \frac{1}{[z_i^*(m)]^3 \ell_i''(z_i^*(m))} > 0, \tag{30}$$

and hence

$$\Delta_i'(m) = \frac{A_{i1} - A_{i0}}{z_i^*(m)} > 0, \qquad \Delta_i''(m) = \frac{A_{i1} - A_{i0}}{[z_i^*(m)]^3 \ell_i''(z_i^*(m))} > 0. \tag{31}$$

Thus increasing differences in organizational capability and demand follow from the general attention structure on the interior domain. This conclusion uses the same loss function ℓ_i in both organizational states. If upgrading also changes its shape, rather than only effective capacity, the common attention level and the convexity result need not survive.

The body specializes the loss function to

$$\ell_i(z) = \frac{\theta_i}{2z}, \qquad \theta_i > 0. \tag{32}$$

Here $h_i(z) = \theta_i / z$, so the interiority condition holds for every $m > 0$ and $z_i^*(m) = \theta_i / m$. This specialization yields closed forms rather than the qualitative increasing-differences result itself. Equations (23) and (32) give

$$C_{iq}(x) = \frac{\theta_i}{2A_{iq}} x^2, \qquad c_{iq} = \frac{\theta_i}{A_{iq}}, \qquad \psi_i = \frac{A_{i1} - A_{i0}}{2\theta_i}. \tag{33}$$

At any common ordinary-demand index $\mu > 0$, optimal attention per operating unit is $z_i^*(\mu) = \theta_i / \mu$ in both organizational states, and $x_{i1}^*(\mu) / x_{i0}^*(\mu) = A_{i1} / A_{i0}$. The upgraded firm therefore uses its additional capacity to operate a proportionally larger ordinary book rather than holding idle attention. At a fixed scale, the cost saving is $\theta_i x^2 (1/A_{i0} - 1/A_{i1})/2$ and

therefore increases with θ_i. Once the firm reoptimizes scale, however,

$$\frac{\partial \psi_i}{\partial A_{i1}} = \frac{1}{2\theta_i} > 0, \qquad \frac{\partial \psi_i}{\partial A_{i0}} = -\frac{1}{2\theta_i} < 0, \qquad \frac{\partial \psi_i}{\partial \theta_i} = -\frac{A_{i1} - A_{i0}}{2\theta_i^2} < 0. \tag{34}$$

The last comparison differs from the fixed-scale comparison because a larger loss parameter contracts optimal scale in both organizational states. These parameters are latent and are not estimated separately. The capacity interpretation also requires ordinary and remediation work to draw on the same delivery organization. If Y2K changes A_{iq}, θ_i, or the investment cost through learning, buyer standards, or financing, that is an additional channel rather than the baseline scale effect.

The next subsection states the corresponding conditions directly for operating costs not generated by the attention problem.

C.2 General convex costs and increasing differences

Let operating profit at organizational state $q \in \{0, 1\}$ be

$$v_{iq}(m) = \max_{x \geq 0} \{mx - C_{iq}(x)\}, \tag{35}$$

where C_{iq} is twice continuously differentiable. For every relevant m, assume $C'_{iq}(0) < m$, $C'_{iq}(x) \to \infty$ as $x \to \infty$, and $C''_{iq}(x) > 0$ on the relevant domain. The unique interior optimizer $x_{iq}(m)$ is then defined by $C'_{iq}(x_{iq}(m)) = m$. The optimized return to organizational capability is

$$\Delta_i(m) \equiv v_{i1}(m) - v_{i0}(m). \tag{36}$$

The envelope theorem and implicit-function theorem give

$$\Delta'_i(m) = x_{i1}(m) - x_{i0}(m) \tag{37}$$

and

$$\Delta''_i(m) = \frac{1}{C''_{i1}(x_{i1}(m))} - \frac{1}{C''_{i0}(x_{i0}(m))}. \tag{38}$$

Two restrictions are distinct for arbitrary operating costs. If organizational upgrading weakly lowers marginal operating cost,

$$C'_{i1}(x) \leq C'_{i0}(x) \quad \text{for all relevant } x, \tag{39}$$

then $x_{i1}(m) \geq x_{i0}(m)$ and equation (37) implies that the return to capability is nondecreasing in project demand. Convexity of that return requires the additional response condition

$$C''_{i1}(x_{i1}(m)) \leq C''_{i0}(x_{i0}(m)), \tag{40}$$

which makes $\Delta''_i(m) \geq 0$. Lower marginal cost alone does not imply this second condition.

The attention family in the preceding subsection satisfies the response condition automatically. At the common optimal attention level $z^*_i(m)$,

$$C''_{iq}(x^*_{iq}(m)) = \frac{[z^*_i(m)]^3 \ell''_i(z^*_i(m))}{A_{iq}}. \tag{41}$$

Because $A_{i1} > A_{i0}$, its upgraded-state curvature is strictly lower. The extra response condition is needed for an arbitrary pair of convex cost functions, but is implied by the organizational-capacity structure.

For ordinary-demand index μ and temporary Y2K-demand index $y \geq 0$, define

$$\phi_i(\mu, y) = \Delta_i(\mu + y) - \Delta_i(\mu). \tag{42}$$

Equations (37)–(40) imply

$$\frac{\partial \phi_i}{\partial y} = \Delta'_i(\mu + y) \geq 0, \qquad \frac{\partial \phi_i}{\partial \mu} = \Delta'_i(\mu + y) - \Delta'_i(\mu) \geq 0. \tag{43}$$

Thus capability supports greater efficient scale, and the return created by a given temporary demand increment is larger at greater ordinary demand when Δ_i is convex. The quadratic specialization in Section 4, $C_{iq}(x) = c_{iq}x^2/2$, satisfies these restrictions and gives $\Delta_i(m) = \psi_i m^2$.

C.3 The phase-level upgrading cutoff

The option to postpone a sunk choice follows the export-participation tradition (Baldwin and Krugman 1989; Roberts and Tybout 1997; Das et al. 2007); the finite-window normalization below is specific to this application.

Set the window opening to $t = 0$. Let N denote the no-Y2K path and Y the Y2K path. The optimized operating profits are

$$\Pi^N_{it}(q) = \frac{\mu^2_{it}}{2c_{iq}}, \qquad \Pi^Y_{it}(q) = \frac{(\mu_{it} + y_{it})^2}{2c_{iq}}, \tag{44}$$

with $\Pi_{it}^{Y}(q) = \Pi_{it}^{N}(q)$ after T. Upgrading values are measured relative to remaining at $q = 0$ under the same demand path. This normalization removes operating profits that do not depend on upgrade timing.

Assume discounted normal capability returns are summable:

$$\sum_{t=0}^{\infty} \lambda^t \Delta_i(\mu_{it}) < \infty. \tag{45}$$

Under the quadratic specialization, $\sum_t \lambda^t \mu_{it}^2 < \infty$ is sufficient. This condition also makes the unrestricted values finite and, with never upgrading assigned value zero, ensures that the maxima below are attained: postponed-date values converge to zero as the upgrading date diverges, while the in-window date set is finite.

Policy W pays K_i at $t = 0$ and holds $q = 1$ thereafter. Policy P remains at $q = 0$ through the window and chooses its best later upgrading date $\tau > T$, including never. Their values when Y2K demand is absent are

$$V_i^{W,N} = -K_i + \sum_{t=0}^{\infty} \lambda^t \left[\Pi_{it}^{N}(1) - \Pi_{it}^{N}(0)\right] \tag{46}$$

and

$$V_i^{P,N} = \max_{\tau \in \{T+1, T+2, \ldots, \infty\}} \left\{ -\lambda^\tau K_i + \sum_{t=\tau}^{\infty} \lambda^t \left[\Pi_{it}^{N}(1) - \Pi_{it}^{N}(0)\right] \right\}, \tag{47}$$

where the value at $\tau = \infty$ is zero. All flows and any later investment cost are therefore discounted to the window opening.

Because policy P remains unupgraded throughout the window, its capability value is unchanged by Y2K in this relative normalization: $V_i^{P,Y} = V_i^{P,N}$. Policy W instead receives the discounted Y2K upgrading premium

$$\begin{aligned} \Gamma_i &= \sum_{t=0}^{T} \lambda^t \Big\{ [\Pi_{it}^{Y}(1) - \Pi_{it}^{Y}(0)] - [\Pi_{it}^{N}(1) - \Pi_{it}^{N}(0)] \Big\} \\ &= \sum_{t=0}^{T} \lambda^t \psi_i (2\mu_{it} y_{it} + y_{it}^2) \end{aligned} \tag{48}$$

in addition to $V_i^{W,N}$. Using $\mathcal{R}_i$ from equation (13), the final line is $\Gamma_i = \psi_i \mathcal{R}_i$. In total-profit levels, both policies also earn the common $q = 0$ remediation return

$$\Xi_i^0 = \sum_{t=0}^{T} \lambda^t [\Pi_{it}^{Y}(0) - \Pi_{it}^{N}(0)]. \tag{49}$$

It cancels from their comparison; it is not absent from the postponed policy.

Define the value advantage of postponement when Y2K demand is absent as

$$\Omega_i = V_i^{P,N} - V_i^{W,N}. \tag{50}$$

A firm strictly chooses P without Y2K when $\Omega_i > 0$. The Y2K difference between upgrading at $t = 0$ and postponement is

$$V_i^{W,Y} - V_i^{P,Y} = \Gamma_i - \Omega_i. \tag{51}$$

Resolving indifference toward earlier upgrading, the shock changes the choice from P to W if and only if

$$0 < \Omega_i \leq \Gamma_i, \tag{52}$$

which proves equation (14). Without the early-upgrading tie break, the upper inequality is strict. The proof holds other durable states fixed: Y2K does not alter the post-deadline branch through cash, learning, financing, or participation in the baseline.

Under the organizational-capacity specialization, the same cutoff yields an exact selection result in ψ_i even when policy P includes finite later upgrading. For every $\tau \in \{T+1, T+2, \ldots, \infty\}$, define

$$\mathcal{M}_i(\tau) \equiv \sum_{t=0}^{\tau-1} \lambda^t \mu_{it}^2, \qquad \lambda^\infty = 0. \tag{53}$$

Using equations (46) and (47), the no-Y2K waiting advantage can then be written as

$$\Omega_i(\psi_i) = \max_{\tau > T} \left\{ (1-\lambda^\tau) K_i - \psi_i \mathcal{M}_i(\tau) \right\}. \tag{54}$$

Upgrading at $t = 0$ weakly dominates every postponed date without Y2K when ψ_i exceeds the largest pairwise threshold. With Y2K, upgrading at $t = 0$ receives the additional return $\Gamma_i = \psi_i \mathcal{R}_i$. Define the two thresholds

$$\overline{\psi}_i^N \equiv \max_{\tau > T} \frac{(1-\lambda^\tau) K_i}{\mathcal{M}_i(\tau)}, \qquad \underline{\psi}_i^Y \equiv \max_{\tau > T} \frac{(1-\lambda^\tau) K_i}{\mathcal{M}_i(\tau) + \mathcal{R}_i}. \tag{55}$$

With indifference resolved toward upgrading at $t = 0$, the firm postpones without Y2K exactly when $\psi_i < \overline{\psi}_i^N$ and upgrades at $t = 0$ under Y2K exactly when $\psi_i \geq \underline{\psi}_i^Y$. Y2K

therefore induces early upgrading exactly for

$$\underline{\psi}_i^Y \leq \psi_i < \overline{\psi}_i^N. \tag{56}$$

For $\mathcal{R}_i > 0$, every ratio defining $\underline{\psi}_i^Y$ is smaller than its no-Y2K counterpart, so $\underline{\psi}_i^Y < \overline{\psi}_i^N$. A larger $\mathcal{R}_i$ lowers the Y2K threshold and expands the induced interval toward firms with weaker complementarity, holding K_i and the ordinary-demand path fixed. When $L_i = 0$, $\mathcal{R}_i = 0$, the thresholds coincide, and the interval is empty. Because ψ_i, K_i, and ordinary demand may covary across firms, this conditional result is not an unconditional prediction by observed firm size.

In the now-or-never special case, $V_i^{P,N} = 0$. Let $\mathcal{M}_i \equiv \mathcal{M}_i(\infty)$. Then

$$\Delta V_i^N = \sum_{t=0}^{\infty} \lambda^t [\Pi_{it}^N(1) - \Pi_{it}^N(0)] = \psi_i \mathcal{M}_i, \tag{57}$$

so $\Omega_i = K_i - \psi_i \mathcal{M}_i$, and equation (52) becomes

$$\psi_i \mathcal{M}_i < K_i \leq \psi_i(\mathcal{M}_i + \mathcal{R}_i) \quad \Longleftrightarrow \quad \frac{K_i}{\mathcal{M}_i + \mathcal{R}_i} \leq \psi_i < \frac{K_i}{\mathcal{M}_i}. \tag{58}$$

If $\Delta_i(\mu_{it})$ is constant over time, every strict postponer chooses never to upgrade, so this collapse applies. If that return rises sufficiently over time, finite post-deadline upgrading can instead yield $V_i^{P,N} > 0$, and the general cutoff is required.

Unrestricted upgrading and implementation. Suppose a firm can launch at any date $s \in \mathbb{Z}_+$, pays K_i then, and becomes operational after an integer implementation lag $d_i \in \mathbb{Z}_+$. The branch derivation below takes $d_i \leq T$; if $d_i > T$, no launch can make capability operational inside the window and the baseline mechanism has no in-window upgrading branch. The no-Y2K incremental value is

$$U_i^N(s) = -\lambda^s K_i + \sum_{t=s+d_i}^{\infty} \lambda^t \Delta_i(\mu_{it}), \tag{59}$$

with $U_i^N(\infty) = 0$. The window return captured by that launch date is

$$H_i(s) = \sum_{t=s+d_i}^{T} \lambda^t \phi_{it}, \tag{60}$$

with an empty sum equal to zero. Let $\mathcal{S}_W = \{s : s + d_i \leq T\}$ be launch dates that make capability operational inside the window, and let $\mathcal{S}_P$ contain the remaining dates and never upgrading. The optimized no-Y2K values are

$$\widehat{V}_i^{W,N} = \max_{s \in \mathcal{S}_W} U_i^N(s), \qquad \widehat{V}_i^{P,N} = \max_{s \in \mathcal{S}_P} U_i^N(s). \tag{61}$$

Under Y2K the in-window branch is instead

$$\widehat{V}_i^{W,Y} = \max_{s \in \mathcal{S}_W} \{U_i^N(s) + H_i(s)\}. \tag{62}$$

The correct optimized window increment is therefore

$$\widetilde{\Gamma}_i = \max_{s \in \mathcal{S}_W} \{U_i^N(s) + H_i(s)\} - \max_{s \in \mathcal{S}_W} U_i^N(s). \tag{63}$$

It is generally incorrect to evaluate $H_i(s)$ at an independently optimized no-Y2K launch date and add it to that branch. With $\widetilde{\Omega}_i = \widehat{V}_i^{P,N} - \widehat{V}_i^{W,N} > 0$, unrestricted timing produces the parallel cutoff $0 < \widetilde{\Omega}_i \leq \widetilde{\Gamma}_i$.

For fixed d_i, $H_i(s)$ is weakly decreasing in s. This property and earliest-date tie-breaking imply that temporary Y2K demand cannot delay the globally optimal upgrading date: any later date that was weakly worse without the shock receives a weakly smaller window bonus. If project load changes implementation time or the cost of managerial attention, that monotonicity need not hold. Such implementation congestion is a genuine counterforce omitted from the baseline firm problem.

Operational capability and observed certification. For a realized launch date s_i, let $d_i^{\text{cert}} \geq 0$ be the additional audit and reporting lag. A binary observation rule is

$$\mathsf{Cert}_{it} = \mathbf{1}\{t \geq s_i + d_i + d_i^{\text{cert}}\}. \tag{64}$$

Operational routines can therefore earn a remediation return before the corresponding certificate appears in the directory. Audit delay can reconcile window implementation with post-deadline certification; implementation that occurs only after T cannot generate equation (48). The baseline algebra keeps $q = 1$ after upgrading. Allowing depreciation or abandonment replaces its infinite normal-return stream with an expected-duration stream; interpreting the baseline over the empirical horizon does not require literal permanence.

C.4 Multi-tier organizational capability

Let latent capability take ordered values $q \in \{0, 1, \ldots, Q\}$, with effective organizational capacity $A_i(q+1) > A_i(q)$ and $c_i(q) = \theta_i / A_i(q)$. Conditional on q being the firm's current tier, the local return to the increment $q \to q+1$ is

$$\Delta_i^{q \to q+1}(m) = \psi_i^q m^2, \qquad \psi_i^q = \frac{1}{2}\left[\frac{1}{c_i(q+1)} - \frac{1}{c_i(q)}\right] = \frac{A_i(q+1) - A_i(q)}{2\theta_i} > 0, \tag{65}$$

and its Y2K increment is

$$\phi_{it}^{q \to q+1} = \psi_i^q (2\mu_{it} y_{it} + y_{it}^2). \tag{66}$$

The binary cutoff applies to an isolated increment, or locally when continuation values from higher tiers are held fixed. Sequential and non-skippable investments alone are not sufficient: moving from q to $q+1$ may unlock continuation value from later rungs. If firms can choose a terminal tier directly, they solve one global problem over tiers; increasing differences makes the chosen tier weakly increase with the Y2K demand index under the usual single-crossing condition, but the binary interval need not apply separately to every rung.

A continuous version makes the upgrading-intensity comparison explicit. Suppose $q \in [0, Q]$ is chosen once at the window opening, effective capacity $A_i(q)$ is differentiable and increasing, and the sunk upgrading cost is $K_i(q)$ with $K_i(0) = 0$. In the now-or-never benchmark, after subtracting the payoff at $q = 0$, the firm's objective is

$$V_i(q; \mathcal{R}_i) = \frac{A_i(q) - A_i(0)}{2\theta_i}(\mathcal{M}_i + \mathcal{R}_i) - K_i(q). \tag{67}$$

It has increasing differences in organizational capability and the Y2K demand contribution:

$$\frac{\partial^2 V_i(q; \mathcal{R}_i)}{\partial q \, \partial \mathcal{R}_i} = \frac{A_i'(q)}{2\theta_i} > 0. \tag{68}$$

Under the usual second-order condition, an interior unique choice satisfies

$$\frac{dq_i^*}{d\mathcal{R}_i} = \frac{A_i'(q_i^*)/(2\theta_i)}{K_i''(q_i^*) - A_i''(q_i^*)(\mathcal{M}_i + \mathcal{R}_i)/(2\theta_i)} > 0. \tag{69}$$

Thus a stronger temporary opportunity can raise the chosen degree of organizational upgrading as well as induce a first upgrade. This result conditions on an opening choice. Combining a continuous tier with unrestricted launch dates requires additional conditions on the joint timing and intensity problem.

The empirical certification ladder is an ordinal measurement device for latent process maturity. This extension does not assert that ISO and CMM are one institutional sequence, that firms must obtain ISO before CMM, or that observed categories cannot be skipped. It allows stronger opportunities to produce deeper, graded capability differences that remain observable after the deadline. It does not by itself generate dynamic post-deadline widening, which would require a joint timing–intensity problem or an observation lag.

C.5 Prior access and deadline search

Let r_i record realized pre-shock foreign reach and let υ_t denote deadline urgency. The Y2K demand facing the firm can be written

$$y_{it} = L_i S_t \chi(r_i, \upsilon_t), \qquad \chi(r_i, \upsilon_t) \geq 0. \tag{70}$$

The body normalizes $\chi = 1$. Prior reach remains distinct from technical match: L_i says whether the firm can perform the work, while χ describes how much of the global Y2K demand it can reach.

The cited interventions assign export orders or subsidize fixed market-access expenditures (Atkin et al. 2017a; Ali et al. 2025). By contrast, r_i records realized pre-shock engagement; it is neither an assigned buyer connection nor an assigned reduction in an access cost.

For a matched firm facing positive window demand, $\chi_r > 0$ raises Γ_i holding μ_{it} and Ω_i fixed. That is not a prediction that high- or low-reach firms must have the larger observed upgrading response. Reach may also raise ordinary demand, change upgrading costs, or bring forward the no-Y2K upgrading date. The relevant switching index is $\Gamma_i(r) - \Omega_i(r)$, whose sign need not be monotone in r. Deadline pressure may also alter χ_r by widening buyer search, but neither the model nor the observed reach proxy signs that interaction. The extension therefore accommodates the empirical bounded null without deriving equality or identifying a search mechanism.

C.6 A persistent participation state

Hold organizational capability fixed at $\bar{q}$ and let $e_{it} \in \{0, 1\}$ be a participation asset acquired at sunk cost K_i^e. Write the demand index without the participation asset as $m_{it}^0 = \mu_{it} + y_{it}$.

While retained, the asset raises that index by $\eta_i > 0$:

$$m_{it}(e) = m_{it}^0 + \eta_i e. \tag{71}$$

Conditional on $\bar{q}$, the flow return to participation, evaluated at the demand index m^0 available when $e = 0$, is

$$\Delta_i^e(m^0) = \frac{(m^0 + \eta_i)^2 - (m^0)^2}{2c_{i\bar{q}}}. \tag{72}$$

The shock's incremental return to participation is

$$\begin{aligned} \phi_{it}^e &= \Delta_i^e(\mu_{it} + y_{it}) - \Delta_i^e(\mu_{it}) \\ &= \frac{\eta_i y_{it}}{c_{i\bar{q}}} \geq 0. \end{aligned} \tag{73}$$

For the parallel cutoff, the opening policy acquires $e = 1$ and retains it throughout the window, while the postponed policy remains at $e = 0$ throughout the window. Values are normalized relative to $e = 0$ under the same demand path, so the common $e = 0$ remediation return cancels. Any post-deadline maintenance, exit, or depreciation rule is the same under the no-Y2K and Y2K paths and is already incorporated in the no-Y2K waiting advantage Ω_i^e. Under these restrictions, defining $\Gamma_i^e = \sum_{t \in \mathcal{W}} \lambda^t \phi_{it}^e$ yields the exact parallel switching condition $0 < \Omega_i^e \leq \Gamma_i^e$. If retention is endogenous inside the window, the optimized branch increment must replace this simple sum.

If the participation asset remains valuable after the deadline, it supports post-window scale without any change in q. This establishes the observational nonuniqueness used in the body. The benchmark requires durability only over the relevant horizon, not an absorbing state. Baldwin and Krugman allow maintenance and exit (Baldwin and Krugman 1989); Roberts and Tybout allow re-entry costs to depend on time out and find rapid depreciation (Roberts and Tybout 1997); Das, Roberts, and Tybout impose complete loss after one year out as a tractable abstraction while retaining a per-period fixed cost (Das et al. 2007).

C.7 Buyer assurance and a procurement gate

The baseline assigns the same Y2K demand index y_{it} to both organizational states and lets internal congestion generate increasing differences. Demanding buyers can reinforce that mechanism. Suppose routine remediation available at either state is $a_{it} = L_i S_t h_0$, while capability unlocks an additional assurance-sensitive tranche $b_{it} = L_i S_t h_1$, with $h_0 > 0$ and

$h_1 \geq 0$. The two organizational states then face

$$m_{it}(0) = \mu_{it} + a_{it}, \qquad m_{it}(1) = \mu_{it} + a_{it} + b_{it}. \tag{74}$$

Under the quadratic operating technology, the window increment to the return from capability decomposes as

$$\begin{aligned} \phi_{it}^{\text{gate}} &= \psi_i(2\mu_{it}a_{it} + a_{it}^2) \\ &\quad + \frac{2(\mu_{it} + a_{it})b_{it} + b_{it}^2}{2c_{i1}}. \end{aligned} \tag{75}$$

The first line is the internal scale–capability complementarity created by demand available at either state. The second is the added return from buyer assurance or a procurement threshold. It is nonnegative and is strictly positive when the gated tranche is positive. If internal organizational costs were identical across states, only the second term would remain.

This extension assumes that some buyers value reliable delivery more on large, deadline-sensitive projects or impose a minimum process-maturity threshold. Foreign-market demand for quality motivates this class of alternatives (Verhoogen 2008; Verhoogen 2023). Buyer interaction can also transmit production knowledge (Atkin et al. 2017a), but that mechanism would require an experience-dependent capability state and is not represented by equation (75). Neither the cited quality-demand papers nor this extension establishes that Y2K contracts generally required certification. The paper does not observe buyer-level orders, requirements, feedback, or a direct measure of technical efficiency needed to distinguish buyer assurance from learning. Buyer assurance is therefore a potential alternative microfoundation or additional channel, not a second identified baseline channel.