The Diffusion State: AI Pilot Zones, Smart Factories, and the Hub Architecture of China's Industrial AI Diffusion

Mateo H. Petel*

Draft prepared for the NBER Economics of Artificial Intelligence conference, Economic Perspectives on AI in China, September 23–24, 2026

Abstract

The economics of artificial intelligence has developed increasingly precise measures of frontier capability, including model performance, compute scale, private investment, talent, and laboratory output. This paper studies diffusion capacity as a distinct economic margin, defined as the capacity of a political economy to embed AI into production systems. I introduce the concept of the AI diffusion state to describe the institutional architecture through which national technology priorities become usable, visible, and scalable across local industrial systems. China provides a useful setting because its AI pilot-zone system and smart-factory recognition programs make a policy-legible part of this adoption architecture observable.

The paper constructs a reproducible city-level measurement pipeline linking 17 National New Generation AI Innovation and Development Pilot Zone units to 509 Ministry of Industry and Information Technology excellence-level smart-factory projects listed in 2024 and 2025. Each project is assigned to a city through an evidence-classified procedure that separates official-location exact assignments, rule-based city inferences, and externally verified assignments. Pilot-zone cities account for 192 listed projects across 16 cities, compared with 317 projects across 143 non-pilot cities. Baseline city-year models show a positive pilot-zone association, and hub-exclusion estimates show that the association attenuates when Beijing, Shanghai, Shenzhen, Hangzhou, direct-admin municipalities, or the top smart-factory cities are removed.

*Stanford University. Email: mateopetel@stanford.edu. The author thanks the contributors who supported the reproducible evidence pipeline underlying this paper. No specific grant from public, commercial, or nonprofit funding agencies supported this research. The author reports no competing interests. Replication materials are available in the accompanying submission bundle at the git commit cited in the cover letter.

The evidence supports an architectural interpretation of the policy-legible portion of China's industrial AI diffusion. Pilot-zone designation marks part of the adoption geography, while municipal capacity, administrative status, industrial ecosystems, and sectoral compatibility shape where recognized adoption appears. Sectoral and export-relevance analyses show that listed smart-factory projects appear in manufacturing domains where AI-enabled production systems are technologically plausible and industrially salient. Appendix analyses add partial public city controls and a tiered patent-geography robustness extension under explicit limits on causal interpretation. The paper contributes a framework for measuring AI diffusion capacity and shows why AI advantage depends on institutions that translate capability into production.

Keywords. artificial intelligence, technology diffusion, industrial policy, China, smart factories, innovation zones, general-purpose technology, recognized adoption, manufacturing
JEL codes. O14, O25, O30, O33, O38, R11, L52

1 Introduction

The economics of artificial intelligence needs a measure of diffusion capacity alongside its increasingly sophisticated measures of frontier capability. Current debates over AI competition often focus on who can build the strongest models, assemble the largest compute clusters, attract the deepest talent pools, raise the most capital, or produce the most visible technical breakthroughs (Agrawal et al., 2019; Maslej et al., 2025; Agrawal et al., 2022). These measures capture the production of capability, yet they leave unresolved a second question that is equally central to economic performance. The consequences of AI will depend on whether firms, cities, sectors, and institutions can reorganize production around AI-enabled systems. A general-purpose technology becomes economically consequential through complementary capital, organizational change, sectoral knowledge, and implementation capacity (Bresnahan and Trajtenberg, 1995; Brynjolfsson et al., 2021; Comin and Hobijn, 2010).

This paper studies China's industrial AI strategy through the lens of diffusion capacity. China is a revealing case because its AI strategy combines frontier ambitions with a large institutional apparatus for embedding digital and intelligent technologies into manufacturing (Naughton, 2021; Juhász et al., 2024). National pilot zones, municipal implementation programs, industrial parks, firm upgrading initiatives, smart-factory recognition, and sector-specific policies form a system for moving AI from national priority into production practice. I define this system as the AI diffusion state, meaning the institutional architecture through which a political economy makes general-purpose technologies usable, visible, and scalable across local production systems.

The paper's central claim is that AI competition has a production-side margin that remains undermeasured. Frontier indicators reveal who produces capability. Diffusion indicators reveal who can absorb capability into factories, supply chains, industrial routines, and local implementation systems. This distinction matters because AI becomes economically consequential through adoption environments, complementary assets, and organizational change. The diffusion-state framework therefore treats policy designation, municipal capacity, administrative recognition, sectoral compatibility, and industrial geography as part of the economics of AI adoption.

The empirical challenge is that industrial AI diffusion is difficult to observe directly. Many deployments are embedded inside factories, production lines, enterprise software, robotics systems, inspection routines, scheduling tools, and process-control systems. Firm-level adoption is fragmented, administrative programs differ across jurisdictions, and industrial AI often appears through operational redesign before it appears in standard innovation statistics. This paper uses a measurable administrative trace of diffusion by linking China's National

New Generation AI Innovation and Development Pilot Zones to the Ministry of Industry and Information Technology's excellence-level smart-factory recognition lists for 2024 and 2025 (Center for Security and Emerging Technology, 2020; Xinhua News Agency, 2021; Ministry of Industry and Information Technology of the People's Republic of China, 2024, 2025). The resulting design treats smart-factory recognition as a visible indicator of policy-legible industrial upgrading while preserving the distinction between recognized adoption, complete adoption, and causal evaluation.

The paper's measurement choice defines its empirical scope. Smart-factory recognition measures the portion of industrial upgrading that becomes visible through administrative recognition, policy classification, and public reporting. The empirical object is policy-legible recognized adoption, which is narrower than total adoption and well suited for studying how state institutions make diffusion visible. The paper uses this administrative trace to study the architecture of diffusion, while leaving the full universe of firm-level industrial AI use to future work with different data.

The dataset connects 17 AI pilot-zone units to 509 listed smart-factory projects and resolves every project to a city through an evidence-classified geocoding procedure. The city-resolution pipeline separates official-location exact assignments, rule-based city inferences from project text, and externally verified assignments, then reports a stratified audit of sampled cases. This structure is central to the paper because the object of interest is geographic and institutional. If the diffusion state operates through hubs, municipal capacity, and industrial ecosystems, the empirical credibility of the paper depends on whether the adoption trace can be assigned to cities in a transparent and reproducible way.

The first empirical pattern is a strong concentration of listed smart-factory recognition in pilot-zone cities and high-capacity hubs. Pilot-zone cities account for 192 listed smart-factory projects across 16 cities, while non-pilot cities account for 317 projects across 143 cities. The average pilot-zone city has 12.00 listed projects, compared with 2.22 projects in the average non-pilot city. Baseline city-year models show a positive association between pilot-zone status and listed smart-factory counts. This establishes the descriptive overlap between national AI policy designation and recognized industrial adoption.

The central result is that the pilot-zone association attenuates when the largest hubs and administratively distinctive cities are removed. The estimated association falls when Beijing, Shanghai, Shenzhen, and Hangzhou are excluded, and it also weakens when direct-admin municipalities or the top smart-factory cities are removed. This attenuation pattern is the key diagnostic for a hub-centered interpretation. Pilot-zone designation marks part of the diffusion architecture, while much of the observed adoption geography is shaped by municipal scale, administrative capacity, industrial depth, firm concentration, and sectoral opportunity.

The sectoral distribution of smart-factory recognition further indicates that industrial AI diffusion depends on compatibility between AI capabilities and production systems. AI is especially relevant where manufacturing relies on machine vision, quality inspection, predictive maintenance, industrial robotics, smart logistics, digital twins, process control, semiconductor production, battery manufacturing, automotive systems, machinery, chemicals, and pharmaceuticals. The paper therefore studies ex ante industry AI exposure and export relevance as mechanism and strategic-context evidence. These analyses show that recognized adoption is concentrated in sectors where AI-enabled production systems are technologically plausible and industrially salient, while leaving productivity outcomes and changes in export outcomes outside the completed evidence base.

The paper contributes to the economics of AI by shifting attention from capability production to capability absorption. A country may produce capable models and still diffuse them slowly across firms, sectors, and production systems. Another political economy may gain advantage through institutions that coordinate adoption, demonstration, standard setting, municipal implementation, industrial finance, and administrative recognition. The diffusion-state framework makes this second channel observable and creates a way to compare AI strategies through adoption architecture alongside frontier capability.

The paper also contributes to the study of industrial policy by treating policy designations as institutional markers embedded in preexisting capacity. Pilot zones are selected through processes shaped by prior municipal capacity, industrial depth, administrative status, and strategic policy priorities. The empirical goal is to measure whether designated cities sit inside a broader geography of recognized adoption and to characterize the architecture of that geography. This distinction is essential because the strongest pattern in the data is spatial concentration, and spatial concentration is exactly what a diffusion-capacity perspective would expect when complementary assets are unevenly distributed.

The paper's evidence hierarchy preserves the central distinction between measured architecture and causal evaluation. The main text relies on smart-factory recognition, city-resolution evidence, hub-exclusion results, city typologies, sectoral exposure, and export relevance. Appendix analyses report partial public city controls from China City Statistical Yearbook data and a tiered patent-geography robustness extension based on IIDS patent records. These appendix results strengthen the measurement exercise while preserving the paper's claim discipline. Strict EPS and NBS controlled models remain unavailable in the open pipeline, and exact publication-number patent geography remains a separate data requirement.

The argument proceeds through the conceptual framework, the Chinese institutional setting, the measurement pipeline, and the empirical architecture of recognized adoption. Section 2 positions the paper in the literature on AI economics, general-purpose technology

diffusion, industrial policy, state capacity, and economic geography. Section 3 develops the diffusion-state framework. Sections 4 and 5 describe the pilot-zone and smart-factory setting and construct the city-level measurement pipeline. Sections 6 through 10 present the empirical strategy, the descriptive overlap, the hub architecture, sectoral selectivity, and export relevance. Sections 11 through 14 interpret the results, report robustness checks, clarify the scope of inference, and draw implications for measuring AI competition through diffusion capacity.

2 Related literature

This paper contributes to the economics of artificial intelligence by treating diffusion capacity as a distinct object of measurement. A large and growing literature studies AI as a general-purpose technology with implications for prediction, innovation, productivity, labor demand, firm organization, and growth (Agrawal et al., 2019, 2022; Aghion et al., 2019; Cockburn et al., 2019). Much of this work measures either the production of AI capability or the exposure of tasks, firms, and occupations to AI-enabled automation and augmentation (Eloundou et al., 2023; Brynjolfsson et al., 2023; Babina et al., 2024). This paper studies a complementary margin. It asks how a political economy converts AI capability into recognized industrial adoption through local institutions, production systems, administrative categories, and sectoral complements.

The paper builds on the literature on general-purpose technologies, adoption lags, and complementary investment. Classic work on technological diffusion shows that new technologies generate economic effects through slow, uneven, and institutionally mediated adoption (Bresnahan and Trajtenberg, 1995; Comin and Hobijn, 2010; Rogers, 2003). Recent work on the productivity J-curve emphasizes intangible investment, organizational redesign, and delayed measurement of productivity gains (Brynjolfsson et al., 2021). The diffusion-state framework extends this perspective to AI by making municipal implementation capacity, industrial ecosystems, administrative recognition, and sectoral compatibility part of the empirical object. The relevant question is therefore whether the surrounding economy can reorganize around AI once the technical frontier advances.

The paper connects to research on automation, robotics, and production tasks by focusing on the industrial settings where AI becomes embedded in equipment, routines, and process control. AI-enabled manufacturing systems often operate through inspection, prediction, scheduling, robotics, predictive maintenance, logistics, digital twins, and quality management (International Federation of Robotics, 2023). These applications differ from consumer software and frontier-model benchmarks because they depend on sensors, machine interfaces, industrial

data, production expertise, and managerial change. The empirical design therefore measures recognized adoption through smart-factory projects and studies the city and sector conditions under which AI-enabled production systems become visible.

The paper contributes to the literature on industrial policy by treating policy designation as an institutional marker embedded in prior capacity. Recent work on industrial policy emphasizes coordination, learning, firm selection, state capacity, and the difficulty of evaluating policies targeted toward already-capable sectors and locations (Juhász et al., 2024; Rodrik, 2004; Aghion et al., 2015; Brandt et al., 2023; Criscuolo et al., 2019). China's AI pilot zones provide a setting where national designation, municipal implementation, and industrial ecosystems interact. The paper therefore interprets pilot-zone status as part of an adoption architecture, while leaving causal evaluation of designation to designs with stronger counterfactual variation.

The paper also relates to work on procurement, public demand, and state capacity as channels through which governments shape technology adoption (Decarolis et al., 2020; Besley and Persson, 2010). In this setting, the relevant state function extends beyond fiscal support. Pilot-zone designation and smart-factory recognition can coordinate attention, certify exemplary projects, create demonstration channels, make industrial upgrading legible to ministries, and signal priority areas to firms and local governments. The diffusion-state framework treats the state as an organizer of adoption environments, a coordinator of implementation channels, and a producer of policy-legible industrial categories.

The paper speaks to economic geography by showing that industrial AI adoption is spatially concentrated. Innovation and technology adoption are rarely distributed evenly across space because labor markets, supplier networks, universities, firm density, infrastructure, logistics, and administrative capacity differ across cities (Syverson, 2011; Duranton and Puga, 2014; Glaeser and Resseger, 2010). The hub-attenuation results show that China's visible industrial AI adoption is concentrated in cities with exceptional implementation capacity. This spatial concentration is central to the paper's interpretation because it reveals how national AI strategy becomes legible through local industrial systems.

The paper finally contributes to empirical work on China's technology policy and manufacturing upgrading (Naughton, 2021; Brandt et al., 2023; Li, 2024). Much of the debate on China and AI focuses on frontier models, chips, firms, data, research output, and geopolitical competition. This paper studies a different layer of Chinese AI strategy by linking AI pilot-zone designation to MIIT excellence-level smart-factory recognition. The resulting evidence shows how national AI priorities become connected to recognized manufacturing projects, sectoral production complementarities, and export-relevant industrial domains. The contribution is a measurement architecture for China's AI diffusion state, with a conservative

claim hierarchy that separates measured adoption traces, descriptive associations, mechanism evidence, robustness checks, and blocked causal claims.

3 Conceptual framework

The economic importance of artificial intelligence depends on the institutions that convert technical capability into deployed production capacity. Frontier capability describes the ability to build models, assemble compute, produce algorithms, and improve technical performance. Diffusion capacity describes the ability to reorganize firms, sectors, cities, and production systems around those capabilities. A general-purpose technology becomes economically consequential when complementary capital, organizational routines, workforce skills, standards, equipment interfaces, software integration, and sectoral knowledge form around it. The central object of this paper is therefore diffusion capacity, understood as the ability of a political economy to embed AI into production at scale.

The AI diffusion state is the institutional architecture through which national AI priorities become usable, visible, and scalable inside industrial systems. This architecture links policy designation, local government coordination, industrial parks, firm upgrading programs, smart-factory recognition, technical standards, procurement channels, demonstration projects, and sector-specific implementation capacity. The relevant state function extends beyond subsidy allocation. It includes the construction of adoption environments in which firms receive signals, legitimacy, technical support, and coordination mechanisms that lower the cost of reorganizing production around AI-enabled systems.

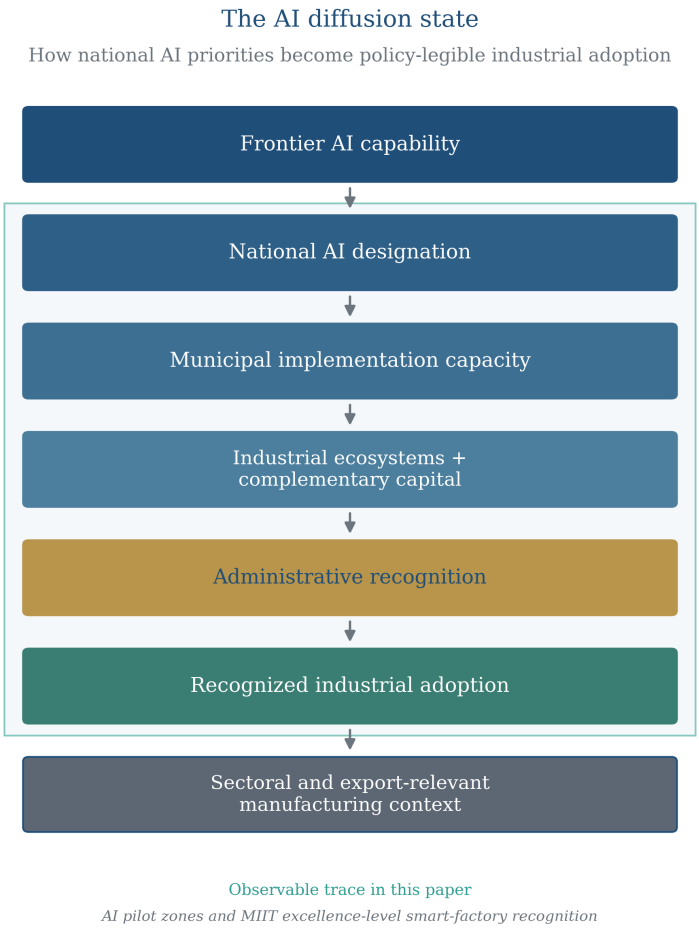


Figure 1: Conceptual architecture of the AI diffusion state. The framework connects frontier AI capability to national designation, municipal implementation capacity, industrial ecosystems, administrative recognition, and sectoral and export-relevant industrial adoption.

The distinction between capability generation and capability absorption is especially important for artificial intelligence because AI often enters the economy through tasks and workflows inside existing systems. In manufacturing, AI may appear as machine vision, predictive maintenance, scheduling optimization, robotics coordination, automated inspection, process control, warehouse management, or digital-twin simulation. These applications are embedded inside production systems, which makes them harder to observe than model releases, chip purchases, research publications, or venture investment. A country can look strong on frontier indicators while remaining weak at industrial absorption, and another country can accumulate advantage through deployment channels that remain largely invisible to frontier rankings.

Diffusion capacity has a spatial structure because the complements required for adoption are unevenly distributed across cities. Large manufacturing hubs often contain skilled engineering labor, supplier networks, universities, technical service firms, industrial parks, logistics systems, leading manufacturers, and local governments experienced in implementing upgrading programs. These assets reduce the cost of adoption and increase the likelihood that a recognized project can be planned, executed, and reported. A hub-centered diffusion architecture should therefore produce concentration in cities that combine formal designation with preexisting industrial and administrative capacity.

Pilot zones matter in this framework as institutional markers inside a broader geography of capacity. Designation can coordinate attention, concentrate resources, generate demonstration effects, and align local governments with national priorities. Designation also tends to select places that already possess the industrial depth and administrative capacity needed for implementation. A positive association between pilot-zone status and smart-factory recognition therefore carries the most economic meaning when interpreted through the geography of complementary capacity. It reveals where national AI policy and local adoption conditions coincide.

Administrative recognition provides a measurable trace of diffusion because many forms of industrial AI adoption are difficult to observe directly. Firms may deploy AI in inspection, maintenance, scheduling, process control, robotics, logistics, and quality management through internal systems that remain outside public records. Smart-factory recognition creates a visible administrative signal that a project has entered a policy-recognized category of digital and intelligent industrial upgrading. This signal captures the portion of adoption that becomes legible to the industrial policy apparatus. The empirical object is therefore recognized adoption, while the broader conceptual object remains diffusion capacity.

Sectoral compatibility is the final layer of the framework. AI diffusion in manufacturing should be stronger where production processes generate data, require inspection, depend on complex scheduling, use expensive equipment, face quality-control constraints, or benefit from robotics and automation. This implies that diffusion capacity should vary across industries as well as across cities. A city with deep activity in electronics, batteries, automotive systems, machinery, chemicals, pharmaceuticals, or semiconductor production may provide a more favorable environment for AI-enabled industrial upgrading than a city whose production base offers fewer AI-complementary processes.

The diffusion-state framework yields three empirical expectations for China's industrial AI adoption architecture. First, listed smart-factory recognition should be concentrated in cities that combine pilot-zone designation with high municipal and industrial capacity. Second, the pilot-zone association should weaken when major hubs and administratively distinctive

cities are excluded, because designation is embedded in a broader capacity hierarchy. Third, recognized adoption should appear more often in sectors where AI has clearer production complementarities and strategic relevance. The empirical sections evaluate these expectations through city-level overlap, hub-exclusion models, city typologies, ex ante industry exposure, and export relevance.

The framework defines the paper's scope of inference by separating the architecture of diffusion from causal evaluation. The evidence can show that pilot-zone designation, smart-factory recognition, hub concentration, sectoral compatibility, and export-relevant manufacturing align in a coherent adoption architecture. Stronger claims about causal effects, productivity gains, complete adoption, changes in export outcomes, and exact patent geography require different evidence. The contribution lies in making an undermeasured margin of AI competition observable through disciplined evidence classes, transparent audit procedures, and claim-specific robustness checks.

4 Institutional background

China's industrial AI strategy makes part of the diffusion state observable because national technology priorities are translated through local implementation systems. The National New Generation AI Innovation and Development Pilot Zones identify cities and city clusters expected to organize AI development, application, experimentation, and demonstration around local industrial strengths. These designations sit inside a wider policy environment in which municipal governments build industrial platforms, coordinate firms and research institutions, attract investment, and align local upgrading programs with national technology priorities. The relevant institutional object is a system of places, programs, firms, and administrative signals through which AI becomes connected to industrial production.

Pilot-zone designation gives local governments a formal role in translating national AI strategy into city-level implementation. A designated city can use the pilot-zone label to coordinate industrial parks, universities, public procurement, financing channels, service providers, and demonstration projects. The designation can also signal that AI-related industrial activity will receive administrative attention. These channels make pilot zones useful for studying diffusion architecture, while designation itself remains embedded in prior municipal capacity and local industrial strength.

Smart-factory recognition provides the complementary adoption trace. The Ministry of Industry and Information Technology recognizes excellence-level smart-factory projects as part of China's broader manufacturing upgrading agenda, and these lists make visible a subset of industrial projects that have entered a formal category of digital and intelligent

transformation. Recognition captures the portion of factory upgrading that becomes legible to the state through project classification, reporting, and administrative validation. It therefore offers a policy-facing trace of industrial AI adoption, while the wider universe of internal firm deployments remains outside the observable list.

The overlap between pilot zones and smart-factory recognition connects two layers of the diffusion state. Pilot zones represent a national-to-local policy architecture for AI development and application. Smart-factory projects represent recognized industrial upgrading activity inside manufacturing systems. When these two layers overlap in cities, the pattern reveals where AI-oriented policy designation and industrial adoption traces coincide. When the overlap concentrates in major hubs, the pattern indicates that implementation capacity and industrial ecosystem depth shape the geography of recognized adoption.

The institutional setting also explains why the paper interprets pilot-zone status as a marker of diffusion architecture. Chinese pilot zones were selected from cities and regions with different industrial bases, administrative status, research systems, firm networks, and implementation capacities. Many designated places already had substantial manufacturing ecosystems, digital infrastructure, universities, leading firms, or local policy experience before designation. The empirical analysis therefore treats pilot-zone status as an institutional signal embedded in a broader geography of capability, using its association with smart-factory recognition to characterize the architecture of diffusion.

The smart-factory lists sit close to the production margin because industrial AI often appears inside operational systems. In manufacturing, these systems may involve inspection, predictive maintenance, robotics coordination, scheduling optimization, digital twins, warehouse automation, process control, and enterprise software integration. These applications rarely generate clean public firm-level AI variables. Recognition lists provide an observable administrative trace of projects that have entered a formal category of intelligent manufacturing, which allows the paper to measure where recognized adoption appears and how that geography relates to AI policy designation.

The institutional background motivates the city-level design. If AI diffusion operates through local implementation capacity and industrial ecosystems, cities become the relevant unit for observing the alignment between designation and recognized adoption. A national policy label matters through the places that implement it, the firms that absorb it, and the industrial systems that make it useful. The analysis therefore resolves each smart-factory project to a city and studies whether recognized adoption is concentrated in pilot-zone cities, major hubs, and sectors with stronger AI-production complementarities.

5 Data and measurement

The empirical design builds a city-level measure of policy-legible industrial AI diffusion by connecting national AI policy designation to recognized smart-factory adoption. The first component is the set of National New Generation AI Innovation and Development Pilot Zones, which identifies cities and city clusters assigned a formal role in China's AI development and application strategy. The second component is the Ministry of Industry and Information Technology's excellence-level smart-factory recognition lists for 2024 and 2025, which identify manufacturing projects that entered an administratively recognized category of digital and intelligent upgrading. The resulting dataset links 17 pilot-zone units to 509 listed smart-factory projects.

The unit of analysis is the city-year because the paper studies diffusion through local implementation capacity and industrial ecosystems. Pilot-zone designation operates through municipal governments, local firms, industrial parks, universities, service providers, and sector-specific production systems. Smart-factory recognition appears at the project level, while the theoretical object requires a city-level geography of recognized adoption. The measurement pipeline therefore resolves each project to a city, assigns project-year and industry information, and constructs city-year smart-factory counts that can be compared across pilot-zone and non-pilot cities.

The smart-factory lists are treated as administrative traces of recognized industrial adoption. Recognition by the Ministry of Industry and Information Technology indicates that a project has entered a formal category of intelligent manufacturing or digital industrial upgrading. This trace captures the visible and policy-legible portion of adoption, especially projects that firms, local officials, and ministries classify as exemplary or advanced. The analysis therefore interprets listed projects as recognized adoption events, while acknowledging that many firm-level uses of AI, industrial software, robotics, inspection systems, scheduling tools, and process-control systems remain outside public recognition lists.

The city-resolution procedure separates location evidence into distinct classes to preserve measurement credibility. Projects with an official location stated in the registry or source material are classified as official-location exact assignments. Projects assigned through structured project text, firm names, industrial park names, or location-bearing descriptions are classified as rule-based text inferences. Projects supported by non-list external sources are classified as external-evidence verified assignments. This classification allows the empirical analysis to use all resolved projects while reporting the evidentiary strength behind the spatial claims.

The final city-resolution register assigns all 509 projects to cities and reports the evidence

Table 1: City-resolution evidence quality

Resolution class	Projects	Share of projects
External evidence verified	50	9.8%
Official location exact	102	20.0%
Rule-based text inference	357	70.1%

Production table source: `table_B_city_resolution_quality.csv`. This table provides evidence-tiered measurement support for the city-resolution procedure. The full evidence-type breakdown appears in Appendix A.

Table 2: Dataset counts and evidence classes

Measure	Count	Role
AI pilot-zone units	17	Policy designation layer
Listed smart-factory projects	509	Recognized adoption trace
Official-location exact assignments	102	Strongest city evidence
Rule-based text inferences	357	Expanded city coverage
External evidence verified assignments	50	Non-list verification layer

Production table source: `table_A_dataset_counts.csv`. This table provides the measured counts used to define the paper-facing evidence base.

tier associated with each assignment. The paper-facing split contains 102 official-location exact projects, 357 rule-based text-inference projects, and 50 externally verified projects. The external verification layer uses non-list sources such as annual reports, company pages, project registries, and industrial-park records to support selected city assignments. This split is central to the paper's credibility because the main claims concern the spatial distribution of recognized adoption.

A stratified audit checks the reliability of the city-resolution procedure. The audit samples official-location exact rows and rule-based inference rows separately because these classes carry different evidentiary risks. The official-location sample confirms 20 of 20 assignments. The rule-based sample confirms 20 assignments, marks 30 as insufficient evidence, and records zero incorrect assignments. The audit gives the strongest support to the official-location layer and records greater uncertainty in the rule-based layer. The results section and appendix preserve this distinction when presenting the geography of recognized adoption.

The paper constructs several city-level outcomes from the resolved project data. The main outcome is the count of listed smart-factory projects in a city-year. Additional measures aggregate projects by pilot-zone status, city type, industry group, and ex ante AI exposure. City typologies distinguish major hubs, pilot-zone cities, non-pilot adopters, and cities with limited or no recognized adoption. These aggregations allow the paper to study whether

recognized adoption appears evenly across the urban system or concentrates in a small set of high-capacity cities.

The industry classification links smart-factory projects to the sectoral layer of the diffusion-state framework. Each project is assigned to an industry category that can be compared with ex ante AI-production compatibility and export relevance. Ex ante exposure is constructed from sectoral characteristics associated with industrial AI use, including machine vision, quality inspection, robotics, predictive maintenance, digital twins, logistics optimization, scheduling, process control, semiconductor production, battery manufacturing, automotive systems, machinery, chemicals, and pharmaceuticals. This exposure measure is defined independently of realized smart-factory counts so that the sectoral analysis reflects technological compatibility instead of outcome-derived labeling.

The export-relevance measures provide strategic context for the industrial distribution of recognized adoption. The paper compares smart-factory sectors with export-relevant manufacturing categories to evaluate whether recognized adoption appears in sectors connected to China's position in global production. These measures are descriptive. They show the industrial importance of the sectors in which smart-factory recognition appears, while leaving claims about changes in export outcomes, productivity gains, and firm-level performance effects outside the completed evidence base.

The empirical tables are organized by claim tier so that each result carries the appropriate interpretation. Dataset counts, evidence-class splits, city-resolution quality, pilot-zone overlap, city typology, sectoral exposure, and export relevance are treated as measured or validated descriptive claims. Baseline associations between pilot-zone status and smart-factory counts are treated as associations. Hub-exclusion results are treated as architecture diagnostics for the hub-centered interpretation. Public city controls and patent-based triangulation are reported as appendix-only robustness exercises because their coverage and measurement constraints place them outside the main identification layer.

The appendix extends the measurement architecture while preserving the main evidence hierarchy. Public city controls from China City Statistical Yearbook tables provide a partial 2024 robustness check using GDP, population, secondary-industry structure, foreign trade, telecom, and industrial-output proxies, while missing FDI and fixed-asset investment prevent an EPS-equivalent controlled design. A tiered IIDS patent-geography extension provides additional triangulation for industrial AI activity, with coverage reported by geography-confidence tier and exact publication-number applicant-address geography treated as unavailable. These appendix layers strengthen the transparency of the project while keeping the main paper centered on smart-factory recognition and audited city-level diffusion evidence.

5.1 Measurement under incomplete observability

The paper's measurement contribution is to make policy-legible industrial AI diffusion observable under incomplete, uneven, and administratively mediated data. Industrial AI adoption is difficult to measure because many deployments occur inside production lines, enterprise software, quality-control systems, robotics cells, warehouse systems, and process-control routines. These deployments may affect production without appearing in standard innovation data, firm surveys, or patent records. The empirical strategy therefore uses administrative recognition as an adoption trace, while explicitly separating what is observed, what is inferred, what is externally verified, and what remains outside the data.

The first measurement problem is the distinction between adoption and recognized adoption. Smart-factory lists identify projects that have entered a formal category of intelligent manufacturing, which makes them valuable for studying policy-legible industrial upgrading. These lists measure a policy-legible trace of visible adoption. This distinction protects the empirical interpretation and turns the smart-factory lists into a disciplined measurement object.

The second measurement problem is spatial assignment. The paper studies diffusion architecture at the city level, so project geography must be resolved with evidence that can be inspected and audited. The city-resolution pipeline assigns every project to a city and classifies the supporting evidence into official-location exact assignments, rule-based text inferences, and externally verified assignments. This tiering matters because different rows carry different degrees of evidentiary strength. The paper reports the evidence classes and uses a stratified audit to make location uncertainty visible.

The third measurement problem is claim discipline. Empirical AI research often combines administrative data, scraped records, inferred geography, patent records, and policy labels, then presents them as if they supported a single level of inference. This paper instead organizes the evidence into claim tiers. Dataset counts and evidence splits are measured facts. Pilot-zone and smart-factory overlap is a descriptive association. Hub attenuation is an architecture diagnostic. Sectoral exposure and export relevance are mechanism and context evidence. Public controls and tiered patent geography are appendix-only robustness exercises. Causal treatment effects, productivity outcomes, exact patent geocoding, and strict production-control estimates remain outside the completed evidence base.

The fourth measurement problem is reproducibility. The paper is accompanied by a pipeline that regenerates the main submission package, validates the evidence gates, exports paper tables and figures, and blocks unsupported claims. This matters because the empirical object is constructed from multiple administrative and derived sources. Reproducibility is therefore part of the measurement design, since it allows the reader to distinguish stable

paper-facing artifacts from exploratory layers, blocked specifications, and robustness-only extensions.

The resulting approach can be used beyond the Chinese case. Many AI-adoption settings will have incomplete observability because deployment often appears inside workflows and production systems before it appears in public statistics. A general measurement strategy for AI diffusion should therefore combine visible administrative traces, evidence-tiered geography, audit procedures, claim-specific robustness, and transparent limits on causal interpretation. The paper contributes such a strategy by showing how an adoption architecture can be measured while preserving the distinction between the observable trace and the full adoption universe.

6 Empirical strategy

The empirical strategy characterizes the architecture of recognized industrial AI adoption. The main outcome is the number of listed excellence-level smart-factory projects in city c and year t, denoted Y_{ct}. The main explanatory variable is an indicator $Pilot_c$ equal to one when a city belongs to a National New Generation AI Innovation and Development Pilot Zone. The baseline specification estimates the association between pilot-zone status and listed smart-factory recognition in the observed public-list window,

$$Y_{ct} = \alpha + \beta Pilot_c + \lambda_t + \varepsilon_{ct}, \tag{1}$$

where λ_t denotes year effects for the 2024 and 2025 smart-factory lists. The coefficient β measures the average difference in listed smart-factory counts between pilot-zone and non-pilot cities within the recognition window. Standard errors are reported in the generated paper tables, and all estimates are interpreted as descriptive associations.

The log-count specification evaluates whether the pilot-zone association is robust to the skewed distribution of recognized projects across cities. The estimating equation is

$$\log(1 + Y_{ct}) = \alpha + \beta Pilot_c + \lambda_t + \varepsilon_{ct}. \tag{2}$$

The transformation preserves zero-count cities while reducing the influence of the largest adoption hubs. The coefficient β in this specification is interpreted as an association between pilot-zone status and the transformed intensity of recognized adoption.

The hub-exclusion design tests whether the aggregate pilot-zone association is carried by cities with exceptional implementation capacity. Let S_k denote a prespecified exclusion set, such as Beijing, Shanghai, Shenzhen, Hangzhou, direct-admin municipalities, or the top

smart-factory cities. For each exclusion rule, the paper estimates

$$Y_{ct} = \alpha_k + \beta_k Pilot_c + \lambda_t + \varepsilon_{ct}, \qquad c \notin S_k. \tag{3}$$

The sequence of coefficients $\{\beta_k\}$ is the paper's central empirical diagnostic. A stable coefficient across exclusions would indicate a broad pilot-zone association distributed across designated cities. An attenuating coefficient indicates that the full-sample association depends on the largest and most capable implementation hubs.

The hub-exclusion coefficients are interpreted as architecture diagnostics. Hub status, administrative rank, industrial structure, and pilot-zone designation are jointly shaped by prior economic capacity, so the exclusions do not identify formal mediation effects. Their value is descriptive and diagnostic. They show whether policy-adoption alignment is broadly distributed across designated cities or concentrated in places with unusually dense complements for AI-enabled industrial upgrading.

The city-typology analysis provides a second diagnostic for spatial concentration. Cities are classified ex ante by pilot-zone status, administrative status, and major-hub indicators before using listed smart-factory outcomes. The paper then compares recognized adoption across these categories. This design separates ex ante city categories from realized smart-factory outcomes and allows the analysis to distinguish pilot-zone hubs, direct-admin municipalities, non-pilot industrial hubs, and lower-intensity adoption cities.

The sectoral analysis links the city-level architecture to the production economics of AI adoption. Let $Exposure_i$ denote an ex ante industry-level measure of AI-production compatibility for industry i. The paper aggregates smart-factory projects by city, year, and industry exposure category, then examines whether recognized adoption is more common in sectors where AI-enabled systems have plausible production complementarities. The sectoral analysis is mechanism evidence. It evaluates whether adoption appears in industries where machine vision, predictive maintenance, robotics, logistics, digital twins, process control, and automated inspection are technologically meaningful.

The export-relevance analysis adds strategic context while preserving the paper's claim hierarchy. Smart-factory sectors are matched to export-relevant manufacturing categories using public trade data and industry mappings. The resulting tables ask whether recognized adoption appears in sectors connected to China's manufacturing position in global markets. Export performance depends on demand, prices, exchange rates, firm capabilities, supply-chain organization, and preexisting industrial specialization, so the tables provide industrial context for the diffusion architecture.

The appendix reports two robustness extensions that remain outside the main identification

Table 3: Pilot-zone and non-pilot smart-factory overlap

City group	Projects	Cities	Mean projects per city
Pilot-zone cities	192	16	12.00
Non-pilot cities	317	143	2.22
Difference			9.78

Production table source: `table_C_pilot_overlap.csv`. This table provides descriptive overlap evidence outside the causal-identification layer.

layer. The public-control exercise uses a restricted 2024 cross-section with available city covariates from the ChinaUTC public pathway. This design includes GDP, population, secondary-industry structure, and limited proxies for foreign trade, telecom, and industrial output, while the strict EPS/NBS specification remains blocked because FDI and fixed-asset investment are unavailable in the public bundle. The tiered patent Atlas uses IIDS industrial-AI patent records with a coverage-qualified applicant-geography layer. It is reported by geography-confidence tier and excluded from publication-ready patent F1 estimates.

The paper requires listed smart-factory recognition, once resolved to cities and reported with evidence tiers, to provide an informative administrative trace of recognized industrial AI adoption. The empirical results are interpreted as evidence about the geography and institutional architecture of diffusion. Causal evaluation of pilot-zone designation is left to future work with richer adoption histories, firm-level outcomes, stronger controls, and credible counterfactual variation.

7 Descriptive overlap between pilot zones and smart-factory recognition

The first empirical pattern is a strong city-level overlap between AI pilot-zone designation and listed smart-factory recognition. Across the 2024 and 2025 excellence-level smart-factory lists, the dataset contains 509 projects resolved to cities. Pilot-zone cities account for 192 of these projects across 16 cities, while non-pilot cities account for 317 projects across 143 cities. The average pilot-zone city has 12.00 listed smart-factory projects, compared with 2.22 projects in the average non-pilot city. This difference establishes the descriptive fact that policy-legible industrial upgrading is substantially more concentrated in cities that also carry national AI pilot-zone status.

The overlap is economically meaningful because the two institutional layers come from different parts of the diffusion architecture. Pilot-zone designation identifies places selected

for AI development, application, experimentation, and demonstration. Smart-factory recognition identifies industrial projects that entered an administratively recognized category of digital and intelligent manufacturing. Their city-level alignment indicates that national AI policy geography and recognized manufacturing-upgrading geography coincide in observable administrative data.

The descriptive contrast also shows why a simple pilot-zone comparison is insufficient. Pilot-zone cities represent a small subset of the urban system, yet they contain a large share of listed projects. Non-pilot cities contain more listed projects in absolute count because the group includes many more cities. The relevant empirical question is therefore about the structure of the geography that produces this difference. The key issue is whether the observed association is broadly distributed across designated cities or concentrated in a smaller set of high-capacity hubs.

The city distribution points toward concentration in places with unusually deep implementation capacity. A small number of cities account for a large share of recognized projects, and several of these cities combine pilot-zone status with manufacturing ecosystems, administrative capacity, research institutions, leading firms, and digital infrastructure. This pattern is consistent with the diffusion-state framework because industrial AI adoption requires complementary assets that are unevenly distributed across cities. Formal designation matters most where firms, municipal governments, technical service providers, and sectoral production systems can absorb AI-enabled upgrading.

The descriptive overlap serves as the starting point for the paper's main empirical interpretation. It shows that AI policy designation and smart-factory recognition are aligned, while leaving open whether the alignment reflects broad pilot-zone diffusion, capacity selection, administrative status, or hub concentration. The next section examines this structure directly by estimating how the pilot-zone association changes when the largest hubs and administratively distinctive cities are removed.

8 Hub architecture

The central empirical diagnostic is that the pilot-zone association attenuates when the largest hubs and administratively distinctive cities are removed. This pattern is central to the diffusion-state interpretation. A broad pilot-zone architecture would produce a relatively stable association across exclusion rules. A hub-centered diffusion architecture would produce a strong full-sample association that depends heavily on cities with exceptional industrial, administrative, and technological complements. The observed pattern supports the hub-centered interpretation.

The baseline city-year models show that pilot-zone status is associated with higher listed smart-factory counts. This association is consistent with the descriptive overlap documented above, where pilot-zone cities have substantially more recognized projects on average than non-pilot cities. The baseline result establishes that national AI policy geography and recognized industrial upgrading geography are aligned in the observed data. Its interpretation depends on whether the association remains stable after removing the cities most likely to dominate implementation capacity.

Table 4: Hub-exclusion summary

Specification	Coefficient	Relative to full sample
Full sample	4.55	100%
Exclude Beijing, Shanghai, Shenzhen, and Hangzhou	3.67	81%
Exclude direct-admin municipalities	2.90	64%
Exclude top five smart-factory cities	2.95	65%

Production table source: `table_D_hub_exclusion.csv`. This table reports architecture diagnostics based on descriptive association estimates.

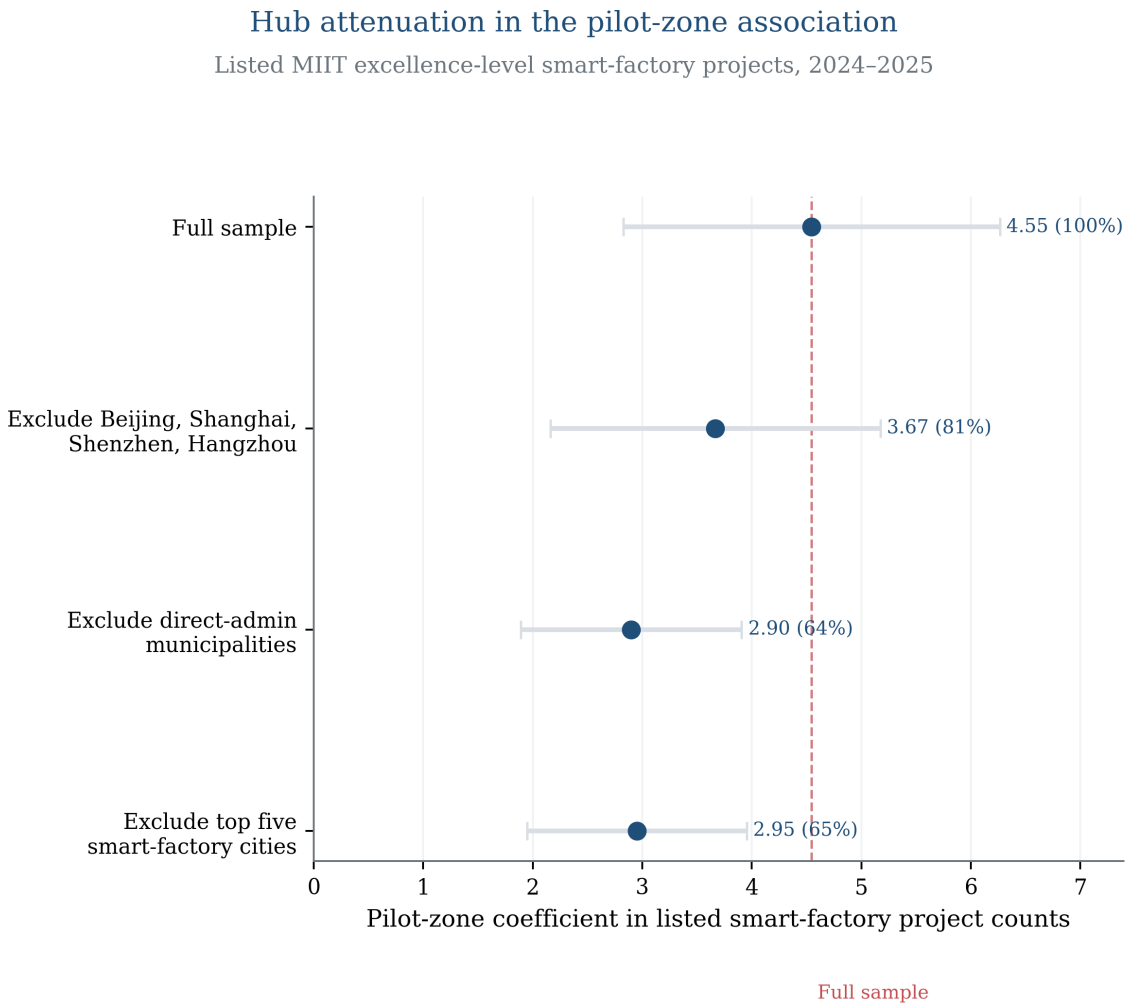


Figure 2: Hub attenuation in the pilot-zone association. The figure reports pilot-zone coefficients for listed smart-factory counts under the full sample and under successive hub-exclusion rules, with 95 percent confidence intervals.

The hub-exclusion estimates show that the full-sample association is partly carried by China's largest AI and manufacturing hubs. When Beijing, Shanghai, Shenzhen, and Hangzhou are excluded, the coefficient falls from about 4.55 to about 3.67. When direct-admin municipalities are removed, the estimate falls to about 2.90. When the top five smart-factory cities are removed, the estimate falls to about 2.95. These changes leave a positive association in place, while showing that a substantial share of the full-sample pattern depends on cities with exceptional administrative status, industrial depth, firm concentration, and digital-economy capacity.

The attenuation pattern is more informative than the baseline association alone. Beijing and Shanghai combine administrative status, research capacity, headquarters density,

and policy visibility. Shenzhen and Hangzhou combine digital firms, manufacturing networks, entrepreneurial ecosystems, and strong municipal implementation capacity. Other high-performing cities reflect similar combinations of industrial specialization, local upgrading programs, and firms capable of implementing intelligent manufacturing systems. The geography of recognized adoption therefore looks like a hierarchy of implementation capacity.

The hub-centered pattern is consistent with how general-purpose technologies diffuse through industrial systems. AI adoption in manufacturing requires complementary investments in equipment, data infrastructure, software integration, process redesign, engineering talent, and managerial routines. These complements are more abundant in cities with dense supplier networks, universities, industrial parks, leading firms, logistics systems, and local governments experienced in manufacturing upgrading. Pilot-zone designation can amplify coordination in such places, while the ability to translate designation into recognized adoption depends on the surrounding ecosystem.

The results also explain why pilot-zone status remains economically informative under a conservative interpretation. A designation can mark the places where national AI strategy and local industrial capacity converge. It can identify cities that national authorities expect to function as demonstration sites, experimentation platforms, or diffusion nodes. The observed association therefore reveals the geography through which the state's AI ambitions become connected to recognized industrial projects, while keeping causal evaluation of designation outside the current design.

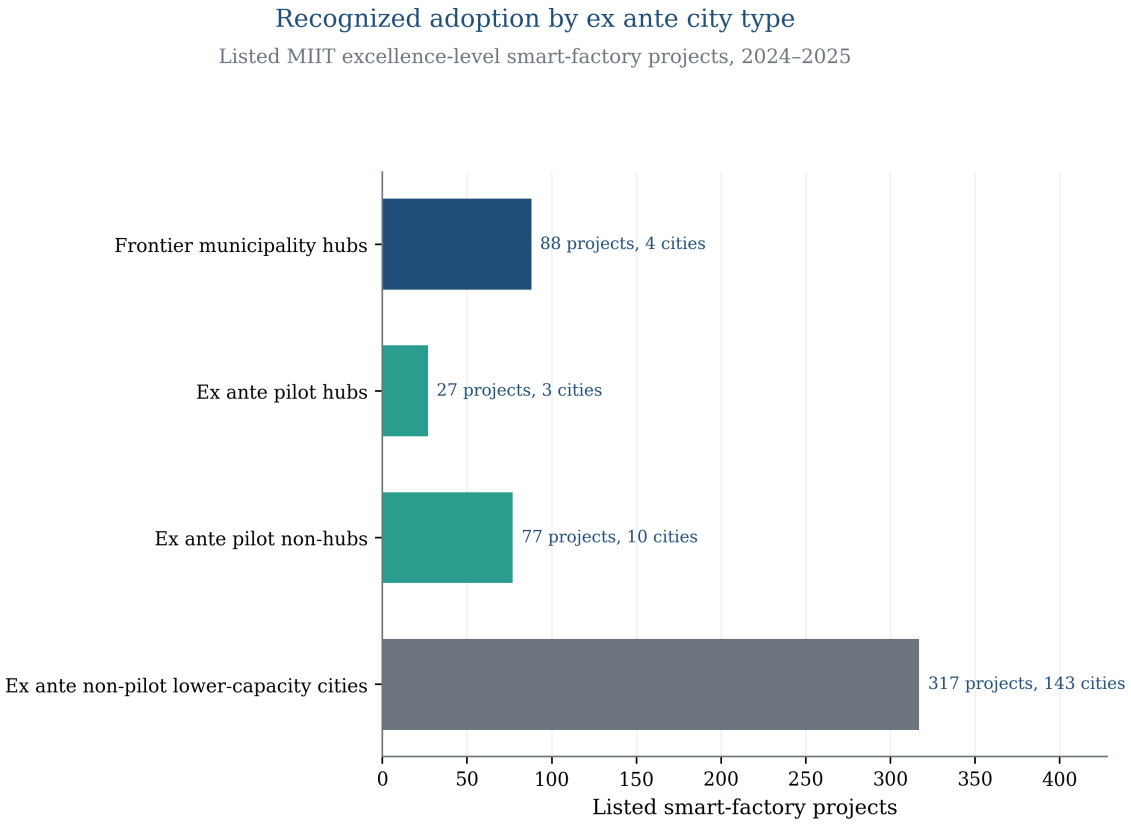


Figure 3: Listed smart-factory projects by ex ante diffusion-state city type. City classifications are defined before observing smart-factory outcomes and separate direct-admin municipalities, pilot industrial hubs, non-pilot industrial hubs, and lower-intensity adoption cities.

City typologies reinforce the interpretation by distinguishing several forms of adoption geography. Some pilot-zone cities appear as major adoption hubs with many listed projects. Some non-pilot cities also show substantial recognized adoption because they have strong manufacturing bases or sectoral specialization. Many cities have limited or no listed smart-factory recognition during the observed period. This distribution indicates that industrial AI diffusion is organized through a layered geography of policy designation, industrial capacity, administrative status, and sectoral opportunity.

The hub architecture result defines the paper's main empirical contribution. Pilot-zone designation is aligned with recognized industrial adoption, and this alignment is partly carried by cities with the deepest implementation capacity. The evidence therefore supports an architectural and measurement-based interpretation of China's AI diffusion state. National AI strategy becomes visible in a concentrated, hub-mediated pattern of recognized industrial adoption, where formal designation, municipal capacity, industrial ecosystems, and sectoral opportunity interact.

Table 5: Listed projects by diffusion-state city type

Diffusion type	Listed projects
Non-pilot low adoption	277
Frontier municipality hub	88
Pilot non-hub	77
Non-pilot industrial hub	40
Pilot industrial hub	27

Production table source: `table_E_city_typology.csv`. This table provides descriptive city-typology evidence. Labels use realized smart-factory counts for hub ranking.

Table 6: Ex ante city typology and listed projects

Ex ante diffusion type	Cities	Listed projects
Ex ante non-pilot low capacity	143	317
Frontier municipality hub	4	88
Ex ante pilot non-hub	10	77
Ex ante pilot hub	3	27

Production table source: `table_E_ex_ante_city_typology.csv`. This table provides ex ante city-typology evidence using pilot status, direct-admin status, and mega-hub flags before outcome-based hub ranking.

9 Sectoral selectivity

Industrial AI diffusion depends on the fit between AI capabilities and sector-specific production systems. The city-level evidence shows where recognized smart-factory adoption appears, while the economic meaning of that geography depends on the industries in which adoption is visible. AI-enabled systems are more likely to matter in production settings where firms generate machine-readable data, operate complex equipment, face quality-control constraints, manage high-throughput logistics, coordinate multi-stage production, or rely on inspection, prediction, scheduling, and process optimization. The paper therefore studies whether recognized smart-factory projects are concentrated in sectors where AI has plausible production complementarities.

The ex ante industry exposure measure captures technological compatibility before observing smart-factory outcomes. Sectors receive higher exposure when their production processes are naturally connected to machine vision, predictive maintenance, industrial robotics, digital twins, process control, warehouse automation, smart logistics, semiconductor production,

Table 7: Ex ante industry exposure heterogeneity

Term	Coefficient	Std. err.	p-value
Intercept	0.83	0.13	< 0.001
Pilot zone	-0.17	0.10	0.093
High ex ante exposure	-0.06	0.08	0.434
Pilot zone $\times$ high ex ante exposure	0.26	0.08	< 0.001

Production table source: `table_F_ex_ante_industry_heterogeneity.csv`. This table provides suggestive mechanism evidence outside the causal-identification layer.

battery manufacturing, automotive systems, machinery, chemicals, and pharmaceuticals. This measure is defined independently of the realized distribution of listed projects, which preserves the distinction between the mechanism under study and the outcome used to measure recognized adoption.

The sectoral distribution of recognized smart-factory projects supports the view that policy-legible diffusion is selective across manufacturing. Projects appear more often in industries where AI-enabled systems can plausibly improve inspection, reliability, throughput, energy use, scheduling, or process stability. Electronics and equipment-related sectors provide natural environments for machine vision and automated quality control. Battery and automotive supply chains provide settings where process consistency, traceability, and production optimization are economically valuable. Chemicals, pharmaceuticals, and advanced materials provide settings where monitoring, control, and predictive systems can matter for quality and yield.

The positive interaction between pilot-zone status and high ex ante exposure indicates that recognized adoption is more visible where policy geography and production compatibility coincide. This result is consistent with the diffusion-state framework because AI adoption in industrial production depends on instrumented machines, usable data, equipment interfaces, industrial software, workers and managers able to redesign routines, and firms with incentives to invest in integration. These complements differ across sectors, which means that diffusion capacity should be measured through the interaction of policy geography and industrial composition.

The sectoral results help explain why hubs matter. Major adoption cities often combine local implementation capacity with sectoral opportunity. A hub with electronics, advanced equipment, automotive, battery, or semiconductor activity may provide the firms, suppliers, engineers, and service providers needed to deploy AI-enabled manufacturing systems. City-level concentration and sectoral selectivity are therefore connected. The diffusion state operates through the fit between local industrial structure and the production problems that

AI can address.

The sectoral analysis provides mechanism-oriented evidence for the diffusion-state interpretation. Recognized adoption appears in industries where AI-enabled production technologies have clear operational use cases, including inspection, control, prediction, scheduling, logistics, and automation. This sectoral coherence strengthens the view that the observed smart-factory projects capture a meaningful trace of industrial AI diffusion.

The evidence connects three layers of the paper's argument. Pilot zones and smart-factory recognition reveal an institutional overlap. Hub-exclusion results show that this overlap is concentrated in cities with deep implementation capacity. Sectoral exposure shows that recognized adoption appears in industries where AI has plausible production complementarities. Together, these patterns suggest that China's industrial AI diffusion is organized through the interaction of national designation, municipal capacity, and sectoral fit.

10 Export relevance

The sectoral pattern of smart-factory recognition matters because industrial AI diffusion is connected to the composition of national production and trade. AI-enabled manufacturing systems are most strategically relevant when they appear in sectors that shape a country's position in global supply chains, advanced manufacturing, and export competition. The paper therefore examines whether recognized smart-factory projects appear in sectors that are relevant to China's export structure. This analysis provides industrial context for the adoption geography while treating trade outcomes as outside the completed empirical design.

The export-relevance analysis links smart-factory sectors to manufacturing categories that matter for China's global production position. These include electronics, electrical equipment, machinery, automotive systems, batteries, chemicals, pharmaceuticals, and other production domains where process quality, scale, inspection, logistics, and equipment utilization affect production performance. In these sectors, AI-enabled systems can plausibly support quality control, predictive maintenance, production scheduling, energy optimization, defect detection, inventory coordination, and automated inspection. The relevance of recognized adoption therefore depends partly on whether listed projects appear in sectors where such capabilities could matter for production.

The descriptive evidence shows that smart-factory recognition is connected to export-relevant manufacturing sectors. This pattern links the observed projects to industrial capacity as well as administrative classification. Recognized adoption appears in parts of the economy where production systems are complex, internationally exposed, and strategically

important. The export-relevance table positions smart-factory recognition within China's broader manufacturing base and shows that the diffusion architecture is connected to sectors that matter for industrial strategy.

Table 8: Smart-factory recognition and the 2024 export basket

Sector group	Projects	SF share	Export share	Share gap
Industrial machinery	247	48.5%	23.1%	25.4%
Petrochemicals	53	10.4%	4.9%	5.5%
Steel and metals	45	8.8%	6.3%	2.5%
Shipbuilding	19	3.7%	1.2%	2.5%
Pharmaceuticals	6	1.2%	0.6%	0.6%
Batteries	14	2.8%	2.8%	0.0%
Autos and NEV	37	7.3%	7.6%	-0.4%
Semiconductors and electronics	4	0.8%	8.1%	-7.4%
AI servers and computing equipment	48	9.4%	45.3%	-35.8%
Other	36	7.1%	—	—

Production table sources: `table_G_export_relevance.csv` and `table_H_export_sector_share_comparison.csv`. Export shares use BACI HS17 descriptives (CEPII, 2024). This table provides descriptive industrial context and leaves changes in export outcomes to future outcome-level designs.

The export analysis helps distinguish diffusion capacity from frontier AI capability. Frontier capability is often measured through model performance, research output, talent, investment, or compute scale. Diffusion capacity becomes visible when AI-related systems enter sectors where production constraints matter. A country can accumulate frontier capability while moving slowly into export-relevant industrial practice, and a country can strengthen manufacturing capabilities through adoption channels that receive little attention in frontier-model rankings. The export-relevance evidence therefore supports the paper's broader claim that AI competition has a production-side dimension that requires separate measurement.

The interpretation follows the paper's evidence hierarchy. The analysis asks whether the sectors in which recognized adoption appears are economically and strategically meaningful. Export performance reflects global demand, prices, exchange rates, firm capabilities, supply-chain organization, industrial policy, and preexisting sectoral advantages. The paper therefore uses export relevance as context for the diffusion architecture, while leaving export growth, export quality, firm productivity, and trade upgrading for future work with outcome-level

designs.

The export-relevance results complete the paper's core empirical chain. Pilot-zone status and smart-factory recognition overlap in city-level data. The association is concentrated in hubs with stronger implementation capacity. Recognized projects appear in sectors where AI has plausible production complementarities. Those sectors are connected to China's export-relevant manufacturing base. Taken together, these facts suggest that the AI diffusion state is an institutional architecture linking national AI priorities to local production systems and strategically important sectors of the industrial economy.

11 Economic interpretation

The empirical results identify the policy-legible portion of AI diffusion as an institutional and spatial process. Pilot-zone cities have higher listed smart-factory counts, yet the association weakens when the largest hubs and administratively distinctive cities are removed. This pattern indicates that policy designation and local capacity operate together. National AI strategy becomes economically meaningful when it reaches cities with the firms, suppliers, engineers, infrastructure, administrative routines, and industrial sectors able to absorb production technologies.

The hub architecture is consistent with the economics of general-purpose technology adoption. General-purpose technologies require complementary investments that are costly to coordinate and unevenly distributed across space. AI in manufacturing requires sensors, robotics, data infrastructure, machine interfaces, industrial software, process knowledge, managerial redesign, and workers able to operate new systems. Cities with deep industrial ecosystems can combine these complements more easily than cities with thinner production bases. The observed attenuation under hub exclusions therefore indicates that much of the pilot-zone association is carried by places where complementary capital and implementation capacity are already concentrated.

The results imply that national AI advantage can diverge across the capability and diffusion margins. A country may lead in frontier models while moving slowly into production deployment, because frontier systems do not automatically reorganize factories, supply chains, procurement systems, or local industrial routines. A country may also accumulate advantage through the ability to embed available AI capabilities across manufacturing systems when frontier leadership is distributed across multiple actors. The diffusion-capacity margin is therefore central to AI competition because it governs the conversion of technical possibility into productivity-relevant use.

China's case illustrates how state capacity can shape the environments in which firms

coordinate, finance, implement, and display AI-enabled industrial upgrading. The relevant state function is the construction of adoption environments through city designations, demonstration platforms, industrial parks, standards, recognition systems, and municipal implementation channels. These instruments can reduce uncertainty for firms, concentrate attention, create exemplars, and make industrial upgrading legible to local officials and national ministries. The diffusion state therefore operates as a coordination architecture for complementary investments.

Administrative recognition matters economically because it turns dispersed firm-level upgrading into a visible policy category. Many industrial AI deployments are internal to production systems and difficult to observe from outside the firm. Recognition lists create a public trace of projects that have become legible to the industrial policy apparatus. This legibility can affect incentives, reputation, local government attention, and the circulation of implementation models across firms and regions. The paper uses recognition as a measurement trace, while also recognizing that classification systems can shape what firms and officials treat as exemplary adoption.

The sectoral results show that diffusion capacity is constrained by production compatibility. AI is more easily embedded in sectors where production generates data, uses expensive equipment, requires inspection, depends on scheduling, and benefits from predictive control. This means that the economic geography of AI adoption depends on the industrial composition of cities. Pilot-zone status has different implications in a city with electronics, batteries, machinery, automotive systems, chemicals, or pharmaceuticals than in a city whose production base offers fewer AI-complementary processes. Diffusion capacity therefore has a sectoral structure as well as a spatial structure.

The export-relevance evidence gives the diffusion architecture strategic significance without turning the paper into a trade-effects study. Recognized smart-factory projects appear in sectors connected to China's export-relevant manufacturing base, which means that adoption traces are linked to domains where production capability affects international competitiveness. The paper leaves export effects, productivity gains, and quality upgrading for future outcome-level designs. The evidence here shows that the architecture of recognized adoption is connected to sectors where such effects would matter if AI-enabled production systems improve performance. This connection makes diffusion capacity a relevant object for the economics of industrial strategy.

The broader implication is that AI competition should be analyzed as a two-margin problem. The first margin is frontier capability, which concerns who can build the most capable systems. The second margin is diffusion capacity, which concerns who can absorb those systems into productive structures. Economists have long studied adoption lags,

complementarities, productivity dispersion, and spatial concentration. AI makes these problems more urgent because the technology is general, fast-moving, and deeply dependent on organizational redesign. The diffusion-state framework provides a way to measure this second margin with administrative traces, city-level evidence, sectoral compatibility, and disciplined scope of inference.

12 Robustness and extensions

The robustness exercises evaluate whether the hub-centered interpretation remains visible across alternative samples, partial public controls, and supplementary evidence on industrial AI activity. Each exercise remains in its appropriate claim tier. Hub exclusions provide the main architecture diagnostic. Public controls provide an appendix-only check in a restricted 2024 sample. Tiered patent geography provides supplementary triangulation under explicit coverage limits.

The first robustness layer examines dependence on major hubs and administratively distinctive cities. The pilot-zone coefficient attenuates when Beijing, Shanghai, Shenzhen, and Hangzhou are removed, when direct-admin municipalities are excluded, and when the top smart-factory cities are removed. These checks strengthen the interpretation that the observed pattern reflects a hub-centered diffusion architecture. They also discipline the interpretation by showing that the full-sample association is partly carried by cities with unusual administrative status, industrial scale, research capacity, firm concentration, and digital-economy capacity.

The public-controls appendix adds a partial city-level control exercise using China City Statistical Yearbook tables obtained through the public ChinaUTC path. This specification includes GDP, population, secondary-industry structure, and limited 2024 coverage for foreign trade, telecom, and industrial-output proxies. In the 2024 public-control subset, pilot-zone status remains positively associated with listed smart-factory counts in OLS count and log-count specifications, while the Poisson count specification is statistically indistinct from zero. This result is useful as a public-data robustness check because it shows that the pilot-zone association remains visible in the restricted 2024 sample with available city covariates.

The public-controls appendix remains separate from the strict production-control design. The public ChinaUTC bundle lacks foreign direct investment and fixed-asset investment at the coverage required for the strict EPS/NBS specification, and the available controls cover only a smaller 2024 subset of cities. The paper therefore reports this exercise as Table I in the appendix and avoids treating it as equivalent to a full controlled panel. This distinction increases transparency without overstating the strength of the control design.

The tiered patent Atlas provides a supplementary robustness extension for industrial AI activity. The project imported a real OpenXLab IIDS-derived patent layer and constructed a tiered applicant-geography layer that reaches approximately 65.4 percent city coverage across the IIDS patent keys. The geography layer combines high-confidence institutional and corporate location rules, curated applicant mappings, and applicant-name city tokens, then reports coverage by confidence tier. Because exact publication-number applicant-address geography remains unavailable, this extension is used as coverage-qualified triangulation and excluded from the main identification layer.

The patent extension connects the smart-factory adoption architecture to a second empirical trace of industrial AI activity. Patent records capture inventive and technical activity, while smart-factory recognition captures administratively recognized industrial adoption. These two traces describe different margins of the AI economy. When the patent layer is reported by geography-confidence tier, it can help assess whether industrial AI activity appears in a spatial structure consistent with the smart-factory evidence. Its value comes from triangulation under transparent uncertainty.

The patent extension also demonstrates the importance of claim-tiering in empirical AI research. The project constructs a large IIDS patent layer and a tiered city assignment, while the exact applicant-address geography required for a publication-ready patent evidence chain remains unavailable. The paper therefore excludes patent F1 estimates from the main results, describes the patent layer as tiered geography, and prevents the patent extension from supporting causal pilot-zone claims. This restraint preserves the distinction between robustness evidence and identification evidence.

The robustness results support the main interpretation while clarifying its boundaries. Hub exclusions show that the pilot-zone association is embedded in a hierarchy of high-capacity cities. Public controls show that the association remains visible in a restricted 2024 sample with available city covariates, while leaving the strict EPS/NBS design for future work. The tiered patent Atlas provides additional triangulation for industrial AI activity, while keeping exact geography and publication-ready patent estimates outside the paper's main claims. Together, these exercises strengthen the measurement contribution because they show how the diffusion-state architecture can be examined without collapsing different forms of evidence into a single claim.

13 Scope of inference

The paper's empirical contribution depends on a clear separation between measured architecture, descriptive association, mechanism evidence, robustness, and causal identification. The

analysis documents where recognized industrial AI adoption appears, how that geography aligns with national AI pilot-zone designation, how the association changes when major hubs are removed, and which sectors are most compatible with AI-enabled production systems. These patterns support an interpretation of China's industrial AI diffusion as a hub-centered architecture.

Pilot-zone designation is interpreted as an institutional marker embedded in prior capacity. Designated cities often possess manufacturing depth, administrative status, digital firms, universities, industrial parks, and local policy experience before designation. The empirical design therefore interprets pilot-zone status as part of a policy-adoption geography. The baseline association between designation and smart-factory recognition is informative because it reveals where national AI policy and recognized industrial upgrading align.

Smart-factory recognition is interpreted as a visible administrative trace of industrial upgrading. Recognition indicates that a project entered an official category of intelligent manufacturing, which makes it useful for measuring the legible portion of adoption. The lists capture the policy-legible portion of recognized adoption within the industrial policy apparatus. Many firms may deploy AI, robotics, inspection systems, scheduling software, or process-control tools without appearing in the recognition lists. The paper therefore studies the visible adoption trace instead of the complete universe of firm-level industrial AI use.

City assignment is treated as an evidence-classified measurement task. Official-location exact assignments provide the strongest location evidence, externally verified assignments add non-list support, and rule-based city inferences expand coverage with a weaker evidence tier. The stratified audit supports the official-location layer and records uncertainty in the rule-based layer. The empirical analysis uses the full resolved dataset while preserving the evidence classes that make the spatial measurement transparent.

Hub attenuation is interpreted as evidence for a concentrated diffusion architecture. The decline in the pilot-zone association after excluding Beijing, Shanghai, Shenzhen, Hangzhou, direct-admin municipalities, or the top smart-factory cities shows that the full-sample pattern depends meaningfully on high-capacity hubs. This evidence supports the interpretation that municipal capacity and industrial ecosystem depth shape recognized adoption. It leaves room for meaningful adoption activity in non-hub cities and for heterogeneity within the pilot-zone group.

Sectoral exposure provides mechanism evidence about technological compatibility. Ex ante AI exposure identifies industries where machine vision, robotics, predictive maintenance, digital twins, process control, smart logistics, or automated inspection are especially plausible. This analysis helps explain why recognized adoption appears in some parts of manufacturing more than others. It leaves estimation of sector-level productivity gains, firm-level returns,

and welfare effects from smart-factory recognition to future work.

Export relevance provides strategic context for the industrial distribution of recognized adoption. The export tables show that many smart-factory sectors are connected to parts of China's manufacturing base that matter for global production and external competitiveness. This finding helps interpret why the diffusion architecture matters for political economy and industrial strategy. Claims about export growth, export quality, and changes in China's trade position require outcome-level designs beyond the scope of the current evidence.

Public city controls provide an appendix-only robustness exercise. The ChinaUTC public-control specification adds GDP, population, secondary-industry structure, and limited 2024 proxies for foreign trade, telecom, and industrial output. The restricted public sample is useful because the pilot-zone association remains visible in OLS count and log-count specifications. The missing FDI and fixed-asset investment controls keep this appendix exercise outside the strict EPS/NBS production-control design.

The patent Atlas provides a coverage-qualified robustness extension. The IIDS patent pipeline supplies a large industrial AI patent layer, and the tiered applicant-geography procedure provides supplementary city assignment for a substantial share of records. The current patent layer uses tiered applicant geography because exact publication-number applicant-address geocoding remains unavailable. The paper therefore reports patent geography by confidence tier and excludes patent F1 estimates from the main results.

The paper's central claim is architectural and measurement-based. China's AI diffusion state becomes visible through the alignment of pilot-zone designation, smart-factory recognition, hub concentration, sectoral compatibility, and export-relevant manufacturing. The evidence shows how national AI strategy becomes legible in the geography of recognized industrial adoption. Full causal evaluation, firm productivity outcomes, exact patent geocoding, and strict production-control specifications remain future work with the required data.

14 Conclusion

The economic significance of artificial intelligence depends on the capacity to convert technical capability into deployed production systems. Measures of frontier capability capture who can build powerful models, assemble compute, attract talent, and produce research advances. Measures of diffusion capacity capture a different margin of AI competition, the ability to embed AI into firms, cities, sectors, and industrial routines. This paper develops a framework for measuring that second margin and applies it to China's industrial AI strategy.

China's AI diffusion state appears in the data as a hub-centered adoption architecture. National AI pilot-zone designation aligns with recognized smart-factory adoption, and

pilot-zone cities have substantially higher listed project counts than non-pilot cities. The association attenuates when major hubs and administratively distinctive cities are removed, which indicates that municipal capacity, industrial ecosystem depth, firm concentration, and sectoral opportunity shape the geography of recognized adoption. The evidence points to an architecture in which national strategy becomes visible through cities that can implement, absorb, and display industrial AI upgrading.

The sectoral and export-relevance results explain why the observed geography matters for industrial political economy. Recognized smart-factory projects appear in sectors where AI-enabled production systems have plausible complementarities, including inspection, predictive maintenance, robotics, logistics, digital twins, process control, batteries, automotive systems, machinery, chemicals, pharmaceuticals, electronics, and advanced equipment. Many of these sectors are connected to China's export-relevant manufacturing base. The diffusion state therefore embeds AI into parts of the economy where production capability and strategic competition are closely connected.

The evidence hierarchy preserves the paper's central distinction between measured architecture and causal evaluation. Smart-factory recognition captures the legible portion of industrial upgrading within the policy apparatus. Pilot zones identify an institutional geography embedded in prior capacity. Public city controls remain partial, and the patent Atlas is reported as a tiered robustness extension because exact publication-number applicant geography remains unavailable. These limits define the interpretation of the results and clarify the contribution. The paper measures an adoption architecture while leaving causal evaluation, productivity outcomes, changes in export outcomes, and exact patent geography to future work with the required data.

The broader implication is that AI competition should be measured through diffusion capacity as well as frontier capability. A country can possess advanced models while lacking the institutions, firms, sectors, and local implementation systems needed to translate AI into production. A state can also build strategic advantage through the administrative and industrial channels that make adoption visible, scalable, and compatible with national production systems. For economists of AI, this suggests that adoption traces, institutional designations, sectoral compatibility, and local implementation capacity should become central objects of empirical study.

The diffusion-state framework opens a path for comparative research beyond China. Similar measurement strategies could study how different political economies convert AI into manufacturing, logistics, energy, health, finance, public administration, and defense production. Future work can connect administrative adoption traces to firm-level productivity, investment, labor demand, export performance, and supply-chain resilience. The central

challenge is to measure the movement from capability to use, because the economic effects of AI will be determined by the institutions that make the frontier productive.

A City-resolution evidence and audit

The city-resolution procedure assigns all 509 MIIT excellence-level smart-factory projects to cities and records the evidence tier supporting each assignment. The evidence split contains 102 official-location exact assignments, 357 rule-based text inferences, and 50 externally verified assignments. Official-location exact assignments are based on locations stated in the registry or source material. Rule-based text inferences use structured project text, firm names, industrial-park names, location-bearing descriptions, and related textual signals. Externally verified assignments use non-list sources such as company pages, annual reports, project registries, and industrial-park records.

Table 9: City-resolution evidence quality

Resolution class	Evidence type	Projects	Share
external evidence verified	all	50	9.8%
official location exact	all	102	20.0%
rule based text inference	all	357	70.1%
external evidence verified	company annual report	11	2.2%
external evidence verified	company site registry	23	4.5%
external evidence verified	industrial park page	1	0.2%
external evidence verified	project registry	15	2.9%
official location exact	miit location field	102	20.0%
rule based text inference	embedded html table	28	5.5%
rule based text inference	firm embedded city token	64	12.6%
rule based text inference	firm parenthetical	22	4.3%
rule based text inference	firm province county	1	0.2%
rule based text inference	firm registry match	221	43.4%
rule based text inference	html p tag rtl location	17	3.3%
rule based text inference	project branch city	4	0.8%

Production table source: `table_B_city_resolution_quality.csv`. This table provides the full evidence-type breakdown for the city-resolution procedure.

The stratified audit samples 20 official-location rows and 50 rule-based rows. The audit confirms 20 of 20 official-location rows, confirms 20 rule-based rows, marks 30 rule-based rows as insufficient evidence, and records zero incorrect assignments. The audit gives the strongest support to the official-location layer and records greater uncertainty in the rule-based layer. The paper uses the full resolved dataset for descriptive measurement, while preserving the evidence classes that allow the reader to evaluate spatial reliability.

B Public fallback controls

The public-controls appendix reports a partial 2024 cross-section using ChinaUTC public city data. The specification includes GDP, population, secondary-industry structure, and limited 2024 coverage for foreign trade, telecom, and industrial-output proxies. The public bundle excludes the FDI and fixed-asset investment controls required for the strict EPS/NBS specification, so Table 10 is reported as an appendix-only robustness exercise.

Table 10: Public fallback controls

Model	Pilot-zone coefficient	p-value	Complete cases
OLS count	1.58	0.020	51
OLS log count	0.50	0.018	51
Poisson count	0.23	0.428	51

Production table source: `table_I_appendix_public_fallback_controls.csv`. This table provides appendix-only robustness evidence. The specification sits outside the strict EPS/NBS specification because FDI and fixed-asset investment controls are unavailable in the public fallback bundle.

The public-control results show that the pilot-zone association remains visible in the restricted 2024 public-control sample under OLS count and log-count specifications. The Poisson count specification is statistically indistinct from zero. These results are useful for transparency because they add observable city covariates to the descriptive association, while preserving the distinction between public-data robustness and a full controlled panel design.

C Tiered patent geography robustness

The IIDS patent extension provides a robustness-only layer for industrial AI activity. Industrial-AI patent records carry a tiered applicant-geography layer frozen at approximately 65.4 percent city fill. The extension uses tiered geography instead of exact publication-number applicant-address geocoding, and it is excluded from publication-ready patent F1 estimates and causal pilot-zone claims.

The patent extension provides a second trace of industrial AI activity under explicit coverage limits. Patent records capture inventive and technical activity, while smart-factory recognition captures administratively recognized industrial adoption. The two traces describe different margins of the AI economy. The tiered geography layer is useful because it allows supplementary triangulation of spatial patterns while keeping exact geocoding and patent-based identification outside the main claim hierarchy.

Table 11: Tiered patent geography robustness

Geography tier	Share of keys	City fill
Official headquarters page	43.9%	100%
University location	12.1%	100%
Applicant-name city token	9.3%	100%
Unresolved	34.6%	0%

Production sources: `table_P14_tiered_geography_coverage_by_confidence.csv` and `table_P17_tiered_geography_tier_breakdown.csv`. This table provides robustness-only patent geography evidence because exact publication-number applicant-address geography remains unavailable.

D Timing and balance diagnostics

The timing and balance diagnostics provide additional context for interpreting the descriptive associations. The balance figure reports standardized mean differences for available public-control variables. The timing figure reports the relationship between pilot-zone designation and the available smart-factory recognition window. These diagnostics provide descriptive support for the available recognition window because the public excellence-level smart-factory lists begin in 2024.

E Reproducibility and evidence package

The paper separates claim tiers to make the scope of inference explicit. Pilot-zone endogeneity is addressed by interpreting designation as an institutional marker embedded in prior capacity. Smart-factory recognition is addressed by treating the lists as a visible administrative trace of intelligent manufacturing within the policy apparatus. City-resolution uncertainty is addressed through evidence classes, external verification, and a stratified audit. Partial public controls are addressed by placing ChinaUTC controls in the appendix and by leaving the strict EPS/NBS specification outside the completed evidence base. Patent geography uncertainty is addressed by reporting the IIDS patent layer only as a tiered robustness extension. Export relevance is addressed by treating export tables as descriptive industrial context.

The reproducibility package documents the evidence pipeline, paper-facing artifacts, and validation steps. The main evidence package is refreshed by `make pcs-paper-owner` and validated by `make validate-submission`. The tiered patent robustness layer is refreshed by `make atlas-iids-tiered-extension` and remains separate from the exact evidence chain. The submission bundle contains the manuscript, paper-facing tables, figures, claim maps, gate reports, data availability statement, and reproducibility statement.

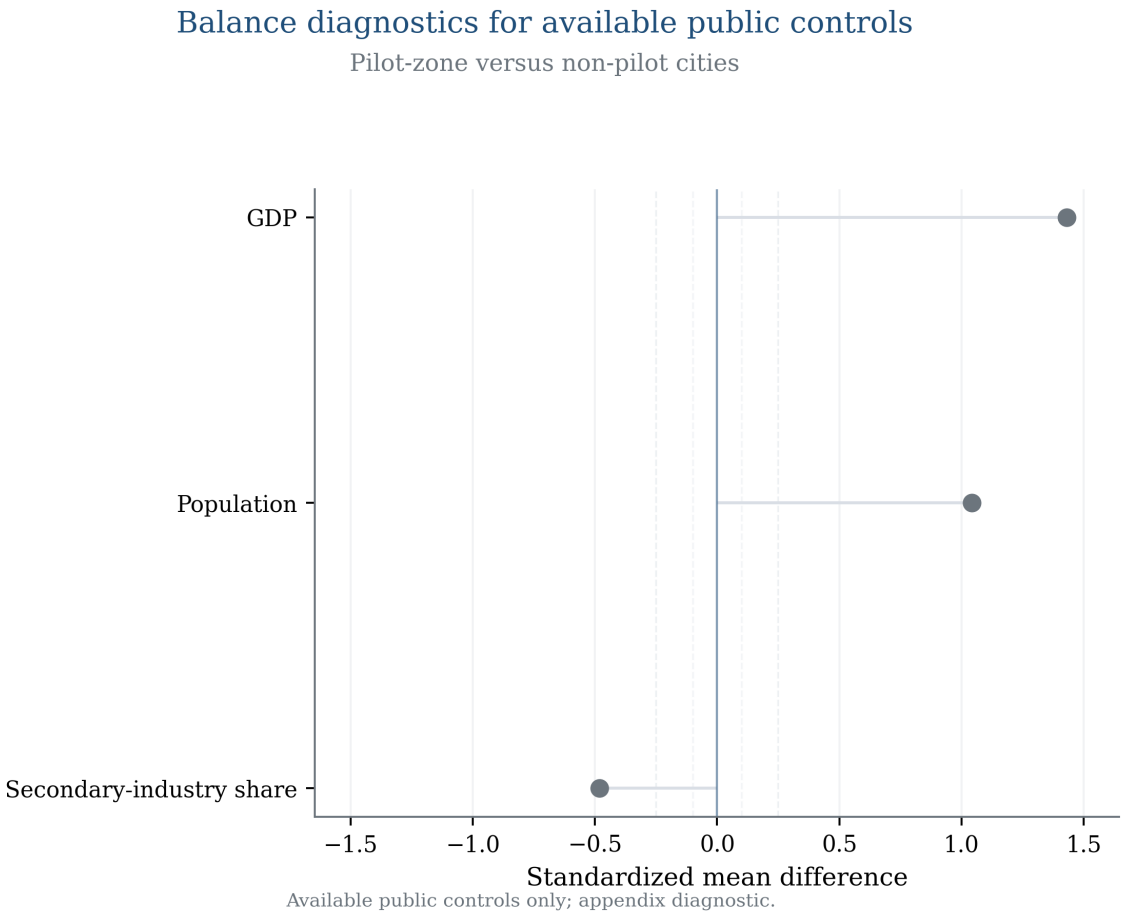


Figure 4: Standardized mean differences for pilot and non-pilot cities. The figure reports balance diagnostics for available public-control variables and should be interpreted as controls-dependent descriptive evidence.

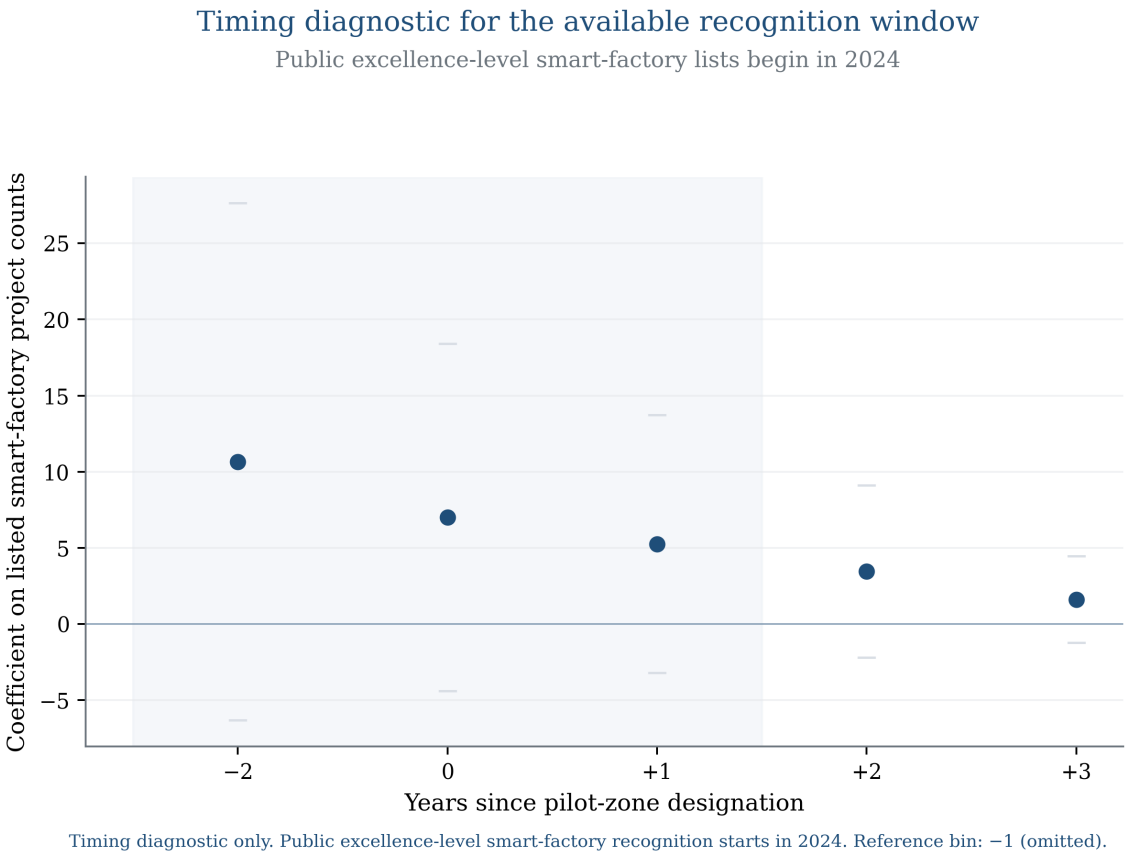


Figure 5: Timing diagnostic for pilot-zone designation and listed smart-factory recognition. Pre-2024 smart-factory counts are zero-filled because the public excellence-level lists begin in 2024. The figure serves as a timing diagnostic for the available recognition window.

References

Aghion, P., Dewatripont, M., Du, L., Harrison, A., and Legros, P. (2015). Industrial policy and competition. *American Economic Journal: Macroeconomics*, 7(4):1–32.

Aghion, P., Jones, B. F., and Jones, C. I. (2019). Artificial intelligence and economic growth. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, pages 237–282. University of Chicago Press, Chicago.

Agrawal, A., Gans, J., and Goldfarb, A., editors (2019). *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, Chicago.

Agrawal, A., Gans, J., and Goldfarb, A. (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press, Boston.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Besley, T. and Persson, T. (2010). State capacity, conflict, and development. *Econometrica*, 78(1):1–34.

Brandt, L., Van Biesebroeck, J., Wang, L., and Zhang, Y. (2023). Wto accession and performance of chinese manufacturing firms. Working Paper 31883, National Bureau of Economic Research. Used as background on Chinese manufacturing and firm-level upgrading.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies "engines of growth"? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative ai at work. Working Paper 31161, National Bureau of Economic Research.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Center for Security and Emerging Technology (2020). China's national new generation artificial intelligence innovation and development pilot zones. https://cset.georgetown.edu/. Background source on China's AI pilot-zone system; accessed 2026-06-01.

CEPII (2024). Baci: International trade database at the product-level. https://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=37. Version used for 2024 HS17 export descriptives; accessed 2026-06-01.

Cockburn, I. M., Henderson, R., and Stern, S. (2019). The impact of artificial intelligence on innovation. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, pages 115–146. University of Chicago Press, Chicago.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.

Criscuolo, C., Martin, R., Overman, H. G., and Van Reenen, J. (2019). Some causal effects of an industrial policy. *American Economic Review*, 109(1):48–85.

Decarolis, F., Giorgiantonio, C., and Giovanniello, V. (2020). Corruption red flags in public procurement: New evidence from italian calls for tenders. *EPJ Data Science*, 9(1):16.

Duranton, G. and Puga, D. (2014). The growth of cities. In Aghion, P. and Durlauf, S. N., editors, *Handbook of Economic Growth*, volume 2, pages 781–853. Elsevier.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv:2303.10130.

Glaeser, E. L. and Resseger, M. G. (2010). The complementarity between cities and skills. *Journal of Regional Science*, 50(1):221–244.

International Federation of Robotics (2023). World robotics 2023: Industrial robots. https://ifr.org/worldrobotics/. Executive summary and statistical report; accessed 2026-06-01.

Juhász, R., Lane, N. J., and Rodrik, D. (2024). The new economics of industrial policy. *Annual Review of Economics*, 16:213–242.

Li, Q. (2024). Report on the work of the government. https://english.www.gov.cn/. Government work report and related policy materials on advanced manufacturing, artificial intelligence, and new productive forces; accessed 2026-06-01.

Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Lotufo, J. B., Shi, A., and Oak, S. (2025). Artificial intelligence index report 2025. Technical report, Stanford Institute for Human-Centered Artificial Intelligence. AI Index annual report.

Ministry of Industry and Information Technology of the People's Republic of China (2024). Excellence-level smart factory recognition list. https://www.miit.gov.cn/. Official MIIT public list used to construct the 2024 smart-factory project register; accessed 2026-06-01.

Ministry of Industry and Information Technology of the People's Republic of China (2025). Excellence-level smart factory recognition list. https://www.miit.gov.cn/. Official MIIT public list used to construct the 2025 smart-factory project register; accessed 2026-06-01.

Naughton, B. (2021). *The Rise of China's Industrial Policy, 1978 to 2020*. Lynne Rienner Publishers, Boulder, CO.

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

Rogers, E. M. (2003). *Diffusion of Innovations*. Free Press, New York, 5 edition.

Syverson, C. (2011). What determines productivity? *Journal of Economic Literature*, 49(2):326–365.

Xinhua News Agency (2021). China builds national new generation artificial intelligence innovation and development pilot zones. http://www.xinhuanet.com/. News source on the expansion of the national AI pilot-zone system to 17 units; accessed 2026-06-01.