The Impact of Public Data on AI Innovation: Open Government Data Initiatives Encourage Local Healthcare AI Investment

Guohao Zhu, Zhejiang University
Aaron Cheng, London School of Economics
Ning Zhao, Zhejiang University

ABSTRACT

Public data infrastructure is often built to enhance transparency and accountability. Yet it is costly, its economic and innovation spillovers remain uncertain, and governments have responded with divergent data-opening strategies. We examine whether and how the launch of open government data portals (OGDPs) affects healthcare AI innovation, which spans upstream vendor development and downstream hospital investment. Healthcare is an apt setting: it is data-intensive, faces substantial regulatory uncertainty, and depends on both data availability and institutional support for AI deployment. Using a panel dataset of Chinese cities and a staggered difference-in-differences design, we measure downstream investment through hospitals' procurement of diagnostic and operational AI systems. We find that OGDP launches generate sustained increases in such investment, a result that holds across heterogeneity-robust estimators, matching, synthetic control, and alternative-explanation tests. Mechanism analyses point to two channels: data provision is a necessary baseline input, but governmental signaling emerges as the more proximate driver, as OGDPs convey credible policy commitment that reduces regulatory uncertainty and accelerates the translation of upstream AI development into downstream investment. The effect is further amplified where governments prioritize digitalization, healthcare demand is acute, and local AI vendors are abundant. Our findings position OGDPs as both digital infrastructure and institutional signals. We contribute to research on public data and AI innovation by distinguishing material from cognitive channels, to work on healthcare AI ecosystems by elevating the institutional layer, and to the literature on directed technical change, while offering guidance on realizing the broader impacts of public data.

Keywords: Artificial Intelligence; Public Data; AI Innovation; Healthcare AI Investment; Digital Infrastructure; Governmental Signaling; Open Government Data; Difference-in-Differences

Working Paper

May 2026

1. Introduction

Governments are the primary collectors of large-scale public data and increasingly release it through Open Government Data Portals (OGDPs), a key form of public data infrastructure (Lnenicka et al., 2024). The principal motivation for opening such data has long been to enhance transparency and accountability (Attard et al., 2015; Jetzek et al., 2019). More recently, however, a second rationale has gained prominence. Because AI development is fundamentally data-complementary (Agrawal et al., 2019; Jones and Tonetti, 2020), public data infrastructure carries the potential to stimulate AI innovation as a valuable spillover beyond its original public-governance purpose.

Yet this potential has not translated into convergent practice. Despite the substantial costs of building and maintaining OGDPs, their innovation benefits remain uncertain, and governments around the world have responded with markedly divergent strategies. The U.S. has recently reduced public data funding and dataset availability,[1] while China continues to expand its public data initiatives.[2] This divergence is puzzling from an institutional economics perspective: inclusive institutions foster innovation (Robinson and Acemoglu, 2012), and OGD initiatives are a paradigmatic example of such arrangements, as they lower barriers to access to a key input for AI. Why, then, do governments diverge so sharply on public data initiatives that theory predicts should encourage innovation?

This puzzle motivates a closer examination of the literature on public data infrastructure and AI innovation. Despite sustained investments in public data collection and release, and growing scholarly and policy attention, theory and evidence on whether and how public data infrastructure shapes AI innovation remain underdeveloped in two respects. First, prior research has treated data largely as a proprietary, field-held input: how privately accumulated data drives AI innovation (Jones and Tonetti, 2020), or how firms enhance AI commercial innovation by providing services to the state and gaining access to government data (Beraja et al., 2023). What remains underexplored is the broader institutional role of public data infrastructure, not as a private input acquired through state-firm exchange, but as an openly accessible resource that may shape AI innovation across an entire local ecosystem. Second, across studies of public data, documented effects on organizational innovation diverge. Some studies document positive effects (e.g., Nagaraj, 2022; Magalhaes and Roseira, 2020), while others suggest that unrestricted data access may erode competitive advantages, dilute returns to private investment, and weaken incentives to invest (e.g., Janssen et al., 2012; Zuiderwijk and Janssen, 2014). Taken together, the limited attention to public data

[1] Multiple federal agencies in the United States have recently reduced the availability of publicly accessible datasets, particularly in areas such as climate change, environmental justice, diversity, and health.

[2] In September 2024, China's central government issued a policy document that explicitly promotes the opening and sharing of public data. The National Data Administration also announced a new batch of public data demonstration scenarios in December 2025.

infrastructure and the inconsistent evidence on its innovation effects leave two questions unresolved: *Does public data infrastructure influence local AI innovation? If so, how?*

We address these questions by examining the impact of OGDPs on healthcare AI innovation. Specifically, we focus on healthcare AI investment as a critical and observable manifestation of AI innovation in a data-intensive and socially important domain. Innovation here extends beyond technological invention to commercialization across the value chain (Nambisan et al., 2017; Schumpeter and Swedberg, 2021). Along this chain, in the healthcare AI ecosystem, AI vendors innovate upstream by developing products, while hospitals contribute downstream by deploying them (Sun and Medaglia, 2019). In this study, we focus on the downstream stage, capturing hospitals' investment in healthcare AI, which is directly observable and reflects realized innovation outcomes rather than mere intentions. Using a panel dataset of 258 Chinese prefecture-level cities from 2012 to 2023, we then examine how local healthcare AI investment changes following OGDP launches.

Econometrically, we first implement a staggered difference-in-differences design exploiting the staggered rollouts of OGDPs across cities, finding a statistically significant and sustained increase in healthcare AI investment in treated cities (Table 2). The result holds across a battery of sensitivity checks (Table 3). We further employ recently developed event study methods from the econometrics literature (e.g., de Chaisemartin and d'Haultfoeuille, 2020; Sun and Abraham, 2021; Cengiz et al., 2019; Callaway and Sant'Anna, 2021), which yield estimates that remain robust even after relaxing the assumption of homogeneous treatment effects across cities and over time (Figure 4). The effect also persists after accounting for several potential confounding factors, including concurrent AI and data policies (Table C-2 in Appendix C), local digitalization capacity (Table C-3), and fiscal capacity (Table C-4). These analyses corroborate the sustained rise in healthcare AI investment in cities that adopted OGDPs.

What explains the baseline effect? Because healthcare AI investment emerges from interactions among key participants in the healthcare ecosystem (i.e., governments, hospitals, AI vendors), we propose two underlying mechanisms grounded in distinct theoretical traditions. First, OGDPs operate through a *data provision* channel. Building on the data and AI innovation literature emphasizing the data-complementary nature of AI (Agrawal et al., 2019; Jones and Tonetti, 2020), we argue that OGDPs alleviate data-related constraints in healthcare AI development by lowering vendors' costs of accessing large-scale, reliable data. This facilitates the development of more accurate and robust AI products and, in turn, strengthens hospitals' incentives to invest. Second, OGDPs operate through a *governmental signaling* channel. Drawing on the economics literature on uncertainty and investment showing that policy uncertainty discourages investment (Gulen and Ion, 2016) and that government signals can reduce such uncertainty (Connelly et al., 2011; Spence, 2002; Vasudeva, 2013; Vasudeva et al., 2018), we argue that OGDPs signal credible policy commitment to data sharing and data-related innovation. This signal eases

regulatory uncertainty for both hospitals and AI vendors, and ultimately encourages healthcare AI investment. Empirically, we find that while data availability is a necessary input, cross-city variation in data quantity and quality does not systematically condition the OGDP effect (Figure 5). Cross-city variation in signaling-related proxies, however, does (Figure 6), and this asymmetry persists when both mechanisms are estimated in a single specification (Figure 7). Governmental signaling thus appears to be the more prominent channel through which OGDPs translate into healthcare AI investment.

Additional analyses examine contextual factors that shape the boundary conditions of the main effect. Specifically, OGDPs are more likely to translate into substantive AI investment when local governments prioritize digitalization, as such contexts make public data initiatives more likely to be interpreted as sustained and consequential policy commitments rather than symbolic gestures. The effect is also stronger where healthcare systems face greater demand pressure, as capacity constraints heighten hospitals' urgency to adopt AI-enabled solutions. Finally, the effect is amplified in cities where local healthcare AI vendors are more abundant, as these firms are better positioned to interpret policy signals, leverage released data, and convert public data into deployable AI solutions (Table F-1 in Appendix F).

This work makes several contributions to scholarship and practice. First, it advances the emerging literature on data and AI innovation by elucidating how public data infrastructure shapes AI investment, complementing prior work on private data (Tilson et al., 2010; Agrawal et al., 2019; Jones and Tonetti, 2020) and firm-state data exchange (Beraja et al., 2023). In doing so, it shifts the focus from firm-held data to openly accessible infrastructure that stimulates AI investment across an entire local ecosystem. Second, it extends research on healthcare AI from intra-organizational adoption barriers to the broader institutional environment (e.g., Abdel-Karim et al., 2023; Liang and Xue, 2022), showing that government-led open data initiatives coordinate ecosystem actors through both material (data) and cognitive (signaling) channels. Third, it contributes to the literature on institutions and directed technical change (Acemoglu et al., 2012) by demonstrating how public data infrastructure can steer investment toward socially valuable AI applications. This positions open data as a policy lever for directing, not merely enabling, the trajectory of AI innovation. Finally, our findings offer practical insights by showing that public data infrastructure, when coupled with credible government commitment, can align incentives across policymakers, AI vendors, and hospitals, thereby accelerating healthcare AI investment.

The remainder of the paper is organized as follows. Section 2 reviews the literature on open government data portals and healthcare AI investment and develops the theoretical framework and hypotheses. Section 3 describes the empirical context, data, and descriptive statistics. Section 4 presents the empirical strategy, main results, robustness checks, and economic significance. Section 5 explores the underlying mechanisms and reports the corresponding empirical tests. Section 6 examines heterogeneous effects, and Section 7 concludes with a discussion of the implications.

2. Literature Review and Theoretical Foundation

2.1. Open Government Data Portals

OGDPs represent a key form of public data infrastructure established and maintained by local governments. Open government data (OGD) refers to machine-readable public data released by public institutions and made accessible through OGDPs for unrestricted use by citizens, firms, and other organizations. Originating from mid-20th-century freedom-of-information initiatives in the United States, the OGD movement has since expanded globally, with adoption in more than 70 national jurisdictions (Lnenicka et al., 2024). OGD encompasses diverse data types, including healthcare and statistical data, geographic information, development plans, environmental records, and materials released by public authorities (Attard et al., 2015; Wirtz et al., 2022).

Governments promote these datasets for two broad reasons: (i) the intended benefits concern transparency and accountability, as public access is expected to strengthen trust in state institutions (e.g., Attard et al., 2015; Jetzek et al., 2019); and (ii) the unintended benefits concern economic value creation, as OGD can improve products and services, spark new market niches, and more broadly stimulate innovation (e.g., Janssen et al., 2012; Nagaraj, 2022). Empirical findings on the innovation benefits, however, remain thin and inconsistent. On the positive side, Beraja et al. (2023) show that government-provided datasets in facial recognition AI markets often exceed private-sector collections in scale, enabling firms to improve innovation performance. Nagaraj (2022) similarly finds that satellite mapping data helps new entrants reduce investment uncertainty and discover new markets. Hughes-Cromwick and Coronado (2019) further document how open government data supports more informed business decisions across multiple sectors. On the negative side, Zuiderwijk and Janssen (2014) argue that universal access enables competitive imitation, compressing returns and discouraging data-driven investment. Janssen et al. (2012) note that low data quality and institutional frictions limit the impact on innovation, and despite the abundance of OGD datasets, only a small proportion are actively used (Zuiderwijk and Reuver, 2021).

While these studies span multiple sectors and use cases, healthcare remains largely unexplored, despite being a setting where data sensitivity, regulatory complexity, and AI adoption pressures jointly shape AI investment decisions. Whether and how OGDPs influence local healthcare AI investment thus remains an open question, and one with growing practical stakes given the rapid expansion of public spending on such portals alongside the mixed evidence on their innovative value documented above.

2.2. Healthcare AI Investment

Healthcare AI is increasingly framed as a strategic lever for hospitals seeking to improve clinical quality while containing costs (Coiera and Liu, 2022). Current deployments concentrate in two complementary domains. Diagnostic AI systems (e.g., image interpretation networks, clinical decision-support modules, and surgical robots) have demonstrated significant gains in diagnostic accuracy, speed, and procedural

precision (Abdel-Karim et al., 2023; Beane, 2023). Operational AI systems apply predictive analytics and natural language processing to automate scheduling, triage, and patient communication, thereby reducing administrative burdens and clinician burnout (Reis et al., 2020). At the population level, widely used algorithms predict future healthcare needs and expenditures, enabling more efficient resource allocation and system-wide cost reductions (Obermeyer et al., 2019).

Hospitals nevertheless face well-documented barriers to AI investment (Baird et al., 2020). Limited trust often provokes resistance from both physicians and patients (Abdel-Karim et al., 2023; Dai and Tayur, 2022). For physicians, cognitive and affective concerns, including liability, control, professional status, and patient safety, drive resistance to AI systems (Reis et al., 2020; Liang and Xue, 2022). Among patients, resistance stems from fears that AI-enabled care may overlook individual characteristics and circumstances (Longoni et al., 2019), as well as from surveillance and delegation anxieties (Park et al., 2022), with skepticism amplified by documented cases of racial bias in model outputs (Obermeyer et al., 2019). Beyond these trust-related frictions, stringent privacy and ethical regimes restrict access to granular clinical records, and many hospitals lack the capacity to curate high-quality datasets required by AI vendors or internal teams (Ellahham et al., 2020; Sun and Medaglia, 2019). Prior work has examined these barriers primarily from an intra-organizational perspective, but this micro focus leaves underexplored the broader institutional environment that governs data flows, shapes investment incentives, and ultimately conditions the feasibility of healthcare AI projects.

Yet hospitals do not innovate in a vacuum. Recent research on healthcare AI ecosystems shows that investment decisions unfold within a web of interdependent actors (Leone et al., 2021; Sun and Medaglia, 2019): *technology vendors* that develop AI systems and generate revenue, *hospitals* that both supply data and deploy AI systems, and *governments* that regulate, reimburse, and manage public datasets. Because healthcare AI development depends on access to high-quality data and compliance with regulatory frameworks, hospital investment is shaped not only by internal organizational capability but also by the external policy environment (Sun and Medaglia, 2019; Dai and Tayur, 2022). Yet IS research has paid limited attention to how specific government actions shape organizational AI investment in this sector. While recent studies examine firm-level AI investment more broadly (e.g., Han et al., 2025; Li et al., 2025), how government-led institutional arrangements shape AI investment in regulated sectors such as healthcare remains underexplored. One such arrangement is the release of health and administrative datasets via OGDPs. Whether such initiatives foster or hinder hospital AI investment remains an open question, warranting both theoretical and empirical investigation.

2.3. Open Government Data and Healthcare AI Investment

Building on the literature above, we develop two competing predictions to understand how OGDP may influence local healthcare AI investment.

OGDP may *promote* healthcare AI investment through two channels. The first draws on the economics of data and AI innovation, which emphasizes data as a core complementary input to AI development (Agrawal et al., 2019; Jones and Tonetti, 2020). In healthcare specifically, AI vendors face binding data constraints: hospitals are reluctant to share patient records due to privacy regulations and competitive concerns, and the data accessible through private channels are typically fragmented, costly to acquire, and difficult to assemble at scale (Beraja et al., 2023). OGDP alleviates these constraints by mandating the disclosure of healthcare and administrative records in machine-readable forms at low or no cost. By expanding affordable access to large-scale and diverse datasets, OGDP increases both the feasibility and the expected returns of healthcare AI R&D (Nagaraj, 2022). This encourages greater vendor engagement in developing and refining AI products, which in turn provides hospitals with a richer portfolio of deployable solutions, ultimately increasing hospitals' willingness to invest.

The second channel draws on the signaling theory and economics of uncertainty and investment. Government actions convey policy preferences and future regulatory intentions, and organizations are particularly attentive to such signals when investing under uncertainty (Connelly et al., 2011; Duran et al., 2019; Estrin et al., 2016; Vasudeva et al., 2018; Lemoine, 2025). Policy uncertainty discourages organizational investment, whereas reductions in uncertainty stimulate it, especially in capital-intensive and irreversible technologies (Gulen and Ion, 2016; Sun and Medaglia, 2019). This signaling role is particularly salient in healthcare AI, where strict privacy regulations, evolving ethical norms, and a documented gap between rapid AI advancement and lagging regulatory frameworks create substantial uncertainty (Ellahham et al., 2020; Sun and Medaglia, 2019). Signaling theory suggests that for a government action to function as a credible signal, it must be costly to send and observable to receivers (Connelly et al., 2011; Spence, 1973). OGDP launches satisfy both conditions: (i) portal construction and maintenance entail substantial fiscal and inter-agency coordination costs, making them difficult to mimic without genuine commitment; and (ii) portal launches are publicly announced and timestamped. Under these conditions, government investment in public data infrastructure can be interpreted as endorsement of data use and AI deployment, indicating that data sharing is both permissible and institutionally supported. Hospitals may thus read OGDP launches as signaling a more supportive regulatory environment, reducing perceived policy risk and encouraging AI investment. Hence, we propose the following hypothesis.

H1a: *The launch of OGDP is positively associated with local healthcare AI investment.*

A contrasting logic suggests that OGDP may *hinder* healthcare AI investment. First, most OGDPs are not designed for healthcare AI applications and typically lack dedicated sections for diagnostic or operational use cases (Silva and Silva, 2025). The datasets they release are generally organized around administrative reporting categories rather than the technical specifications required for AI model training. As a result, only a limited share of available datasets can be effectively utilized by healthcare AI vendors,

constraining OGDP's potential to foster innovation (Zuiderwijk and Reuver, 2021). Second, low data quality and institutional frictions undermine the innovative impact of public data (Janssen et al., 2012). Healthcare AI models rely on high-value, well-curated datasets, yet publicly released data are often fragmented or insufficiently granular, and the widespread sharing of low-quality data can raise training and integration costs, discouraging investment (Acemoglu et al., 2022). Third, exploiting open data is complex and time-intensive, often requiring specialized skills (Zuiderwijk and Janssen, 2014), and the associated capability-building costs may further discourage AI vendors from engaging with OGDP. Beyond these supply-side frictions, open access can also dilute the strategic value of data as a competitive resource: when all vendors gain equal access to the same datasets, the data-based advantages that motivate proprietary investment are weakened (Zuiderwijk and Janssen, 2014). Hospitals anticipating these dynamics may likewise be reluctant to contribute data to OGDPs or partner with AI vendors. Collectively, these frictions raise data integration and coordination costs across the ecosystem, which potentially offsets the benefits of OGDP and dampens AI investment. Therefore, we propose a competing hypothesis.

H1b: *The launch of OGDP is negatively associated with local healthcare AI investment.*

Because OGDP can generate both positive and negative effects as predicted above, its net impact remains ambiguous. An in-depth empirical investigation is needed to determine which effect dominates and to elucidate the mechanisms through which OGDP shapes local healthcare AI investment.

3. Context, Data Sources, and Key Variables

3.1. Empirical Context

OGDPs in China. Following State Council directives in 2012, cities across China began establishing OGDPs to aggregate datasets across domains and improve transparency. Rollout has been rapid (Figure 1): by 2023, 134 cities had launched OGDPs, each organized around a nationally mandated catalog that standardizes metadata and access protocols. Because portals are developed under common national guidelines, their interface design and access functions are largely uniform, leaving limited variation in features such as search interfaces and download procedures. OGDPs serve both as centralized repositories and as distribution infrastructures, with unified data-sharing standards that lower acquisition and transaction costs for external users.

OGDPs in China share several design features relevant to this study: (i) *multisource coverage*: datasets are released by a wide range of public agencies (e.g., health, transportation, environmental monitoring, and social security), providing a broad foundation for data-driven applications; (ii) *raw data release*: files are provided in machine-readable forms (CSV, JSON, XML) with minimal reprocessing, requiring users to integrate and analyze the material themselves; and (iii) *tiered access*: core datasets are freely downloadable, while sensitive data (e.g., de-identified patient records) require application and are

granted at a nominal cost. Illustrations of dataset classification and detailed dataset information provided on the OGDPs are also presented in Figures A-1, A-2, and A-3 in Appendix A.

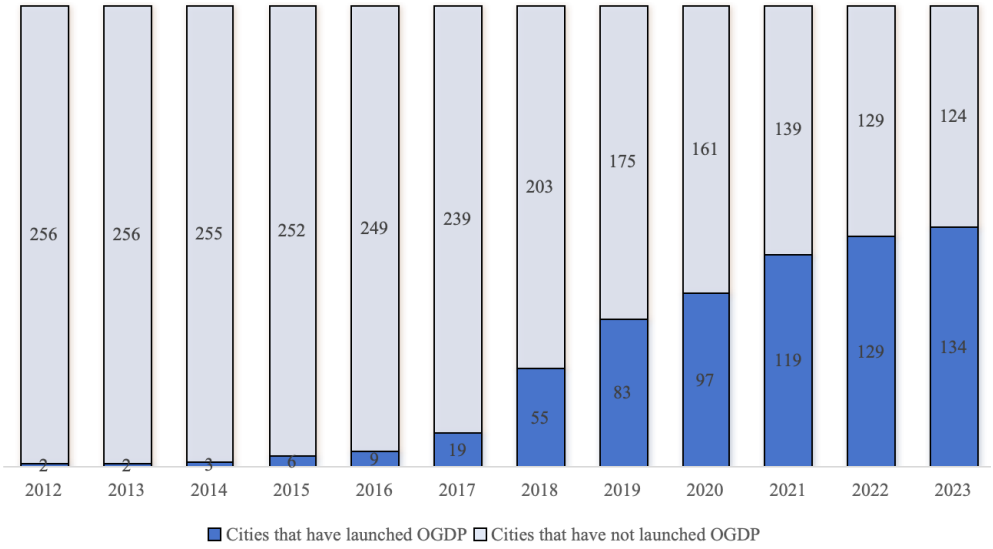


Figure 1. The Staggered Adoption of OGDPs Across Cities over the Period 2012–2023 in China

Several institutional features support a quasi-experimental reading of the OGDP rollout. (i) Chinese OGDPs originate from earlier public data opening policies rooted in transparency and accountability agendas rather than innovation policy, and these policies were introduced well before AI-related policy agendas emerged, i.e., when the potential of public data for AI was neither anticipated nor a policy objective. As such, OGDP launches were unlikely to be driven by local AI strategies or coordinated digitalization plans, and hospitals and AI vendors neither participated in portal design nor had information about rollout timing in advance. (ii) The timing of OGDP launch also does not systematically correlate with local economic development or innovative capacity. For example, despite greater science and technology resources, Nanjing launched its OGDP later than Yangzhou within the same province.[3] (iii) OGDPs operate primarily within city boundaries: released datasets are specific to the issuing municipality, and portal usage is overwhelmingly local (e.g., 98% of visits to the Hangzhou OGDP originate from within the same city, see Figure A-4 in Appendix A).

Healthcare Sector and AI Investment in China. Chinese public hospitals serve as the primary providers of outpatient, emergency, and inpatient care nationwide, yet face mounting pressure from population aging and rising chronic-disease prevalence, which continue to widen the gap between

[3] We cannot, however, fully rule out selection on unobserved city-level characteristics, such as local governments' latent orientation toward digital governance. We address this concern through matching estimators, which balance treated and control cities on observable covariates, and a generalized synthetic control specification, which accounts for unobserved time-varying confounders (Table 3).

healthcare supply and demand.[4] These structural pressures create strong and sustained incentives for hospitals to invest in AI systems that improve diagnostic efficiency and reduce frontline workload. Figure 2 illustrates this trend: the number of healthcare AI procurement contracts awarded by Chinese hospitals rose from 168 in 2012 to 10,655 in 2023, with the pace accelerating markedly after 2019.

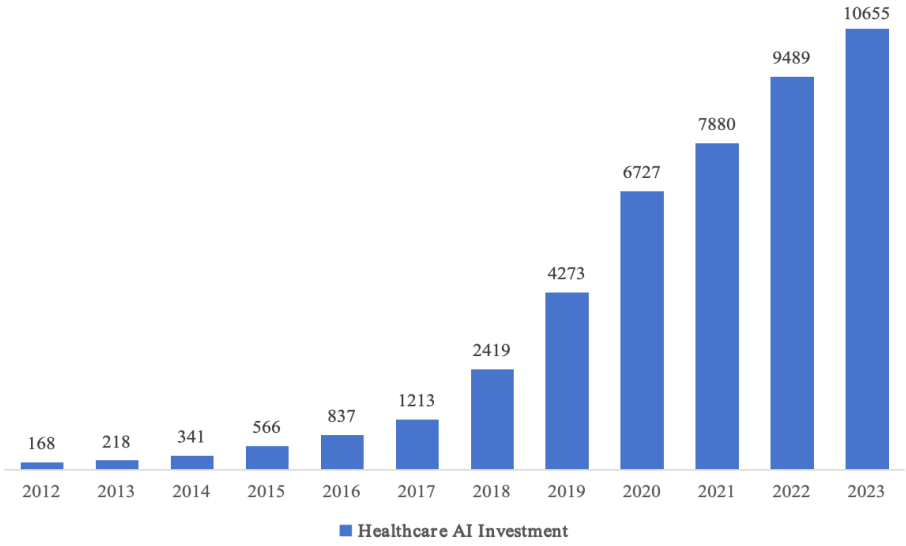


Figure 2. Development of Healthcare AI Investment over the Period 2012–2023 in China

3.2. Key Variables and Data Sources

Independent Variable. Our independent variable, $OGDP_{i,t}$, is a city-quarter indicator equal to one if city i has launched and is operating its OGDP as of year-quarter t, and zero otherwise. For example, Hangzhou launched its OGDP in January 2015; hence, the indicator is coded as 1 from the first quarter of 2015 onward and 0 in all preceding quarters. To identify accurate launch dates, we adopted a multi-step verification procedure: (i) we directly contacted the operators of each city-level OGDP to confirm the dates on which their portals went live; (ii) we then conducted systematic online searches using multiple search engines to supplement and cross-validate the information; and (iii) we cross-checked the resulting dates against the China Local Government Data Open Report (Cities) series published by the Digital and Mobile Governance Lab at Fudan University, treating our directly verified dates as authoritative in cases of discrepancy. To address concerns that OGDPs launched in the final month of a quarter may have insufficient time to generate observable effects within the same period, we construct an alternative measure that postpones the effective timing by one quarter in such cases (Table 3), and results remain consistent (Table B-6 in Appendix B).

[4] In 2022, medical and health institutions recorded 8.42 billion patient visits, with hospital physicians handling an average of 6.2 outpatient visits per day in addition to substantial inpatient workloads, yet the nationwide supply of medical personnel remains limited, with only 3.2 physicians and 3.7 nurses per 1,000 population, well below the OECD average.

Dependent Variable. Our dependent variable captures the extent of local hospitals' investment in healthcare AI systems. We measure this using awarded procurement contracts, which record executed purchases and thus reflect realized technology investment rather than procurement intentions, a measurement approach widely used in prior AI innovation research (e.g., Beraja et al., 2023). We proxy AI investment by the count of awarded procurement contracts rather than their monetary value. Because AI components are frequently bundled into broader healthcare procurements, a value-based measure would conflate AI spending with the full bundle price, leading to overstatement.[5] Nevertheless, we additionally construct a value-based measure as a complementary check (Table 3), and the results remain consistent (Column 2 in Table B-7, Appendix B).

We construct the dependent variables in three steps: (i) we acquired hospital procurement contract data for the period 2012–2023 from a proprietary database of hospital procurement records constructed and maintained by a leading commercial healthcare-industry data vendor (more details in Appendix A.3); (ii) we identified healthcare AI contracts using a structured keyword-based approach, constructing a list of 94 AI-related keywords along two complementary dimensions: a *technical* dimension and an *application-scenario* dimension. The technical dimension captures core AI technologies commonly embedded in healthcare AI systems, including terms such as "Generative AI," "Machine Learning," "Natural Language Processing," and "Visual Image Recognition." The application-scenario dimension reflects healthcare-specific contexts in which AI technologies are deployed, such as "Intelligent Diagnosis," "Virtual Doctor," and "Intelligent Pre-Consultation." This two-dimensional design identifies contracts that reference underlying AI technologies or specific healthcare application scenarios, reducing both false positives and false negatives in classification. The keyword list is compiled from authoritative sources, including the *Artificial Intelligence Index Report 2025*, the National Health Commission's guidelines on AI applications in healthcare, recent academic studies (e.g., Pham et al., 2024), and consultations with senior experts in the field of healthcare AI. To validate the keyword-based classification, two independent coders manually labeled a stratified random sample of 600 contracts across time periods and categories. The classification achieves satisfactory precision, recall, and inter-coder agreement, with stable performance across early and late sample periods (Appendix A.3.2).

(iii) We then classified the 44,786 identified contracts into diagnostic (40.0%) and operational (60.0%) categories based on their functional descriptions and aggregated the city-quarter counts to construct three dependent variables: *Diagnostic Healthcare AI Investment*$_{i,t}$, *Operational Healthcare AI Investment*$_{i,t}$, and *Total Healthcare AI Investment*$_{i,t}$ (details on healthcare AI procurement contracts, the

[5] For example, one contract in our sample awards Guangzhou Medical University Affiliated Cancer Hospital a da Vinci IS4000 surgical robot at ¥24.0 million, where the AI functions account for only a fraction of the bundled price (see Appendix A.3 for details).

list of keywords, and the construction procedure are provided in Figure A-6, Figure A-7, Figure A-8, Figure A-9, and Table A-1 in Appendix A). Because keyword searches may include contracts in which AI-related terms are mentioned only incidentally rather than as a substantive component, we re-estimate our models excluding such contracts as a validation check (Table 3); results are robust.[6]

Table 1. Main Variables, Measurement, and Summary Statistics

Variable	Measurement	Mean	Std. Dev	Min	Max
		(1)	(2)	(3)	(4)
Dependent Variables					
Healthcare AI Investment	Number of healthcare-AI procurement contracts signed by hospitals in a city and quarter	3.616	9.743	0	176
Diagnostic Healthcare AI Investment	Number of diagnostic healthcare-AI procurement contracts signed by hospitals in a city and quarter	1.445	4.469	0	102
Operational Healthcare AI Investment	Number of operational healthcare-AI procurement contracts signed by hospitals in a city and quarter	2.171	5.660	0	110
Independent Variables					
OGDP	Equals 1 if city i has an OGDP in quarter t, otherwise 0	0.192	0.394	0	1
Control Variables					
Registered population	Registered population of city i, in 10,000 persons (year-level)	472.326	332.584	19.800	3416.000
GDP per Capita	GDP per capita of city i, in yuan (year-level)	59454.375	34459.442	8157	206278
# Digital patents	Number of digital patents granted in city i (year-level)	1733.156	5778.092	0	94566
Government R&D spending	Government R&D investment of city i, in 10,000 yuan (year-level)	156687.301	461906.544	519	5549817
# Higher-education institutions	Number of regular institutions of higher education in city i (year-level)	9.418	15.571	1	93
# Doctors	Number of doctors in city i (year-level)	12418.323	12504.859	778	133685
# Hospitals	Number of hospitals and health centers in city i (year-level)	167.651	168.430	8	3052

Notes: Continuous dependent variables and time-varying controls are natural log-transformed in regression analyses.

Control Variables. To account for time-varying city characteristics that might affect both OGDP rollout and local healthcare AI investment, we include control variables drawn primarily from the China Statistical Yearbook: *resident population* (market size), since larger cities face greater healthcare demand and are more likely to invest in healthcare AI; *per-capita GDP* (economic capacity), since economically

[6] We construct a conservative measure that restricts the sample to contracts that explicitly use AI terminology, reducing it from 44,786 to 39,340 contracts. The OGDP coefficient remains positive and statistically significant (Column 1 of Table B-7 in Appendix B).

stronger cities are more likely to adopt OGDPs and invest in AI; *digital patents granted* (digital infrastructure), since digitally advanced cities more readily build OGDPs and adopt AI technologies; *number of regular higher-education institutions* (human capital), which supports both data infrastructure development and AI adoption capacity; *government R&D spending* (technological investment), which may support OGDP construction and local AI activity; and *number of doctors* and *hospitals* (healthcare resources), as greater medical capacity is associated with a higher investment in clinical AI systems.

Our final balanced panel dataset covers 258 prefecture-level cities observed from 2012 through 2023. To preserve comparability, we exclude non-city administrative units (e.g., autonomous prefectures, leagues) and cities that changed administrative status during the study period.[7] Key variables, their measurements, and summary statistics are reported in Table 1.

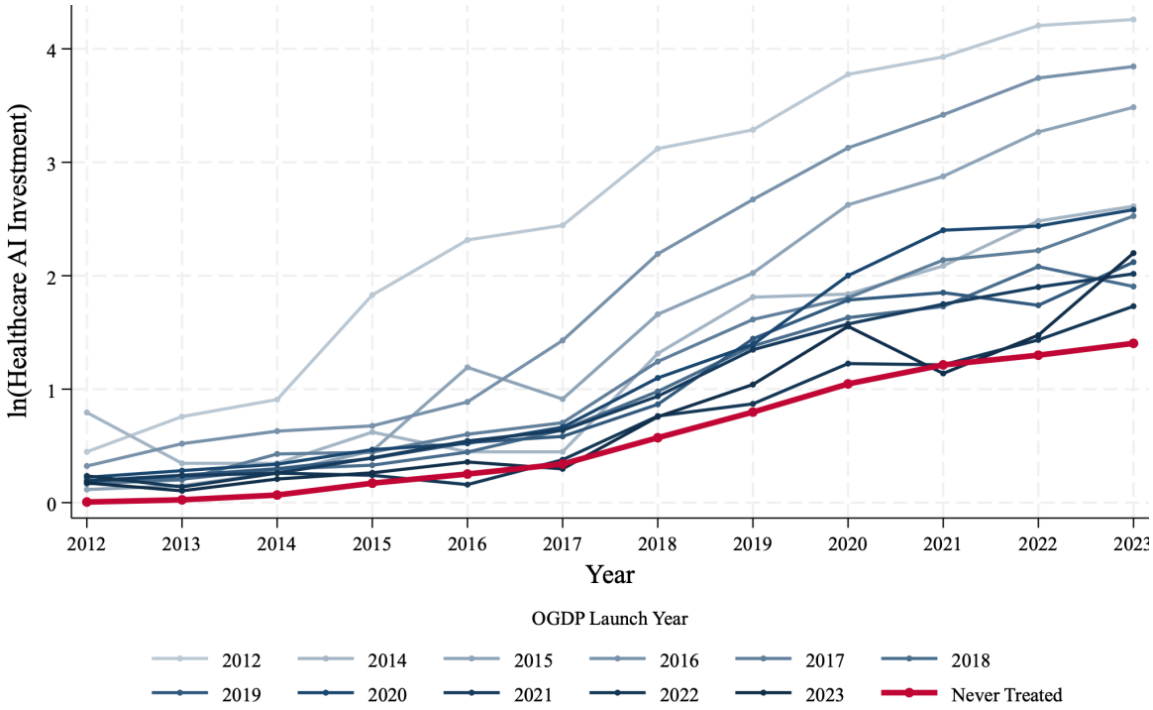


Figure 3. Non-parametric Comparison of Treated and Never-treated Cities

Notes: This figure plots the average log of healthcare AI investment from 2012 to 2023, separately for never-treated cities (red line) and for treated cities grouped by their OGDP launch year (blue lines). The vertical dashed lines mark the OGDP launch timing for each cohort. The figure provides descriptive evidence on the differential trends in healthcare AI investment between cities that did and did not launch OGDPs.

3.3. Descriptive Evidence

Before turning to the formal econometric analysis, we examine non-parametric comparisons of hospitals' AI investment across cities with and without OGDPs, as well as within-city comparisons before and after

[7] To assess whether results are driven by the structural characteristics of China's most developed metropolises, we re-estimated all models excluding Beijing, Shanghai, Tianjin, Chongqing, Guangzhou and Shenzhen. The OGDP coefficient remains stable at 0.290 ($p < 0.01$), virtually identical to the baseline estimate, confirming that our findings are not an artifact of outlier city influence (Table B-5 in Appendix B).

OGDP launches. As Figure 3 shows, treated cities exhibit higher levels of healthcare AI investment than never-treated cities and show further increases relative to their own pre-launch baselines. These patterns are suggestive but raise two identification concerns: (i) OGDP adoption may be endogenous to unobserved city characteristics, and (ii) time-varying unobserved shocks may simultaneously influence both OGDP launches and healthcare AI investment. We address these concerns below through a staggered difference-in-differences design with comprehensive controls and a battery of robustness checks.

4. Empirical Strategy and Results

4.1. Identification Strategy

The phased rollout of city-level OGDPs across China during 2012–2023 provides a quasi-experimental setting. We exploit this staggered adoption to compare changes in hospital AI investment following OGDP launch to contemporaneous changes in not-yet-treated cities. Specifically, we estimate a two-way fixed effects difference-in-differences (TWFE-DiD) model, implemented as a linear OLS regression with covariates and city- and year-quarter fixed effects. The baseline specification is as follows:

$$Healthcare\ AI\ Investment_{i,t} = \beta OGD_{i,t} + Controls_{i,t}^{'}\eta + \mu_i + \lambda_t + \varepsilon_{i,t}, \quad \text{(Eq. 1)}$$

where μ_i captures city fixed effects to account for time-invariant city-level heterogeneity (e.g., geographic location, historical industrial structure, baseline infrastructure). λ_t denotes year-quarter fixed effects that absorb common macroeconomic, technological, and seasonal shocks. *Controls* is a vector of time-varying city-level covariates. β is the estimate of interest to capture the impact of OGDPs on local healthcare AI investment. We apply the log(1+x) transformation to the dependent variables, which accommodates zero values in city-quarter observations with no AI procurement and allows coefficients to be interpreted as percentage changes.[8] Standard errors are clustered at the city level to account for serial correlation and within-city variations over time.[9]

The TWFE specification addresses several threats to causal identification. First, persistent differences across cities, such as baseline digital infrastructure, historical industrial structure, or geographic characteristics, may be correlated with both OGDP launch and healthcare AI investment. City fixed

[8] We assess robustness to functional-form choices in two ways. First, since the log(1+x) transformation can introduce bias when outcomes include many small or zero values (Bellemare and Wichman, 2020), we re-estimate all models using the inverse hyperbolic sine transformation, which remains defined at zero without arbitrary adjustment; results are consistent (Table B-8 in Appendix B). Second, we re-estimate the baseline model using Poisson Pseudo-Maximum Likelihood (PPML), which is robust to distributional misspecification and addresses concerns regarding log-transformation of count variables (Chen and Roth, 2024); results remain consistent (Table B-9).

[9] To assess sensitivity to the choice of clustering level, we re-estimate all models with standard errors clustered at the province level, allowing for arbitrary correlation of residuals across cities within the same province. Results remain consistent (Table B-10).

effects absorb these time-invariant differences, ensuring that identification is driven by within-city variation over time. Second, nationwide shocks, including macroeconomic fluctuations, national AI policies, and broad technological trends, may affect all cities simultaneously and be incorrectly attributed to OGDP effects. Time fixed effects address this concern by absorbing factors that vary over time but are common across cities.

Third, time-varying city-specific shocks may pose a more subtle threat to identification. Our design mitigates this concern by exploiting staggered adoption and comparing treated cities to contemporaneous not-yet-treated cities, such that identification is driven by differential changes around the timing of OGDP launch. Moreover, because identification relies on variation within relatively narrow time windows, confounders that evolve gradually are unlikely to generate sharp discontinuities that align with the OGDP launch. To further address this concern, we estimate a more demanding specification that incorporates city-specific linear and quadratic time trends (Table 3), which flexibly account for smoothly evolving city-level shocks. This specification imposes a stringent identification requirement, as treatment effects are identified only from deviations relative to underlying city-specific trends. The OGDP coefficient remains consistent in both magnitude and statistical significance across these specifications, suggesting that our results are more likely to reflect the impact of OGDP launch than pre-existing city-level trajectories (Column 3 of Table B-1 in Appendix B).

To further guard against time-varying unobserved confounders that may differentially affect treated and control cities, we employ the generalized synthetic control (GSC) method (Xu, 2017), which constructs a synthetic counterfactual for each treated city by reweighting control cities to match pre-treatment covariate profiles and outcome trajectories. The GSC estimate is positive and statistically significant ($p < 0.01$), and somewhat larger in magnitude than the TWFE estimates, as the average treatment effect on the treated in GSC is identified exclusively from treated cities rather than the full sample. This estimate is nonetheless consistent in sign and significance with the baseline results, providing additional assurance that our findings are not driven by unobserved time-varying confounders (Table B-12 and Figure B-5).

Although the TWFE specification has been widely adopted for staggered-adoption program evaluations, it delivers consistent estimates only under the assumption of homogeneous treatment effects (de Chaisemartin and d'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). Under heterogeneous treatment effects, the TWFE coefficient is a weighted average of all possible 2 × 2 DiD comparisons in which some weights can be negative, potentially yielding estimates that do not correspond to any convex combination of true treatment effects (Goodman-Bacon, 2021; Baker et al., 2022). Applying the Bacon decomposition, we find that the identifying weight is concentrated overwhelmingly in clean comparisons between treated and never-treated cities, while "forbidden

comparisons" receive negligible weight. Point estimates across comparison types are directionally consistent and quantitatively similar, suggesting that heterogeneous treatment timing does not meaningfully distort the TWFE estimate (Figure B-3).

Although these diagnostics indicate minimal distortion in our setting, we further verify our results using estimators robust to heterogeneous treatment effects. We report CS-DiD estimates (Callaway and Sant'Anna, 2021) alongside the TWFE results in Table 2, and examine pre-treatment trends and post-treatment dynamics using the heterogeneity-robust estimators of de Chaisemartin and d'Haultfoeuille (2020), Cengiz et al. (2019), and Sun and Abraham (2021) in the event study analysis (Figure 4).

4.2. Main Results

Table 2 reports the TWFE-DiD and CS-DiD estimates, indicating that the launch of OGDPs is associated with a statistically significant increase in local healthcare AI investment ($p < 0.01$), consistent with H1a. Controlling for time-varying city-level covariates and fixed effects, total healthcare AI investment increases by an average of 23.7% (= $e^{0.213}$-1, Column 1), diagnostic healthcare AI investment increases by 24.5% (= $e^{0.219}$-1, Column 2), and operational AI investment increases by 21.3% (= $e^{0.193}$-1, Column 3). A Z-test cannot reject the equality of coefficients when diagnostic and operational AI investment are the dependent variables ($p = 0.79$), suggesting that the launch of OGDPs stimulates both categories of AI investment to a comparable extent.[10]

Table 2. Average Effects of OGDP Launch on Local Healthcare AI Investment

DV:	Total Healthcare AI Investment	Diagnostic Healthcare AI Investment	Operational Healthcare AI Investment
	(1)	(2)	(3)
TWFE-DiD	0.297*** (0.053)	0.290*** (0.048)	0.250*** (0.049)
CS-DiD	0.213*** (0.079)	0.219*** (0.056)	0.193** (0.077)
Controls	Yes	Yes	Yes
City FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
# Observations	12,384	12,384	12,384

Notes: The table reports TWFE-DiD and CS-DiD estimates of the effect of OGDP launches on local healthcare AI investment. All columns include time-varying city-level covariates and city and year-quarter fixed effects. We report both estimators to account for potential heterogeneity in treatment effects across cities and over time. Robust standard errors clustered at the city level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

[10] To address concerns about comparability between treated and control cities, we re-estimate the baseline model using propensity score matching (PSM), coarsened exact matching (CEM), and weighted least squares; the OGDP coefficient remains positive and statistically significant across all specifications (PSM: 0.303, $p < 0.01$; CEM: 0.249, $p < 0.01$; GDP-weighted: 0.297, $p < 0.01$; Tables B-2, B-3, and B-4 in Appendix B). To check spurious correlations, we conduct a placebo test by randomly assigning pseudo treatment status and launch timing to cities across 500 replications; the resulting distribution of placebo coefficients is centered around zero, and our baseline estimate falls far in the right tail, confirming that results are unlikely to reflect chance findings (Figure B-4).

The validity of DiD estimates rests on the parallel-trends assumption: treated and control cities would have followed similar trends in healthcare AI investment in the absence of treatment. We assess this assumption with an event study specification (Autor, 2003) that replaces the single $OGDP_{i,t}$ in Equation (1) with a set of leads and lags, $\sum_j \tau_{i,t} PreOGDP_{i,t}(j)$ and $\sum_k \omega_{i,t} PostOGDP_{i,t}(k)$, defined relative to the quarter immediately preceding the OGDP launch. Figure 4 overlays dynamic effect estimates from five estimators: TWFE-OLS, CS-DiD, and the heterogeneity-robust estimators of de Chaisemartin and d'Haultfoeuille (2020), Cengiz et al. (2019), and Sun and Abraham (2021). Pre-treatment coefficients fluctuate around zero and are jointly insignificant across all estimators, supporting the parallel-trends assumption. Post-treatment estimates turn positive, remain largely significant, and show no long-run attenuation, indicating a sustained increase in healthcare AI investment following OGDP launch.

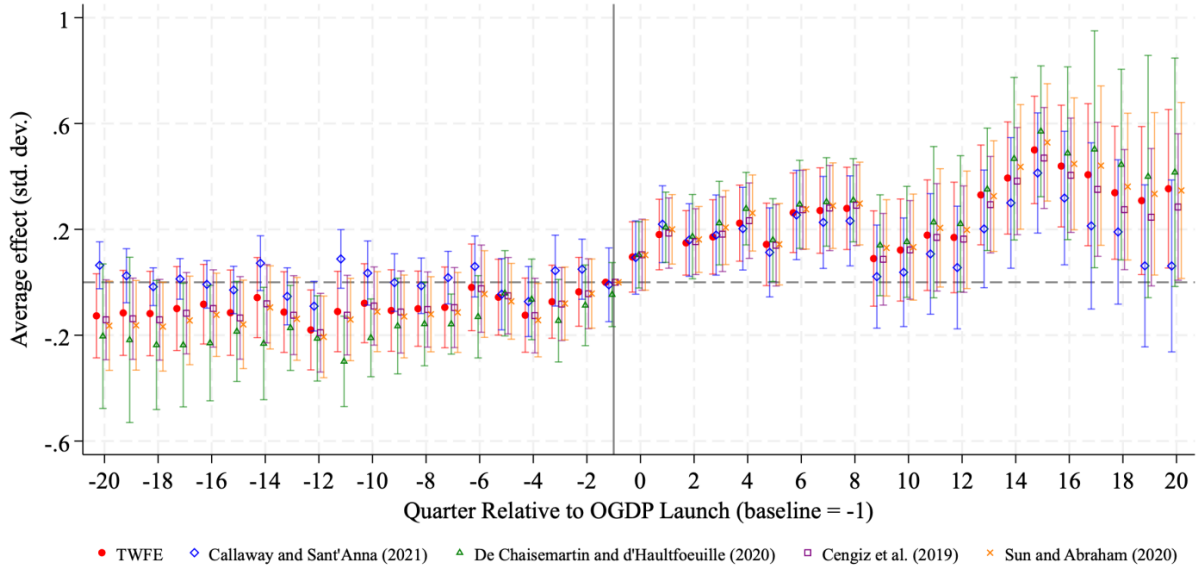


Figure 4. Event Study Estimates on the Dynamic Effects of OGDP Launch

Notes: This figure overlays the event-study plots constructed using five different estimators: a dynamic version of the TWFE model, estimated using OLS (in red with circle markers); Callaway and Sant'Anna (2021) (in blue with diamond markers); De Chaisemartin and d'Haultfoeuille (2020) (in cyan with triangle markers); Cengiz et al. (2019) (in purple with square markers); and Sun and Abraham (2021) (in orange with cross markers). Bars represent 95% confidence intervals. Standard errors are clustered at the city level.

4.3. Robustness Checks

Table 3 reports a battery of robustness checks. We assess sensitivity to (i) the inclusion of alternative covariate sets, including interactions of covariates with year-quarter fixed effects, city-specific intercepts and slopes, and additional controls for medical insurance enrollment; (ii) the sampling strategy, using propensity score matching, coarsened exact matching, GDP- and hospital-weighted regressions, exclusion of outlier municipalities (Beijing, Shanghai, Tianjin, Chongqing, Guangzhou, and Shenzhen), and a Bacon decomposition; (iii) variable measurement, including alternative AI keyword filters that exclude incidental

Table 3. Summary of Empirical Challenges and Robustness Checks

Empirical Challenges	Robustness Checks	Results	Location
1. Are the estimates sensitive to including different sets of covariates?	(1) Drop all covariates and keep the city and quarter fixed effects (2) Maintain all covariates in the baseline model and add the interactions between all covariates and the year-quarter fixed-effects (3) Maintain all covariates in the baseline model and absorb heterogeneous city-specific intercepts and slopes (4) Additionally control for the number of participants in the basic medical insurance scheme for employees	Results remain consistent	Table B-1, Figure B-1
2. Are the estimates sensitive to the sampling strategy?	(1) Perform Propensity Score Matching (PSM)	Results remain consistent	Figure B-2, Table B-2
	(2) Perform Coarsened Exact Matching (CEM)		Table B-3
	(3) Run regressions weighted by city-level per-capita GDP and by the number of hospitals, respectively		Table B-4
	(4) Remove outlier cities (i.e., Beijing, Shanghai, Guangzhou, Shenzhen, Tianjin, and Chongqing) from the sample		Table B-5
	(5) Conduct a Bacon decomposition prior to the main regression analysis		Figure B-3
3. Are the estimates sensitive to different measurements for independent and dependent variables?	(1) Postpone OGDP timing by one quarter if the release month falls in the last month of each quarter	Results remain consistent	Table B-6
	(2) Replace the dependent variable with the log number of explicitly AI-related healthcare procurement contracts, excluding incidental mentions of AI		Table B-7
	(3) Replace the dependent variable with the log number of the value of healthcare procurement contracts		Table B-7
	(4) Apply the inverse hyperbolic sine transformation for cross-validation		Table B-8
4. The distribution of the DV caters to a count data model	Replace linear specification with Poisson Pseudo-Maximum Likelihood (PPML) specification	Results remain consistent	Table B-9
5. Are standard errors different if clustered at a different level?	Cluster standard errors at the province (rather than city level)	Results remain consistent	Table B-10
6. Was COVID-19 shock a confounder?	Exclude the impact of the COVID-19 pandemic by restricting the sample to the period 2012-2019	Results remain consistent	Table B-11
7. Spurious correlation?	Conduct a placebo test by applying randomly assigned pseudo treatments to cities and quarters for 500 times	No evidence for spurious correlation	Figure B-4
8. Unaccounted time-varying confounders	Employ the generalized synthetic control (GSC) method	Results remain consistent	Table B-12, Figure B-5
9. Is the TWFE model with a staggered-adoption design consistent?	Use heterogeneity-robust staggered DiD and event study estimators	Results remain consistent	Table 2, Figure 4

Note: The tables and figures for these robustness checks are provided in the appendix.

mentions, a value-based dependent variable, postponed treatment timing for portals launched in the final month of a quarter, and the inverse hyperbolic sine transformation; (iv) the count nature of the outcome,

using PPML estimation; (v) standard error clustering at the province rather than the city level; and (vi) the COVID-19 period, by restricting the sample to 2012–2019. We further (vii) conduct a placebo test that randomly assigns pseudo-launch quarters across 500 replications; (viii) apply the generalized synthetic control method (Xu, 2017) to address unobserved time-varying confounders; and (ix) re-estimate the model using heterogeneity-robust staggered-DiD estimators (Callaway and Sant'Anna, 2021; de Chaisemartin and d'Haultfoeuille, 2020; Sun and Abraham, 2021). All results remain consistent.

4.4. Alternative Explanations

Beyond the robustness checks, three potential confounders warrant particular attention: bundled policy shocks, pre-existing digital capacity, and local fiscal resources.

The first concern is that OGDP launches may coincide with (and the estimated effects hence reflect) broader policy initiatives targeting AI or data-driven innovation. The 2012–2023 period saw substantial activity on both fronts in China, including national AI strategies and data-governance reforms whose introduction varied across provinces and time. To isolate the OGDP effect, we exploit cross-province and temporal variation in major AI and data policies and include indicators for their introduction in Equation (1). The OGDP coefficient remains consistent, with overlapping confidence intervals across specifications (Table C-2 in Appendix C). While this exercise cannot exhaust the universe of unobserved local initiatives, it provides evidence that the main result is not absorbed by the most salient concurrent policies.

The second concern is that the OGDP rollout may track pre-existing local digital capacity, with such capacity (rather than the OGDP itself) driving healthcare AI investment. If this were the case, controlling for digitalization trends should attenuate the OGDP coefficient, and effects should concentrate in highly digitalized cities. We construct time-varying measures of digital capacity using digital patenting and telecommunications revenue, and conduct subsample analyses by baseline digital intensity. The OGDP coefficient remains stable and significant, and we find no evidence of stronger effects in highly digitalized cities (Table C-3 in Appendix C). The latter pattern is itself informative: a digital-capacity story would predict precisely the opposite, suggesting that OGDP effects operate through channels beyond pre-existing digital infrastructure.

The third concern is that local fiscal capacity, rather than OGDP launches, drives the observed increase, since fiscally stronger cities are better positioned to launch digital public infrastructure and subsidize innovation (Besley and Persson, 2009; Howell, 2017). We construct time-varying measures of local public budget expenditure and the fiscal share allocated to science and technology, and conduct subsample analyses by baseline fiscal strength. The OGDP coefficient again remains stable, with no evidence of stronger effects in fiscally stronger cities (Table C-4 in Appendix C). As with digital capacity, the absence of heterogeneity is informative: a fiscal-resource explanation would predict concentrated effects in wealthier cities, which we do not observe.

These tests do not eliminate every possible confounder. In particular, unobserved local government characteristics correlated with both OGDP launch and AI investment cannot be fully ruled out; we partially address this through the generalized synthetic control specification reported in Table 3, which accounts for unobserved time-varying confounders by constructing synthetic counterfactuals for each treated city. Taken together, the evidence consistently points to an OGDP-specific effect rather than to spurious correlations with concurrent policy, digital, or fiscal conditions.

4.5. Economic Significance

To assess economic significance, we estimate effects across two complementary dimensions, *technical attributes* and *application scenarios*, and benchmark them against public expenditures on OGDPs.

We first group healthcare AI technologies into four categories based on their technical functions: infrastructure, models & algorithms, perception & cognition, and execution & interaction. The estimated coefficient on overall AI investment is 0.895 (Table D-1 in Appendix D), indicating a 144.7% increase relative to the pre-treatment baseline, corresponding to a cumulative incremental procurement of approximately ¥5.67 billion (~$801 million)[11] across treated municipalities. The increase is concentrated in data-intensive analytical systems: models & algorithms and perception & cognition rise substantially (125.0% and 100.2%, respectively), while execution & interaction rises more modestly (22.8%), plausibly reflecting greater organizational adjustments required for clinical deployment (Table D-1).

We then group the AI-related contracts by application scenarios: clinical workflow & services, hospital operations & management, and extended services. The estimated coefficient on overall application-layer AI procurement is 0.540, corresponding to a 71.6% increase (~¥5.1 billion, ~$720 million) in cumulative additional procurement (Table D-2), concentrated in clinical workflow (50.8%) and operations & management (58.6%), with no significant response in extended services. This pattern suggests that OGDP-induced investment flows primarily toward hospitals' core operational and clinical functions rather than peripheral or externally oriented services.

Because these two dimensions represent alternative analytical lenses applied to the same underlying investment, their effects are evaluated separately rather than additively. Benchmarking these increases against the estimated cumulative cost of OGDP construction and operation across treated municipalities (~¥1.34 billion, ~$189 million; see Appendix D.3 in Appendix D) yields leverage ratios of 4.23 (technical attributes) and 3.80 (application scenarios). Notably, we interpret these leverage ratios as suggestive rather than definitive estimates of fiscal leverage: they do not account for OGDPs' broader social returns or full

[11] All RMB-to-USD conversions in this section use the 2023 annual average exchange rate of 1 USD = 7.0774 CNY, sourced from the Federal Reserve Board's H.10 statistical release.

operational costs, nor do they capture spillovers beyond healthcare AI. With these caveats, the magnitudes are consistent with the view that public data infrastructure leverages substantial healthcare AI investment.

5. Exploring Underlying Mechanisms

We examine two candidate channels through which OGDPs may shape healthcare AI investment: data provision and governmental signaling. As developed in Section 2.3, these channels stem from the same intervention but operate through theoretically distinct pathways and target different decision constraints — data constraints versus regulatory uncertainty. Data provision is a material pathway, operating by expanding the tangible inputs available for AI development. Governmental signaling is a cognitive pathway, operating by changing how hospitals and vendors interpret the regulatory environment, independent of the material data flowing through the portal. The two channels generate distinct empirical predictions: data provision implies stronger OGDP effects where portals release more or more actively used healthcare data, whereas signaling implies stronger effects where OGDP-related signals are costlier to send and more observable to ecosystem actors. We assess each in turn and then compare their relative explanatory power in a joint specification.

5.1. Data Provision

The data-provision channel predicts that OGDPs stimulate healthcare AI investment by alleviating data constraints on both sides of the market. On the supply side, developing healthcare AI requires substantial investment in data acquisition, annotation, storage, and model training, yet vendors typically face binding constraints: hospitals encounter regulatory barriers to sharing patient data, limiting the supply of training data, while vendors struggle to obtain datasets of sufficient scale and reliability through private channels. OGDPs alleviate these constraints by making administrative and health-related records available at low or no cost, enabling vendors to train more accurate and generalizable models and thus develop higher-performance healthcare AI products. On the demand side, a richer and higher-quality portfolio of deployable AI products directly increases hospitals' willingness to invest, as better solutions lower adoption risk and strengthen the expected returns from procurement. If this channel drives the OGDP effect, treatment effects should be stronger where portals release a larger volume of healthcare data or where such data are more actively used.

To test this prediction, we construct two proxies, manually collected from each city's portal: (i) *total volume of healthcare data*, the number of healthcare data entries uploaded per year-quarter, which captures the scale of data supply; and (ii) *downloads per healthcare data entry*, total healthcare-data

downloads per quarter divided by total healthcare-data volume, which captures active utilization.[12] For each proxy, we split cities at the sample-period median and estimate Equation (1) separately on the resulting high and low subsamples (each paired with the never-treated control cities), with chi-square tests of the equality of OGDP coefficients across the two groups.

Figure 5 reports the results. The OGDP effect is positive and significant in both subsamples across both proxies. Although point estimates are directionally larger in cities with greater data volumes and higher downloads per entry, the inter-group differences are not statistically significant ($p > 0.10$ in both cases).[13] This pattern suggests that data provision operates as a baseline input on which the OGDP effect rests, rather than as a moderator that systematically amplifies treatment intensity across cities.

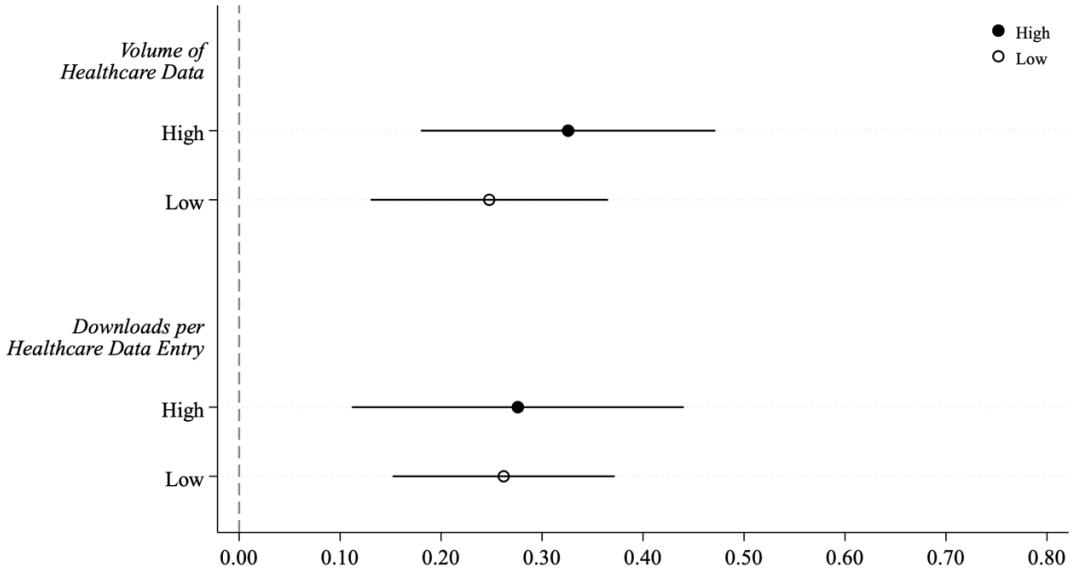


Figure 5. Subsample Analysis Results for the Data Provision Mechanism

Notes: This figure presents subsample tests of the data-provision mechanism. We construct two proxies for healthcare-data availability and use: (i) total volume of healthcare data (healthcare entries uploaded to a city's OGDP per quarter); and (ii) downloads per unit of healthcare data (total quarterly downloads divided by total healthcare data volume), capturing active use. All measures are manually collected from each city's OGDP. We split cities at the sample-period median and estimate Equation (1) separately on the high and low OGDP-treated subsamples (each paired with the never-treated control cities), then compare the OGDP effects across groups.

5.2. Governmental Signaling

The signaling channel predicts that OGDPs reduce regulatory uncertainty for both hospitals and vendors by conveying credible governmental commitment to data-driven innovation. This channel is particularly

[12] Downloads per entry are intended to capture which datasets users find valuable enough to retrieve repeatedly. Like any utilization proxy, it cannot perfectly distinguish active analytical use from repeated retrieval of poorly documented files; we therefore treat it as informative rather than definitive.

[13] We further re-estimate the mechanism using continuous interaction terms, interacting OGDP with each data-provision proxy in Equation (1). The interaction coefficients are not statistically significant across specifications (Figure E-1 in Appendix E).

salient in healthcare AI for two reasons. First, hospitals are especially responsive to government signals because they operate under stringent regulatory oversight and tend to be more risk-averse than firms in less-regulated sectors. The gap between rapidly advancing AI technologies and lagging regulatory frameworks leaves managers uncertain about the legitimacy and future compliance requirements of AI-related data practices. OGDPs signal that data-driven innovation is both permissible and institutionally supported, reducing this perceived regulatory risk and lowering hospitals' hesitation to invest. Second, vendors facing incomplete information about future regulatory conditions and market demand often perceive healthcare AI as a high-risk domain. Credible signals of governmental support reduce perceived policy risk and increase expected returns, encouraging vendors to sustain R&D investment, bring more deployable AI products to market, and engage more actively with local hospitals — collectively raising the overall level of healthcare AI procurement. If the *governmental signaling* mechanism holds, OGDPs that send stronger signals should have larger effects.

To operationalize this channel, we draw on signaling theory, which holds that the strength of a signal depends on two dimensions: (i) how costly it is for the sender, and (ii) how observable it is to receivers (Spence, 2002; Connelly et al., 2011). We construct two proxies to capture variation in the strength of governmental signaling, with each proxy corresponding to a single dimension.

The first proxy captures the cost dimension: the *number of non-healthcare government departments contributing data to the OGDP*, measured as the count of non-healthcare local agencies releasing datasets on the portal. We exclude healthcare regulators (e.g., the Health Commission, the Medical Insurance Bureau, and the Drug Administration) so that the measure reflects the breadth of governmental participation rather than direct healthcare-data supply, thereby separating the signaling channel from the data-provision channel examined above. Mobilizing multiple non-healthcare departments to release data through a unified portal entails substantial organizational and coordination costs, making broad inter-agency participation a hard-to-mimic signal of institutionalized commitment that AI vendors and hospitals can read as a sustained policy direction rather than a single-department initiative.[14]

The second proxy captures the observability dimension: the *intensity of governmental media communication regarding the OGDP*, captured by an index that aggregates standardized counts of OGDP-related posts on the city government's official WeChat public account and Sina Weibo account per quarter.[15] Public communication renders governmental commitment directly observable to ecosystem

[14] Broader inter-agency participation may also reflect general state capacity or inter-agency coordination. The alternative-explanations analyses (Section 4.4) show that OGDP effects do not concentrate in cities with stronger fiscal or digital capacity, attenuating the most direct capacity-based alternative readings.

[15] Both WeChat and Sina Weibo are dominant social-media platforms in China, with the former serving long-form official communication and the latter facilitating broader public discourse.

participants who would otherwise have limited access to internal administrative actions; the more frequently a city government publicly amplifies the OGDP through its official outlets, the more visible and credible its policy stance becomes to hospitals and AI vendors. To construct this measure, we collected all publicly available posts from each prefecture-level city government's verified WeChat public account and Sina Weibo account (where such accounts exist) during 2012–2023, retained posts whose titles or texts referenced open government data or OGDP-related keywords, separately standardized the WeChat and Weibo quarterly post counts (z-score transformation) so as to place them on a common scale despite differences in baseline posting volume across the two media platforms, and summed the two standardized series into a single composite signaling-intensity index.

We then assess the governmental signaling mechanism using subsample analyses. For each proxy above, we split cities at the sample-period median and estimate Equation (1) separately on the high and low subsamples (each paired with the never-treated control cities). We expect a stronger OGDP effect where the government's signal is stronger along the corresponding dimension.

Figure 6 presents the results. Consistent with this prediction, the launch of OGDPs increases local healthcare AI investment in both high and low subsamples, but the effect is substantially and significantly larger in cities with (i) more participating non-healthcare departments and (ii) higher intensity of governmental media communication (inter-group difference $p < 0.01$ in both cases). These patterns provide evidence that OGDPs function as a credible governmental signal that influences local healthcare AI investment along both the cost and observability dimensions.[16]

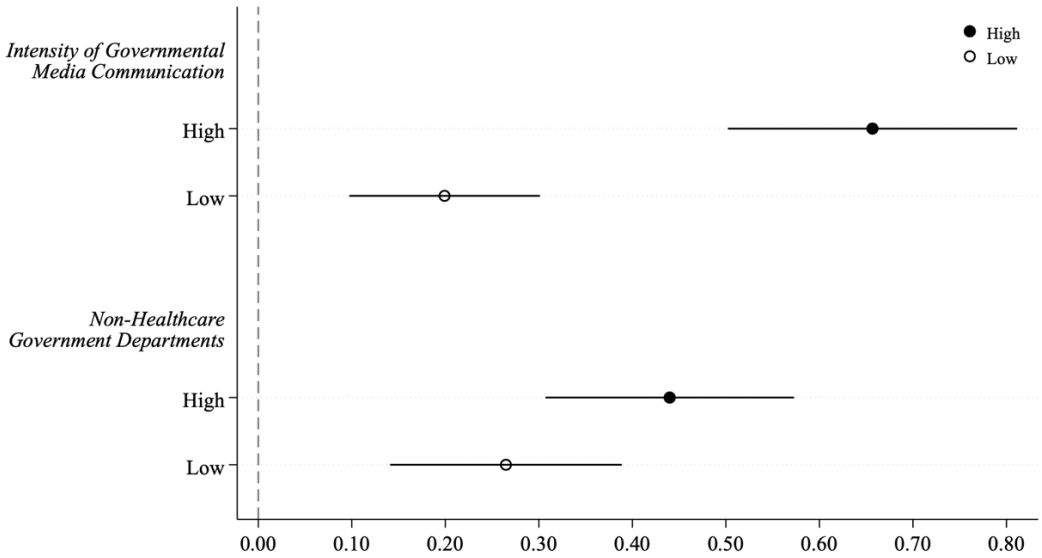


Figure 6. Subsample Analysis Results for the Governmental Signaling Mechanism

[16] We further re-estimate the mechanism using continuous interaction terms, interacting OGDP with each governmental-signaling proxy in Equation (1). The interaction coefficients remain positive and statistically significant across specifications (Figure E-1 in Appendix E).

Notes: This figure presents subsample tests of the governmental signaling mechanism. We construct two proxies, each capturing one dimension of governmental signaling: (i) the number of non-healthcare government departments providing open data (the cost dimension), measured as the count of non-healthcare local agencies releasing datasets on the OGDP; and (ii) the intensity of governmental media communication regarding the OGDP (the observability dimension), measured by a composite index that aggregates standardized counts (z-score) of OGDP-related posts on the city government's official WeChat public account and Sina Weibo account in each year-quarter. The first proxy is manually collected from each city's OGDP; the second is constructed from posts on each city government's verified WeChat and Weibo accounts. For each proxy, we split cities at the sample-period median and estimate Equation (1) separately on the high and low OGDP-treated groups (each paired with the never-treated control cities), then compare the OGDP effects across groups.

5.3. Relative Explanatory Power of Data Provision and Governmental Signaling

To compare the two channels directly, we estimate a single specification that includes all four mechanism proxies interacted with OGDP, jointly with the OGDP indicator and baseline controls. This joint specification allows us to assess whether each mechanism's explanatory power survives the inclusion of the other, addressing the concern that subsample results in Sections 5.1 and 5.2 might be driven by correlation between the two sets of proxies rather than by independent channels. The OGDP coefficient is positive and statistically significant ($\beta = 0.235$, $p < 0.01$), supporting the substantial average effect.

Figure 7 reports the interaction effects. The data-provision interactions are not statistically significant ($p > 0.10$), while both signaling-related interactions remain positive and significant ($p < 0.01$). Cross-city variation in the costliness of inter-agency participation and in the observability of governmental media communication amplifies the OGDP effect; cross-city variation in measured data supply does not. Taken together with the subsample evidence above, these results indicate that data provision is a necessary input that enables the OGDP effect, while governmental signaling is the more proximate driver of cross-city variation in treatment intensity. This pattern is consistent with the dual-channel framework developed in Section 2.3, in which OGDPs function not only as a data conduit but, more centrally in our setting, as a credible signal of policy commitment.

The asymmetry between the two channels aligns with specific institutional features of healthcare AI. While public data infrastructure provides a necessary foundation for AI development, hospitals and AI vendors in healthcare also face substantial regulatory ambiguity, including uncertainty about how regulations will evolve and how strongly governments will support data-driven innovation. In this environment, governmental signals that resolve such uncertainty carry meaningful marginal value beyond the data itself. Healthcare is also a highly regulated and risk-averse sector, where investment decisions tend to be particularly sensitive to perceptions of policy direction (Sun and Medaglia, 2019; Ellahham et al., 2020). These features offer one plausible explanation for the pattern observed in our setting. We acknowledge, however, that our subsample and interaction designs cannot perfectly isolate the two channels, and the patterns we report should be read as suggestive evidence on the relative prominence of the two channels rather than as definitive confirmation of mechanism dominance.

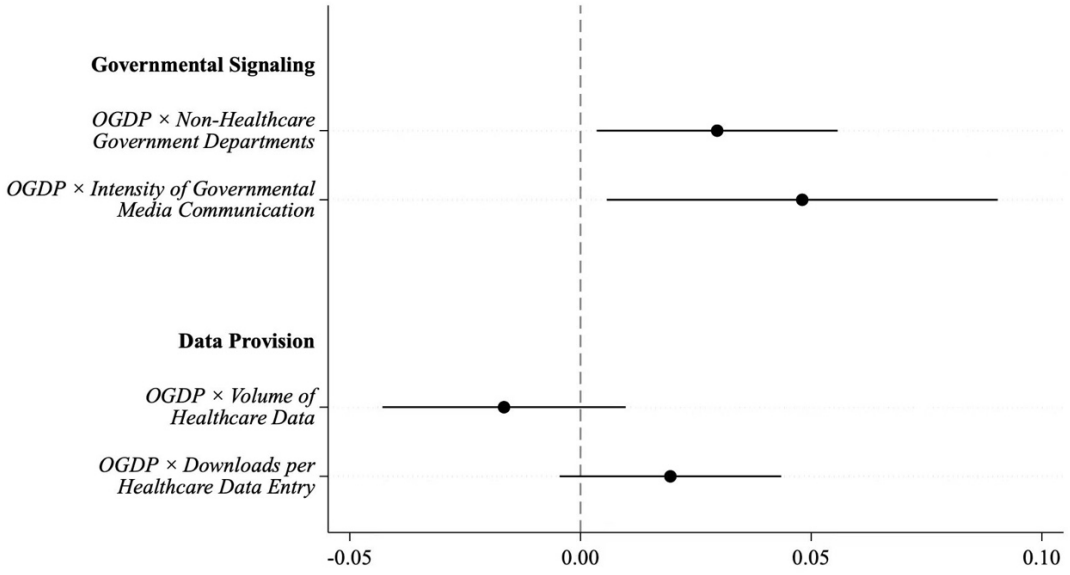


Figure 7. Analysis of Both Mechanisms Underlying the OGDP Effect

Notes: This figure reports coefficient estimates from a regression that jointly examines the data provision and governmental signaling mechanisms underlying the OGDP effect. We interact OGDP with four mechanism-related proxies introduced in Sections 5.1 and 5.2 and simultaneously include these interaction terms in Equation (1). The two Governmental Signaling proxies capture (i) the number of non-healthcare government departments providing open data (the cost dimension) and (ii) the intensity of governmental media communication regarding the OGDP, measured as a standardized composite of post volumes on the city government's official WeChat and Weibo accounts (the observability dimension). The two Data Provision proxies capture (i) the total volume of healthcare data released on the OGDP and (ii) the download volume per unit of healthcare data. All mechanism-related proxies are constructed from each city's OGDP, related procurement records, and the corresponding municipal-government social media accounts.

6. Effect Heterogeneity

Whereas Section 5 examines the internal mechanisms through which OGDPs influence healthcare AI investment, this section examines the external contextual conditions that moderate the strength of the OGDP effect across cities. We organize the analysis around three core participants in the healthcare AI ecosystem whose responsiveness determines whether OGDPs translate into investment outcomes: local governments (the OGDP builder), hospitals (the demand-side investor), and AI vendors (the supply-side developer). The corresponding local conditions are, respectively, government digitalization priorities, healthcare demand pressure, and local AI vendor supply. Results are reported in Table F-1 in Appendix F.

The first concerns local governments' digitalization priorities. Drawing on signaling theory (Spence, 1973; Connelly et al., 2011), policy signals are more credible when aligned with a sustained policy orientation. When digitalization is a core governmental priority, an OGDP launch is more likely to be perceived as part of a broader strategic commitment rather than a symbolic gesture, easing regulatory uncertainty and facilitating AI investment. This is conceptually distinct from pre-existing local digital capacity examined in Section 4.4: digital capacity captures existing technological infrastructure, while digitalization priority signals governmental commitment that shapes expectations about future regulatory

support. Digitalization priority may also reinforce the data provision channel, as governments that prioritize digitalization typically exhibit stronger implementation capacity and cross-departmental coordination, improving data standardization and usability. Subsample analyses confirm significantly larger OGDP effects in cities with higher governmental digitalization priority (Table F-1 in Appendix F).

A second contextual condition is healthcare demand pressure. While this condition is shared by hospitals and vendors operating in the same local healthcare system, hospitals are the focal investing actors whose responses we observe. Consistent with demand-pull theories of innovation, technological investment responds to the size and urgency of local demand (Acemoglu and Linn, 2004). In cities where healthcare systems operate under greater strain, hospitals face greater urgency to expand service capacity and improve operational efficiency, while AI vendors encounter clearer and more immediate application needs. Under these conditions, OGDP launches are more likely to translate into concrete investment responses rather than remain symbolic policy gestures. Accordingly, we find significantly stronger OGDP effects in cities with higher hospital bed occupancy rates (Table F-1).

A third condition is the local supply of healthcare AI vendors. Geographic clustering enhances knowledge spillovers, reduces coordination costs, and improves matching between complementary actors (Porter, 1998). OGDPs release data and convey policy signals, but their effects depend on whether local markets possess sufficient capacity to absorb, interpret, and operationalize these inputs. Local healthcare AI vendors are more familiar with municipal policies and better connected to local hospitals, and are better positioned to transform publicly released data into AI products and to respond strategically to the policy signals embedded in OGD initiatives. We measure local vendor supply as the share of healthcare AI procurement contracts in a city–quarter awarded to vendors registered in the same city as the procuring hospital, identified through legal entity names and Unified Social Credit Codes matched against national and provincial firm-registration databases. Splitting cities at the median value of this share, we find that OGDP effects are stronger in cities with a higher share of local vendors (Table F-1).

Together, across the three ecosystem participants, the returns to public data infrastructure are highest where local governments demonstrate sustained digital commitment, where healthcare systems face acute demand pressure, and where local AI vendors have the capacity to absorb and operationalize public data.

7. Discussion and Conclusion

7.1. Summary of Findings

This study examines how OGDPs influence local hospital investment in healthcare AI. Using a staggered DiD design with a panel of 258 Chinese cities (2012-2023), we find that the launch of city-level OGDPs is associated with a significant and sustained increase in local healthcare AI investment, encompassing both diagnostic and operational AI systems. Mechanism analyses point to two channels: data provision as a

necessary baseline input, and governmental signaling as the channel that systematically conditions cross-city variation in treatment intensity. While both channels plausibly contribute, signaling emerges as the more proximate driver in our setting, suggesting that OGDPs function not only as data conduits but, more centrally, as credible signals of policy commitment that reduce regulatory uncertainty for both hospitals and vendors. Heterogeneous analyses show that the OGDP effect is amplified by local governments' digitalization priorities, healthcare demand pressure, and the local supply of healthcare AI vendors.

7.2. Contributions to Research

This work makes three primary contributions to the information systems (IS) literature and related fields. Collectively, they move the study of public data from a question of access to one of institutional design, reframing open data not as a passive repository but as an active instrument that shapes whether and where AI innovation occurs.

First, we advance the emerging IS and economics literature on digital infrastructure and innovation (Tilson et al., 2010; Nambisan et al., 2017), and, more specifically, the data and AI innovation stream within it, by elucidating how public data infrastructure shapes AI investment. Existing research has examined how privately accumulated data drives AI innovation and how firm-state data exchange shapes AI innovation, but evidence on whether and how openly accessible public data shapes AI innovation at the local level remains limited. By providing evidence on the effect of OGDPs on local healthcare AI investment, and by disentangling data provision from governmental signaling as candidate channels, we extend a literature that has primarily emphasized the direct material role of data in AI innovation (e.g., Beraja et al., 2023; Jones and Tonetti, 2020; Nagaraj, 2022). Our findings also help reconcile the mixed evidence on whether public data assets stimulate organizational innovation: where prior work has documented both positive (Nagaraj, 2022; Magalhaes and Roseira, 2020) and negative (Janssen et al., 2012; Zuiderwijk and Janssen, 2014) effects, we show that public data infrastructure can stimulate downstream investment, with the more prominent channel in our setting being cognitive (signaling) rather than purely material (data access). This resolution is itself substantive: it suggests that the conflicting prior findings may stem not from whether public data matters, but from whether scholars have attended to the cognitive channel through which it operates. By documenting these unintended innovation spillovers, we also extend the OGD literature beyond its traditional focus on transparency and accountability benefits (Jetzek et al., 2019; Attard et al., 2015).

Second, we extend research on healthcare AI by shifting the level of analysis from intra-organizational adoption to the broader local institutional environment. Prior work has identified important micro-level barriers to hospital AI investment, including limited trust, privacy constraints, and capacity gaps (e.g., Abdel-Karim et al., 2023; Liang and Xue, 2022; Reis et al., 2020). Responding to recent calls for research on the "missing macro layer" in healthcare AI investment (Baird et al., 2020; Sun and

Medaglia, 2019), we show that macro-level government action, particularly the establishment of open public data portals, can meaningfully increase local investment in both diagnostic and operational healthcare AI systems. We further extend the ecosystem perspective on healthcare AI by demonstrating how government-led open data initiatives coordinate heterogeneous actors through both material and cognitive channels. While prior studies view governments primarily as orchestrators aligning resources through regulation and funding (Leone et al., 2021; Sun and Medaglia, 2019), our findings suggest that OGDPs operate through a dual mechanism combining data provision (material coordination) with governmental signaling (cognitive coordination), aligning expectations and strengthening investment confidence across technology vendors and hospitals. This dual-channel account offers a more complete explanation of how government action mobilizes ecosystem-wide investment than either resource- or regulation-centered view alone.

Third, we contribute to the literature on institutions and directed technical change (Acemoglu et al., 2012), which emphasizes how institutional arrangements shape the direction of technological development. We conceptualize OGDPs as an inclusive institutional arrangement that lowers data-access barriers, reduces regulatory uncertainty, and thereby steers investment toward socially valuable healthcare AI applications. This perspective extends the literature in two ways. First, we bring it into a new empirical context, showing that institutional arrangements directing technical change can operate through public data infrastructure that simultaneously functions as a material input and credible signal. Second, we provide evidence that such institutional arrangements can plausibly direct AI innovation toward domains of broad social importance, enriching our understanding of how institutions align technological change with societal priorities (Acemoglu, 2021; Taddeo and Floridi, 2018). In an era of intense debate over how to govern AI, this evidence identifies public data infrastructure as a concrete, deployable instrument through which states can shape not only the pace but also the direction of AI development.

7.3. Implications for Practice

This work also offers practical guidance for multiple stakeholders.

For policymakers, our findings offer evidence-based guidance on how to maximize returns from public data infrastructure investments. Realizing the value of open data requires more than expanding data catalogs; it requires credible governmental commitment that signals sustained policy support to hospitals, AI vendors, and other ecosystem actors. Three actions are particularly valuable. First, OGDP launches should be accompanied by visible inter-agency participation, with multiple government departments contributing data — signaling that the initiative reflects institutionalized commitment rather than a single-agency effort. Second, governments should actively communicate about the OGDP through official channels (e.g., social and mass media) to ensure that policy signals reach relevant ecosystem actors. Third,

policymakers should treat open data initiatives as long-term investments rather than one-time launches, sustaining engagement through regular dataset updates and continued institutional support.

For AI vendors, OGDPs provide signals of market readiness that can guide commercialization decisions in environments where regulatory uncertainty has historically slowed adoption. Our results suggest that operating in jurisdictions with visible governmental support through OGDPs provides a strategic advantage by facilitating the translation of R&D investment into local sales. Vendors may accelerate deployment and prioritize hospital engagement in markets where signals are strong, while adopting more cautious, pilot-based approaches where signals are weak or fragmented, until institutional support becomes clearer.

For hospital administrators, government data initiatives can be read as indicators of reduced regulatory uncertainty around AI adoption, warranting a reassessment of investment strategies. Global healthcare systems face persistent resource constraints, creating strong incentives to adopt AI technologies. Hospitals operating in environments with strong OGDP signals may be better positioned to prioritize digital transformation, partner with AI vendors aligned with policy-supported objectives, and gain first-mover advantages across both clinical decision-making and operational domains.

7.4. Limitations and Future Research

To understand the extent to which our findings inform real-world considerations and applications, we highlight several important limitations. The first concerns the unit of analysis. Our city-level design provides clean identification but cannot capture how individual hospitals and AI vendors respond to OGDP initiatives. Future work could construct a complementary organizational-level dataset linking hospital AI investment with vendor R&D and commercial activities, which enables examination of how strategic responses vary across actors with different capabilities and market positions.

A second limitation concerns the outcomes we observe. We measure healthcare AI investment but not its consequences, such as clinical quality or patient outcomes. Future research could draw on hospitals' operational indicators, personnel characteristics, and patient-level treatment data to evaluate broader welfare and performance outcomes.

Finally, our findings should be interpreted in light of the institutional context in which they are generated. The Chinese setting offers strong identification through the staggered rollout of city-level OGDPs, but the relative weight of governmental signaling and data provision may differ in contexts with distinct regulatory baselines, data-governance regimes, or policy-communication norms. Cross-context investigation thus represents an important avenue for future research, both to assess the broader applicability of our findings and to clarify which features of public data infrastructure most reliably translate into innovation outcomes across institutional environments.

8. Concluding Remarks

This study examines how government-led open data initiatives shape local healthcare AI investment. While prior research on public data has largely overlooked its spillovers for AI innovation, our analysis reveals a more nuanced reality: public data infrastructure can do more than provide information — it can actively redirect AI investment. Exploiting temporal and regional variations in OGDP launches across Chinese cities, we show that public data openness promotes healthcare AI investment through two candidate mechanisms, data provision and governmental signaling, with evidence pointing to signaling as the relatively more prominent channel in our setting. These findings contribute to broader debates on the need for sustained investment in public data infrastructure, highlighting how inclusive institutions can steer technological trajectories toward socially valuable applications. Future research should examine how different forms of public data influence AI innovation across sectors and countries, particularly in contexts with varying digital infrastructure and regulatory capacity. Doing so will deepen understanding of how governments can direct AI development toward socially valuable ends.

References

Abdel-Karim, B. M., Pfeuffer N, Carl K. V., Hinz O. (2023). How AI-based systems can induce reflections: The case of AI-augmented diagnostic work. *MIS Quarterly*, 47(7), 1395–1424.

Attard, J., Orlandi, F., Scerri, S., & Auer, S. (2015). A systematic review of open government data initiatives. *Government Information Quarterly*, 32(4), 399-418.

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Exploring the impact of artificial intelligence: Prediction versus judgment. *Information Economics and Policy*, 47, 1-6.

Autor, D. H. (2003). Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing. *Journal of Labor Economics*, 21(1), 1–42.

Acemoglu, D. (Ed.). (2021). *Redesigning AI*. MIT Press.

Acemoglu, D., Makhdoumi, A., Malekian, A., & Ozdaglar, A. (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4), 218-256.

Acemoglu, D., Aghion, P., Bursztyn, L., & Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1), 131-166.

Acemoglu, D., & Linn, J. (2004). Market size in innovation: theory and evidence from the pharmaceutical industry. *The Quarterly Journal of Economics*, 119(3), 1049-1090.

Baird, A., Angst, C., & Oborn, E. (2020). Research Curation on Health Information Technology. *MIS Quarterly*. https://www.misqresearchcurations.org/blog/2018/6/20/health-information-technology

Bellemare, M. F., & Wichman, C. J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, 82(1), 50-61.

Beraja, M., Yang, D. Y., & Yuchtman, N. (2023). Data-intensive innovation and the state: Evidence from AI firms in China. *The Review of Economic Studies*, 90(4), 1701-1723.

Beane, M. (2023). Resourcing a technological portfolio: How Fairtown Hospital preserved results while degrading its older surgical robot. *Administrative Science Quarterly*, 68(3), 691-733.

Besley, T., & Persson, T. (2009). The origins of state capacity: Property rights, taxation, and politics. *American Economic Review*, 99(4), 1218-1244.

Baker, A. C., Larcker, D. F., & Wang, C. C. Y. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2), 370–395.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254-277.
Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200-230.
Coiera, E., & Liu, S. (2022). Evidence synthesis, digital scribes, and translational challenges for artificial intelligence in healthcare. *Cell Reports Medicine*, 3(12), 100860.
Connelly, B. L., Certo, S. T., Ireland, R. D., & Reutzel, C. R. (2011). Signaling theory: A review and assessment. *Journal of Management*, 37(1), 39-67.
Chen, J., & Roth, J. (2024). Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics*, 139(2), 891-936.
Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3), 1405-1454.
Dai, T., & Tayur, S. (2022). Designing AI-augmented healthcare delivery systems for physician buy-in and patient acceptance. *Production and Operations Management*, 31(12), 4443-4451.
Duran, P., van Essen, M., Heugens, P. P., Kostova, T., & Peng, M. W. (2019). The impact of institutions on the competitive advantage of publicly listed family firms in emerging markets. *Global Strategy Journal*, 9(2), 243-274.
De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964-2996.
Ellahham, S., Ellahham, N., & Simsekler, M. C. E. (2020). Application of artificial intelligence in the health care safety context: opportunities and challenges. *American Journal of Medical Quality*, 35(4), 341-348.
Estrin, S., Meyer, K. E., Nielsen, B. B., & Nielsen, S. (2016). Home country institutions and the internationalization of state owned enterprises: A cross-country analysis. *Journal of World Business*, 51(2), 294-307.
Gulen, H., & Ion, M. (2016). Policy uncertainty and corporate investment. *The Review of Financial Studies*, 29(3), 523-564.
Hughes-Cromwick, E., & Coronado, J. (2019). The value of US government data to US business decisions. *Journal of Economic Perspectives*, 33(1), 131-146.
Howell, S. T. (2017). Financing innovation: Evidence from R&D grants. *American Economic Review*, 107(4), 1136-1164.
Han, M., Shen, H., Wu, J., & Zhang, X. (M.). (2025). Artificial intelligence and firm resilience: Empirical evidence from natural disaster shocks. *Information Systems Research*, 36(4), 2116–2133.
Jetzek, T., Avital, M., & Bjørn-Andersen, N. (2019). The sustainable value of open government data. *Journal of the Association for Information Systems*, 20(6), 702-734.
Janssen, M., Charalabidis, Y., & Zuiderwijk, A. (2012). Benefits, adoption barriers and myths of open data and open government. *Information Systems Management*, 29(4), 258-268.
Jones, C. I., & Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9), 2819–2858.
Lemoine D (2025). Big Beautiful Numbers: The Impact of Governmental Data on Firms' Valuation. National Bureau of Economic Research.
Leone, D., Schiavone, F., Appio, F. P., & Chiao, B. (2021). How does artificial intelligence enable and enhance value co-creation in industrial markets? An exploratory case study in the healthcare ecosystem. *Journal of Business Research*, 129, 849-859.
Lnenicka, M., Nikiforova, A., Luterek, M., Milic, P., Rudmark, D., Neumaier, S., & Bolívar, M. P. R. (2024). Identifying patterns and recommendations of and for sustainable open data initiatives: A benchmarking-driven analysis of open government data initiatives among European countries. *Government Information Quarterly*, 41(1), 101898.
Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, 46(4), 629-650.

Liang, H., & Xue, Y. (2022). Save face or save life: Physicians' dilemma in using clinical decision support systems. *Information Systems Research*, 33(2), 737-758.
Li, B., Huang, N., & Shi, W. (2025). Forced to change? Media exposure of labor issues and firm artificial intelligence investment. *Information Systems Research*, 37(1), 156–175.
Magalhaes, G., & Roseira, C. (2020). Open government data and the private sector: An empirical view on business models and value creation. *Government Information Quarterly*, 37(3), 101248.
Nagaraj, A. (2022). The private impact of public data: Landsat satellite maps increased gold discoveries and encouraged entry. *Management Science*, 68(1), 564-582.
Nambisan, S., Lyytinen, K., Majchrzak, A., & Song, M. (2017). Digital innovation management: Reinventing innovation management research in a digital world. *MIS Quarterly*, 41(1), 223-238.
Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453.
Pham, P., Zhang, H., Gao, W., & Zhu, X. (2024). Determinants and performance outcomes of artificial intelligence adoption: Evidence from US Hospitals. *Journal of Business Research*, 172, 114402.
Park, E. H., Werder, K., Cao, L., & Ramesh, B. (2022). Why do family members reject AI in health care? Competing effects of emotions. *Journal of Management Information Systems*, 39(3), 765-792.
Porter, M. E. (1998). Clusters and the new economics of competition (Vol. 76, No. 6, pp. 77-90). Boston: Harvard Business Review.
Reis, L., Maier, C., Mattke, J., Creutzenberg, M., & Weitzel, T. (2020). Addressing user resistance would have prevented a healthcare AI project failure. *MIS Quarterly Executive*, 19(4).
Robinson, J. A., & Acemoglu, D. (2012). *Why nations fail: The origins of power, prosperity and poverty* (pp. 45-47). London: Profile.
Schumpeter, J. A., & Swedberg, R. (2021). *The theory of economic development*. Routledge.
Sun, T. Q., & Medaglia, R. (2019). Mapping the challenges of Artificial Intelligence in the public sector: Evidence from public healthcare. *Government Information Quarterly*, 36(2), 368-383.
Spence, M. (2002). Signaling in retrospect and the informational structure of markets. *American Economic Review*, 92(3), 434-459.
Silva, P. N., & Silva, S. N. (2025). Open government data in health sector: a systematic literature review. BMC Public Health, 25, 47.
Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175-199.
Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, 87(3), 355–374.
Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361(6404), 751-752.
Tilson, D., Lyytinen, K., & Sørensen, C. (2010). Research commentary—Digital infrastructures: The missing IS research agenda. *Information Systems Research*, 21(4), 748-759.
Vasudeva, G. (2013). Weaving together the normative and regulative roles of government: How the Norwegian Sovereign Wealth Fund's responsible conduct is shaping firms' cross-border investments. *Organization Science*, 24(6), 1662-1682.
Vasudeva, G., Nachum, L., & Say, G. D. (2018). A signaling theory of institutional activism: How Norway's sovereign wealth fund investments affect firms' foreign acquisitions. *Academy of Management Journal*, 61(4), 1583-1611.
Wirtz, B. W., Weyerer, J. C., Becker, M., & Müller, W. M. (2022). Open government data: A systematic literature review of empirical research. *Electronic Markets*, 32(4), 2381-2404.
Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1), 57-76.
Zuiderwijk, A., & Janssen, M. (2014, June). The negative effects of open government data-investigating the dark side of open data. *In Proceedings of the 15th Annual International Conference on Digital Government Research* (pp. 147-152).
Zuiderwijk, A., & Reuver, M. D. (2021). Why open government data initiatives fail to achieve their objectives: Categorizing and prioritizing barriers through a global survey. *Transforming government: people, process and policy*, 15(4), 377-395.

The Impact of Public Data on AI Innovation:
Open Government Data Initiatives Encourage Local Healthcare AI Investment

ONLINE SUPPLEMENTARY APPENDICES

Appendix A: Context and Data

A.1. Additional Detail on the Design and Use of Municipal OGDPs

To promote the secondary utilization of public data, Open Government Data Portals (OGDPs) have emerged as important institutional arrangements. These online portals are established and operated by local municipal government authorities, providing various types of data generated, collected, and managed in the course of public administration and public service delivery. The data are typically released in machine-readable formats that support bulk downloading, and they are made available on a non-exclusive basis.

Section 3.1 of the main text introduces the OGDPs studied in this paper and their staggered rollout. This appendix supplements that discussion with four pieces of additional detail not covered in the main text: the overall design and search interface of a representative municipal OGDP (Figure A-1), the multi-dimensional classification system used to organize datasets (Figure A-2), the standardized metadata accompanying each released dataset (Figure A-3), and direct evidence that portal traffic is overwhelmingly local rather than spilling over across municipalities (Figure A-4).

Overall design and search interface. Municipal OGDPs share several common design features established under national guidelines. A central search bar enables keyword-based dataset retrieval, while sidebars and tab structures allow users to browse datasets along multiple categorical dimensions. Figure A-1 shows the homepage of the Hangzhou OGDP, illustrating both functions.



Figure A-1. An Example of the Design Features of Hangzhou OGDP

Notes: This screenshot displays the homepage of the Hangzhou Open Government Data Portal (https://data.hangzhou.gov.cn), captured on December 12, 2025.

Dataset-classification interface. Datasets on each municipal OGDP are organized through a multi-dimensional classification interface that allows users to retrieve data along several complementary axes. Figure A-2 illustrates the dataset-classification interface of the Hangzhou OGDP, which classifies datasets along four dimensions: application scenario, domain (e.g., medical and health, public safety, technological innovation), issuing department, and district within the city. This structure lowers search costs for external users and enables targeted retrieval of healthcare-relevant data.

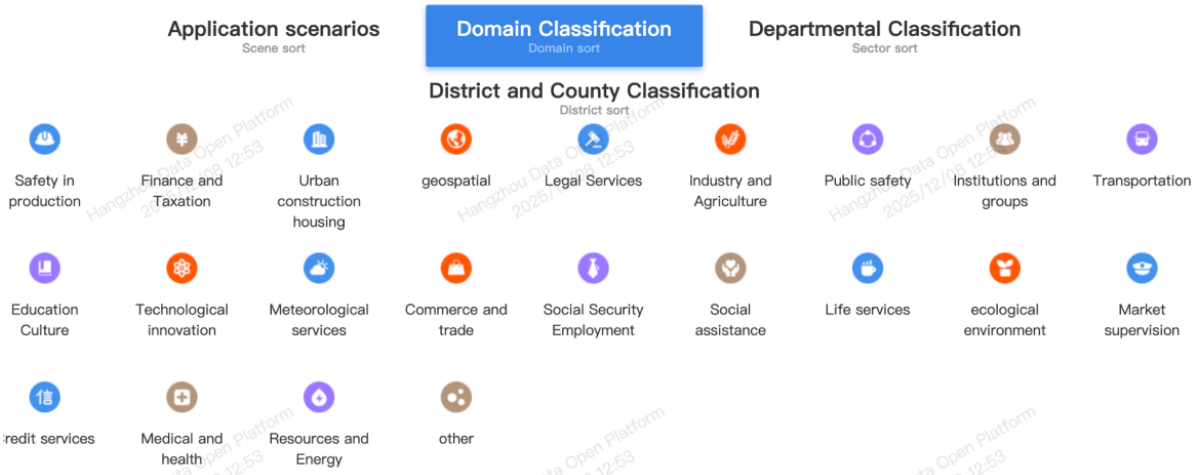


Figure A-2. Illustration of Dataset Classification on a Municipal OGDP

Notes: This screenshot, captured from the Hangzhou Open Government Data Portal (https://data.hangzhou.gov.cn) on December 12, 2025, illustrates its hierarchical classification system. The portal categorizes datasets across four core dimensions: application scenarios, domain classification, departmental classification, and district and county classification, covering key domains such as medical and health, public safety, and technological innovation.

Metadata accompanying released datasets. Each dataset on a municipal OGDP is published with a standardized metadata record that includes the dataset title, the responsible government agency, domain classification, publication and update timestamps, user-engagement metrics (views and downloads), and access conditions (unrestricted or application-based). These structured attributes enhance dataset discoverability and traceability for external users. Figure A-3 shows the metadata page of an illustrative dataset on the Hangzhou OGDP.

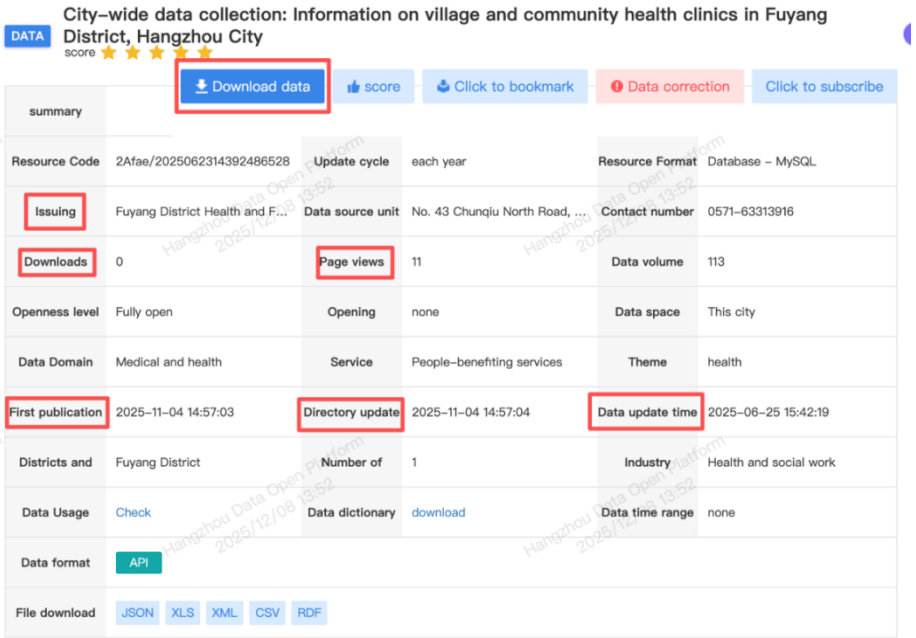


Figure A-3. Illustration of Detailed Dataset Information on a Municipal OGDP

Notes: This screenshot displays the detailed metadata page of a dataset on the Hangzhou OGDP (https://data.hangzhou.gov.cn), captured on December 12, 2025. Key attributes are highlighted, including the issuing authority, download metrics, access statistics, publication/update timestamps, and multi-format download options (JSON, XLS, API). The portal provides standardized, machine-readable, and regularly updated data resources for public and commercial use.

Local concentration of portal traffic. Although Chinese municipal OGDPs are technically accessible nationwide, in practice their usage is overwhelmingly concentrated among users located within the issuing municipality. Figure A-4 illustrates this pattern using publicly displayed visit-origin statistics for the Hangzhou OGDP. Hangzhou-originated visits account for 5,042,867 of the top ten cities' combined 5,144,524 visits (approximately 98%), exceeding visits from the second-ranked city (Shanghai) by a factor of nearly 150. This local concentration supports the within-city interpretation of our identification strategy: OGDP-induced effects on healthcare AI investment are absorbed primarily by hospitals and vendors operating within the issuing city, rather than spilling over to neighboring municipalities.



Figure A-4. Geographic Distribution of Visits to Hangzhou OGDP

Notes: This figure presents the geographic distribution of visits to the Hangzhou Open Government Data Portal (https://data.hangzhou.gov.cn) as displayed on the portal's public usage dashboard on December 4, 2025.

A.2. Acquiring the Launch Dates of OGDPs

According to data released by the National Bureau of Statistics of China, as of December 2023, there were 297 prefecture-level and above cities in China. This figure includes the four municipalities directly under the central government (Beijing, Shanghai, Tianjin, and Chongqing) but excludes special prefecture-level administrative units such as autonomous prefectures and leagues. To ensure comparability across cities, we excluded those located in the Inner Mongolia Autonomous Region, Guangxi Zhuang Autonomous Region, Tibet Autonomous Region, Ningxia Hui Autonomous Region, and Xinjiang Uygur Autonomous Region, as these regions are characterized by substantial data gaps and considerable institutional differences in governmental administration compared with other provinces. In addition, cities that experienced changes in administrative status (e.g., downgrading from prefecture-level to county-level) were also excluded to avoid potential inconsistencies arising from structural breaks in statistical reporting. The final sample consists of 258 cities.

For these 258 cities, we collected official OGDP launch dates through a multi-step, verification procedure designed to maximize the use of first-hand evidence. First, we directly contacted the platform operators of each city-level OGDP via official website messaging systems or contact forms, requesting confirmation of the launch date as recorded by the responsible government agency. This approach provides the most direct, agency-confirmed source of launch dates and constitutes the primary basis of our launch-date dataset.

Second, to complement and cross-validate the records obtained through direct contact, we conducted systematic online searches using multiple search engines, including Google, Baidu, and Bing. Keyword combinations such as "[city name] + data openness," "[city name] + data resources," and "[city name] + public data" were employed, with searches restricted to official government domains (e.g., websites containing "gov.cn") to ensure data reliability. Where direct contact was not feasible or responses were unavailable, launch dates were triangulated using these online sources together with publicly available secondary sources, including official news reports, press releases, government announcements, and recorded dates of initial dataset publication on the platform itself.

Finally, as an additional cross-check, we compared our manually collected dates against the *China Local Government Data Open Report (Cities)* series, published by the Digital and Mobile Governance Lab (DMG) at Fudan University, which systematically tracks the construction and launch of OGDPs across China and has been referenced by multiple local governments. In cases of discrepancy between our records and the DMG series, the manually collected dates were treated as authoritative, given that they originate from direct verification or multi-source triangulation rather than secondary reporting. All verified OGDP launch dates were then compiled into a structured dataset containing each city's name and the corresponding launch date. Figure A-5 plots the staggered adoption of OGDPs across the 258 sample cities from 2012 to 2023, showing that OGDP launches were spread over more than a decade, with substantial variation in adoption timing across cities.

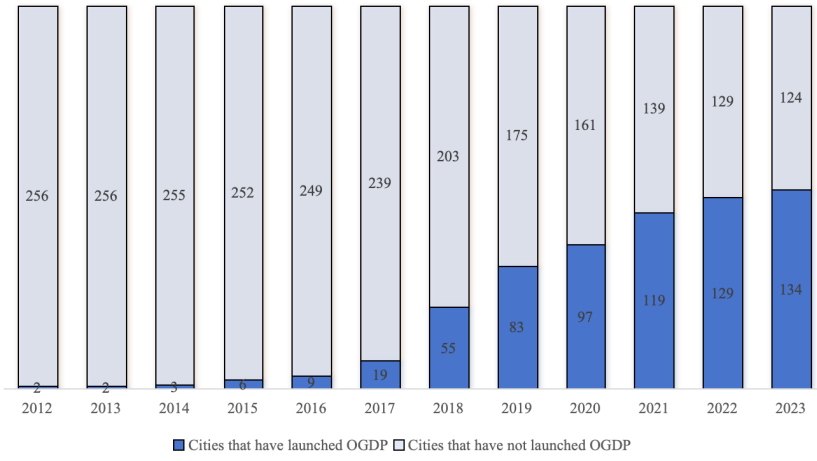


Figure A-5. The Staggered Adoption of OGDPs Across Cities Over the Period 2012–2023

A.3. Acquiring Healthcare AI Investment Data

To measure local hospitals' investment in artificial intelligence (AI), we follow the approach of Beraja et al. (2023) and use awarded procurement contracts for AI-related medical projects as our primary proxy. The construction procedure, including contract restriction, AI identification and validation, and panel aggregation, is summarized in Figure A-6.

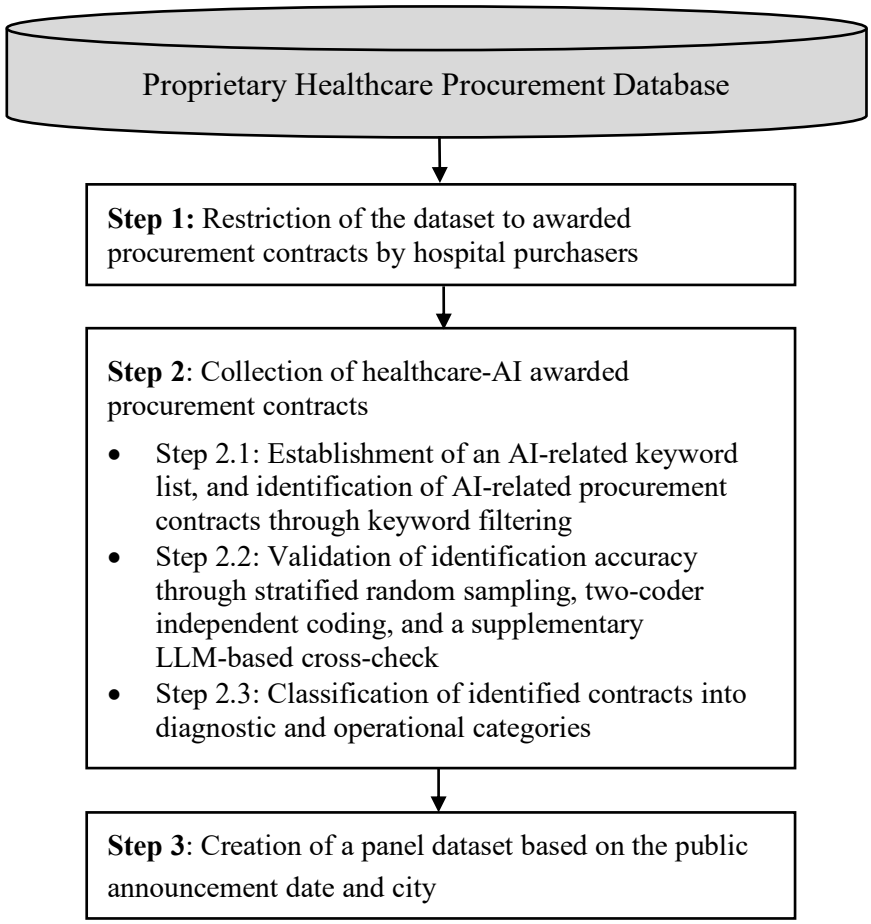
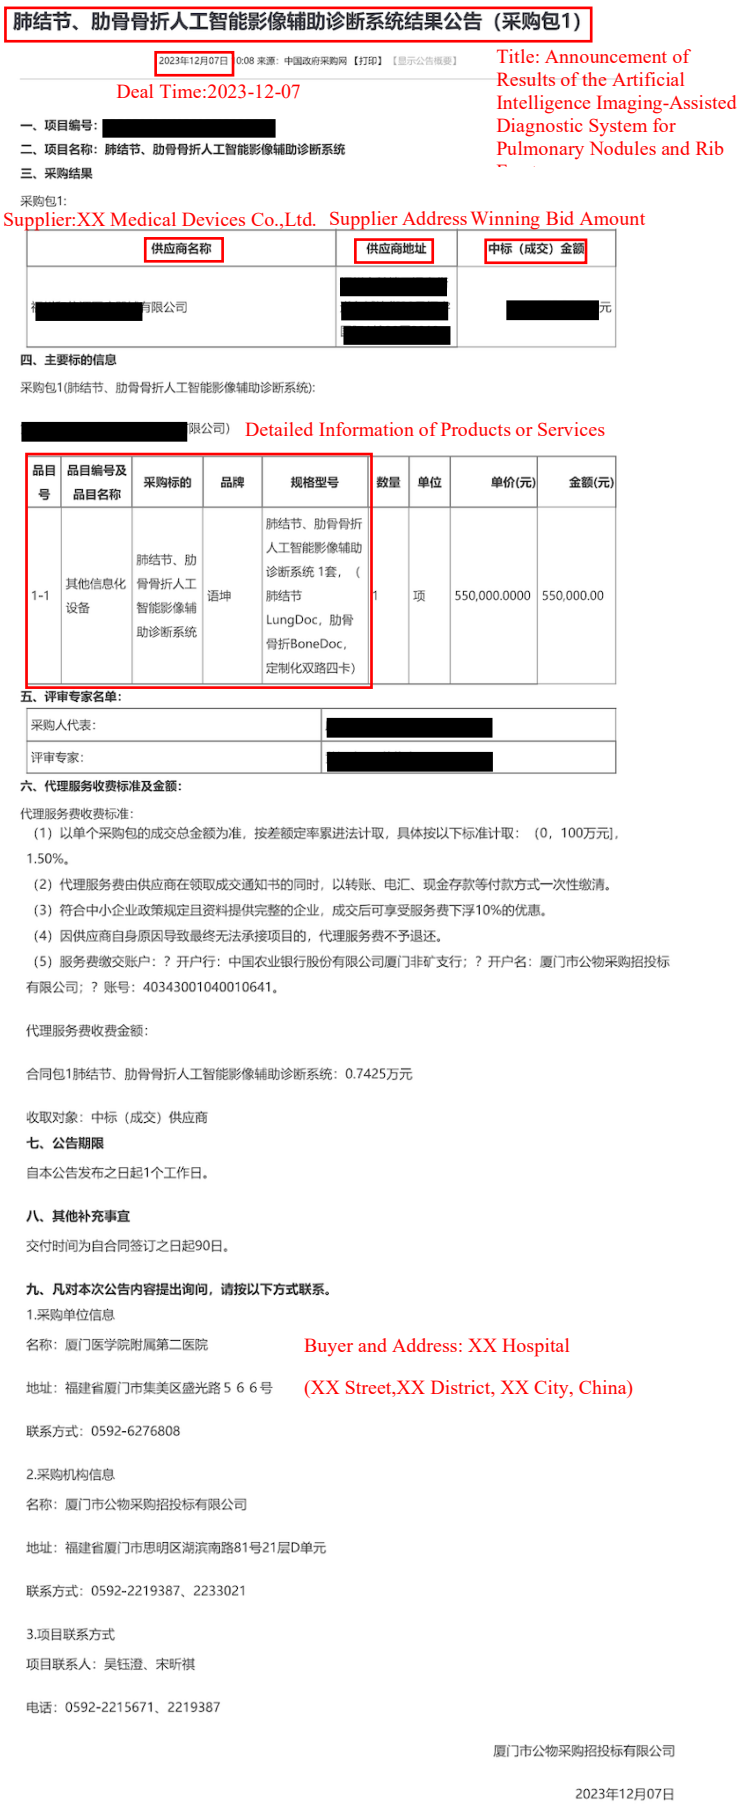


Figure A-6. Construction of the AI Procurement Sample

A sample of a healthcare AI procurement contract is shown in Figure A-7. Specifically, we operationalize hospital AI investment as the count of awarded AI procurement contracts in each city-quarter, rather than their monetary value. This choice reflects a structural feature of healthcare AI procurement: AI components are often embedded within larger procurement bundles that may also include non-AI hardware, system integration, training, and multi-year maintenance services. For these bundled contracts, a value-based measure would attribute the full bundled price to AI investment, mechanically overstating it. Counting contracts, by contrast, treats each AI-containing procurement as a discrete investment decision, independent of bundle composition.



肺结节、肋骨骨折人工智能影像辅助诊断系统结果公告（采购包1）

2023年12月07日 0:08 来源：中国政府采购网 【打印】 【显示公告概要】

Title: Announcement of Results of the Artificial Intelligence Imaging-Assisted Diagnostic System for Pulmonary Nodules and Rib

Deal Time:2023-12-07

一、项目编号：[redacted]

二、项目名称：肺结节、肋骨骨折人工智能影像辅助诊断系统

三、采购结果

采购包1:

Supplier:XX Medical Devices Co.,Ltd. Supplier Address Winning Bid Amount

供应商名称	供应商地址	中标（成交）金额
[redacted]有限公司	[redacted]	[redacted]元

四、主要标的信息

采购包1(肺结节、肋骨骨折人工智能影像辅助诊断系统):

[redacted]限公司) Detailed Information of Products or Services

品目号	品目编号及品目名称	采购标的	品牌	规格型号	数量	单位	单价(元)	金额(元)
1-1	其他信息化设备	肺结节、肋骨骨折人工智能影像辅助诊断系统	语坤	肺结节、肋骨骨折人工智能影像辅助诊断系统 1套，（肺结节LungDoc，肋骨骨折BoneDoc，定制化双路四卡）	1	项	550,000.0000	550,000.00

五、评审专家名单：

采购人代表：	[redacted]
评审专家：	[redacted]

六、代理服务收费标准及金额：

代理服务费收费标准：

（1）以单个采购包的成交总金额为准，按差额定率累进法计取，具体按以下标准计取：（0，100万元]，1.50%。

（2）代理服务费由供应商在领取成交通知书的同时，以转账、电汇、现金存款等付款方式一次性缴清。

（3）符合中小企业政策规定且资料提供完整的企业，成交后可享受服务费下浮10%的优惠。

（4）因供应商自身原因导致最终无法承接项目的，代理服务费不予退还。

（5）服务费缴交账户：？ 开户行：中国农业银行股份有限公司厦门非矿支行；？ 开户名：厦门市公物采购招投标有限公司；？ 账号：40343001040010641。

代理服务费收费金额：

合同包1肺结节、肋骨骨折人工智能影像辅助诊断系统：0.7425万元

收取对象：中标（成交）供应商

七、公告期限

自本公告发布之日起1个工作日。

八、其他补充事宜

交付时间为自合同签订之日起90日。

九、凡对本次公告内容提出询问，请按以下方式联系。

1.采购单位信息

名称：厦门医学院附属第二医院 Buyer and Address: XX Hospital

地址：福建省厦门市集美区盛光路５６６号 (XX Street,XX District, XX City, China)

联系方式：0592-6276808

2.采购机构信息

名称：厦门市公物采购招投标有限公司

地址：福建省厦门市思明区湖滨南路81号21层D单元

联系方式：0592-2219387、2233021

3.项目联系方式

项目联系人：吴钰澄、宋昕祺

电话：0592-2215671、2219387

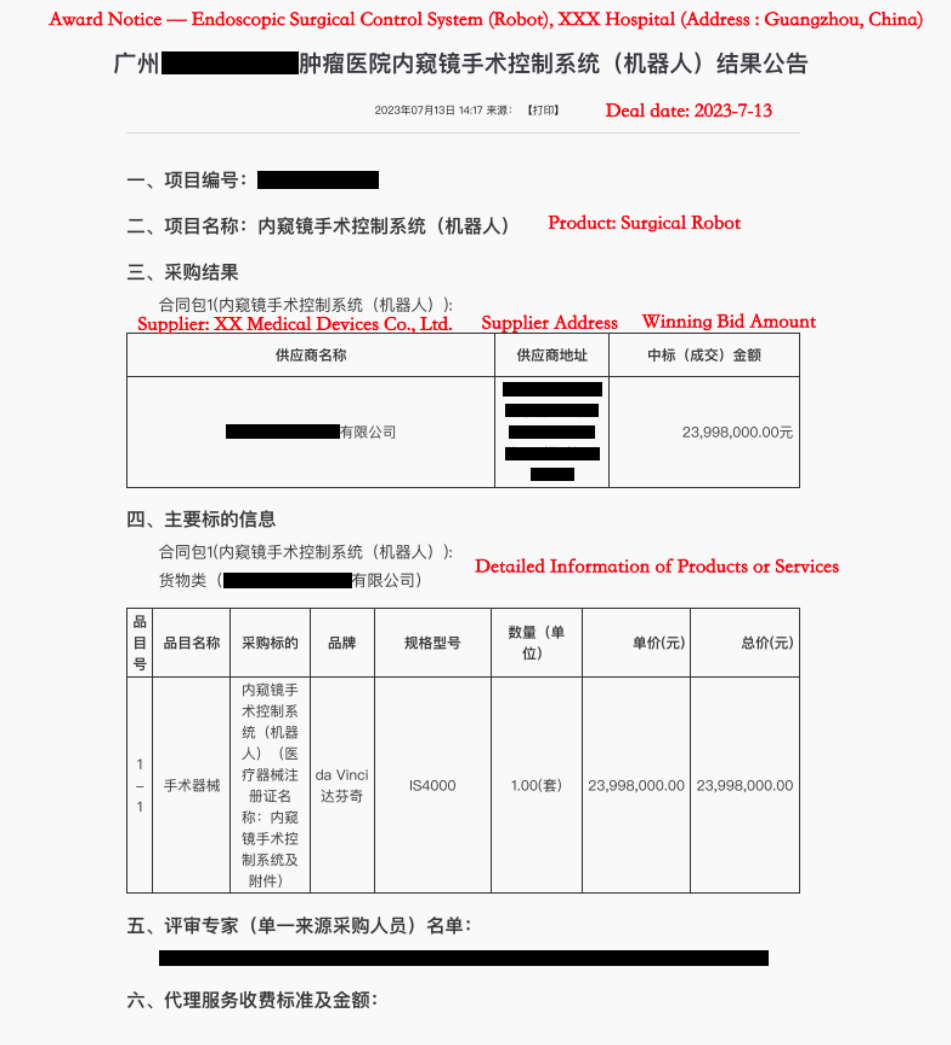
厦门市公物采购招投标有限公司

2023年12月07日

Figure A-7. Sample of a Healthcare AI Procurement Contracts

Notes: The contract record is drawn from the proprietary healthcare procurement database used in this study. The contract is titled "AI-assisted Diagnostic System for Pulmonary Nodules and Rib Fractures" and was awarded on December 7, 2023. The procuring entity is Second Affiliated Hospital of Xiamen Medical College, and the winning supplier is Fuzhou Hechuang Medical Equipment Co., Ltd. The total contract value is ¥ 550,000. The procured item is classified as an AI-assisted medical imaging diagnostic system, specifically including modules for pulmonary nodule detection (LungDoc) and rib fracture detection (BoneDoc). The contract specifies the brand, model, quantity (one system), unit price, and total transaction amount. Additional information is also disclosed, such as evaluation committee members, agency service fees, contract announcement period, delivery timeline (within 90 days after contract signing), and contact details of the procurement agency.

We illustrate this with two contracts from our sample, each representing a different form of AI bundling (*technical* and *application-scenario*). In the first case, a hospital procures a surgical robot system at ¥ 24.0 million, where AI software is central to the system's functionality but accounts for only a small share of the bundled price; the majority funds proprietary mechanical hardware and ancillary services (Figure A-8). In the second case, a hospital procures a multi-technology "Intelligent Hospital Construction" project at ¥ 31.0 million, in which AI is one of four parallel technology components alongside cloud computing, big data, and mobile applications, with much of the contract value funding non-AI components and ongoing implementation services (Figure A-9). Both contracts are correctly identified as AI-related under our keyword list, yet in both, the AI-specific share of the bundled price is substantially smaller than the contract total.



Award Notice — Endoscopic Surgical Control System (Robot), XXX Hospital (Address : Guangzhou, China)

广州[redacted]肿瘤医院内窥镜手术控制系统（机器人）结果公告

2023年07月13日 14:17 来源： 【打印】 Deal date: 2023-7-13

一、项目编号：[redacted]

二、项目名称：内窥镜手术控制系统（机器人） Product: Surgical Robot

三、采购结果

合同包1(内窥镜手术控制系统（机器人）):

Supplier: XX Medical Devices Co., Ltd. Supplier Address Winning Bid Amount

供应商名称	供应商地址	中标（成交）金额
[redacted]有限公司	[redacted]	23,998,000.00元

四、主要标的信息

合同包1(内窥镜手术控制系统（机器人）):

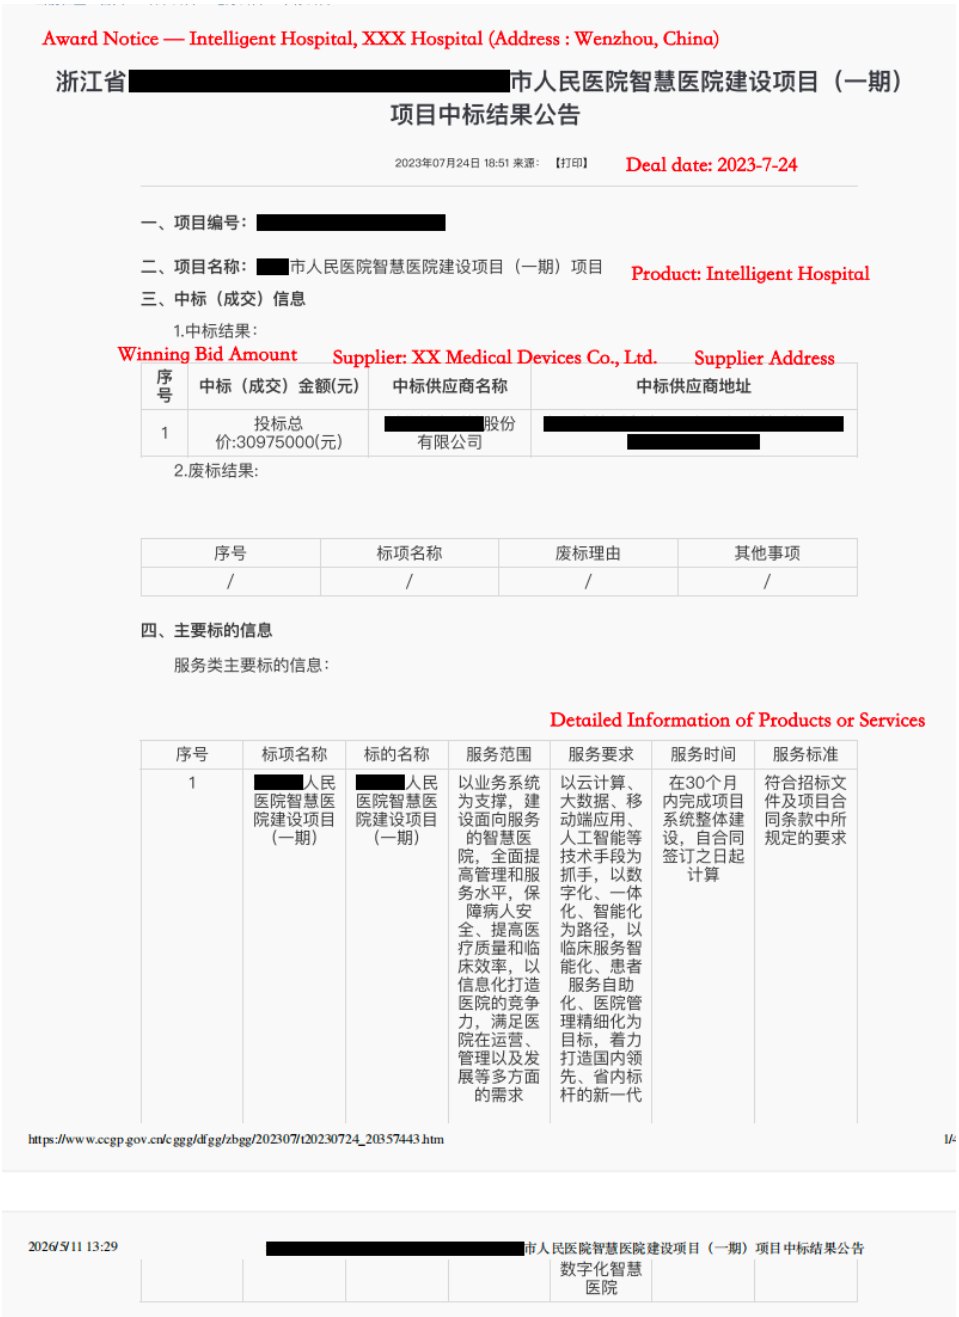
货物类（[redacted]有限公司） Detailed Information of Products or Services

品目号	品目名称	采购标的	品牌	规格型号	数量（单位）	单价(元)	总价(元)
1-1	手术器械	内窥镜手术控制系统（机器人）（医疗器械注册证名称：内窥镜手术控制系统及附件）	da Vinci 达芬奇	IS4000	1.00(套)	23,998,000.00	23,998,000.00

五、评审专家（单一来源采购人员）名单：[redacted]

六、代理服务收费标准及金额：

Figure A-8. An Awarded Healthcare AI Procurement Contract Matched via the Technical Dimension



Award Notice — Intelligent Hospital, XXX Hospital (Address : Wenzhou, China)

浙江省[redacted]市人民医院智慧医院建设项目（一期）项目中标结果公告

2023年07月24日 18:51 来源： 【打印】 Deal date: 2023-7-24

一、项目编号：[redacted]

二、项目名称：[redacted]市人民医院智慧医院建设项目（一期）项目 Product: Intelligent Hospital

三、中标（成交）信息

1.中标结果：

Winning Bid Amount Supplier: XX Medical Devices Co., Ltd. Supplier Address

序号	中标（成交）金额(元)	中标供应商名称	中标供应商地址
1	投标总价:30975000(元)	[redacted]股份有限公司	[redacted]

2.废标结果:

序号	标项名称	废标理由	其他事项
/	/	/	/

四、主要标的信息

服务类主要标的信息：

Detailed Information of Products or Services

序号	标项名称	标的名称	服务范围	服务要求	服务时间	服务标准
1	[redacted]人民医院智慧医院建设项目（一期）	[redacted]人民医院智慧医院建设项目（一期）	以业务系统为支撑，建设面向服务的智慧医院，全面提高管理和服务水平，保障病人安全、提高医疗质量和临床效率，以信息化打造医院的竞争力，满足医院在运营、管理以及发展等多方面的需求	以云计算、大数据、移动端应用、人工智能等技术手段为抓手，以数字化、一体化、智能化为路径，以临床服务智能化、患者服务自助化、医院管理精细化为目标，着力打造国内领先、省内标杆的新一代数字化智慧医院	在30个月内完成项目系统整体建设，自合同签订之日起计算	符合招标文件及项目合同条款中所规定的要求

https://www.ccgp.gov.cn/cggg/dfgg/zbgg/202307/t20230724_20357443.htm 1/4

2026/5/11 13:29 [redacted]市人民医院智慧医院建设项目（一期）项目中标结果公告

Figure A-9. An Awarded Healthcare AI Procurement Contract Matched via the Application-Scenario Dimension

For data collection, we obtained hospital procurement contract records from a leading commercial data vendor specializing in China's healthcare industry. Since 2012, this vendor has continuously aggregated procurement records from China's decentralized regional and institutional platforms into a proprietary database licensed to healthcare-industry clients for commercial use. Our team accesses this database under a long-term academic collaboration agreement permitting scientific use but prohibiting redistribution. We extracted all awarded contracts in which the procuring entity is a Chinese healthcare institution (hospitals, public health centers, and similar entities), yielding approximately 11 million contracts over 2012-2023. We restrict the analysis to awarded contracts rather than tender notices, since awarded contracts are legally binding records of executed transactions and more accurately capture realized investment than tender announcements, which can be canceled or fail to result in an award.

A.3.1. Construction of the Keyword List

We constructed the keyword list from four complementary sources. First, we drew on industry technology taxonomies, particularly the *Artificial Intelligence Index Report 2025* (Stanford University), for an internationally recognized list of AI subfields. Second, we incorporated official policy documents, specifically the *Guidelines for Artificial Intelligence Applications in the Health Industry* issued by the National Health Commission of China, which provide the authoritative regulatory taxonomy of AI use cases in Chinese hospitals. Third, we surveyed recent academic work on healthcare AI (e.g., Pham et al., 2024) for additional terminology used in the literature. Fourth, we conducted structured consultations with senior domain experts from hospitals, medical AI vendors, clinical practice, and academia, who reviewed the draft list, removed terms with low procurement-context relevance, and supplied Chinese-language equivalents.

We reconciled the four sources by removing duplicates, splitting overly broad terms, and consolidating synonymous variants, yielding a final list of 94 keywords organized along two complementary dimensions: a technical-attribute dimension and an application-scenario dimension.

From the technical-attribute dimension, the identified keywords are categorized into four sub-categories according to their functional hierarchy and implementation logic: (i) *Infrastructure*, referring to the underlying hardware and platforms that enable AI operations; (ii) *Models & Algorithms*, representing the core intelligence components that facilitate data processing, pattern recognition, and decision-making; (iii) *Perception & Cognition*, capturing AI's capacity to sense, interpret, and analyze external information; and (iv) *Execution & Interaction*, denoting the translation of cognitive outputs into actionable outcomes or user interactions. Collectively, these sub-categories reflect a bottom-up technological architecture, progressing from foundational computational support to intelligent processing, perceptual capability, and practical deployment.

From the application-scenario perspective, the identified keywords are classified into three functional domains based on the primary healthcare activities they serve: (i) *Clinical Workflow & Services*, encompassing AI applications embedded in core diagnostic and treatment processes, including imaging, prescription, consultation, and clinical decision support; (ii) *Hospital Operations & Management*, referring to AI systems that enhance administrative coordination, operational efficiency, and institutional governance; and (iii) *Extended Services*, comprising AI-enabled preventive, rehabilitative, chronic care, and community-based health services beyond core hospital treatment. These three domains represent parallel functional segments of the healthcare system rather than a hierarchical structure, distinguishing AI applications according to their organizational boundary and service orientation.

The two dimensions are designed to be complementary rather than mutually exclusive: a contract is classified as AI-related if its title or full text matches at least one keyword on either dimension. The complete lexicon is reported in Table A-1.

Table A-1. Keywords Related to Healthcare AI

Dimension	Sub-category	Keywords
Technical Attributes	Infrastructure	Digital Intelligence Technology, Digital Twin, Data Mining, Digitalization & Intelligence/Digital Intelligence Technology, Cloud Computing, Big Data
	Models & Algorithms	Artificial Intelligence (AI), Machine Learning (ML), Reinforcement Learning (RL), Supervised Learning, Deep Learning (DL), Neural Network, Large Language Model (LLM), Large Model, Generative Model, Graph Database, Clinical Knowledge Base, Expert System, Knowledge Engineering, Knowledge Graph
	Perception & Cognition	Computer Vision (CV), Image Processing, Machine Vision, Face Recognition, Image Recognition, Image Segmentation, Image Analysis, Natural Language Processing (NLP), Speech

		Recognition, Machine Translation, Voice Assistant, Text Mining
	Execution & Interaction	Robotics, Human-Computer Interaction(HCI)
Application Scenario	Clinical Workflow & Services	Intelligent Medical Care, Intelligent Pre-consultation, Intelligent Consultation, Virtual Doctor, Intelligent Medical Escort, Intelligent Outpatient Service, Assisted Diagnosis and Treatment, Intelligent Imaging, Automated Diagnosis, Intelligent Diagnosis, Assisted Diagnosis, Intelligent Assistance, Electronic Medical Record (EMR), Intelligent Medical Record, Intelligent Screening, Specialized Disease Database, Disease-Specific Intelligence, Disease Category Database, Intelligent Prescription Review, Assisted Medication, Intelligent Medication Management, Intelligent Prescription, Intelligent Pharmaceuticals, Intelligent Pharmacy, Intelligent Bed, Intelligent TCM, Intelligent TCM Pharmaceuticals, Intelligent Consultation, Clinical Decision Support System (CDSS), Intelligent Surgery, Clinical Intelligence, Intelligent Decision Making, Intelligent Follow-up, Intelligent Patient Monitoring, Intelligent Assistant, Virtual Assistant, Intelligent Herbal Paste Formulation, Intelligent Meridian System
	Hospital Operations & Management	Intelligent Management, Intelligent Health Insurance, Intelligent Hospital, Intelligent Medicine, Intelligent Healthcare, Intelligent Monitoring, Intelligent Testing/Detection, Intelligent Quality Control, Intelligent Prevention and Control, Intelligent Logistics
	Extended Services	Intelligent Nursing, Intelligent Prediction, Intelligent Prevention, Intelligent Elderly Care, Intelligent Scientific Research, Intelligent Health, Intelligent Rehabilitation, Intelligent Psychology, Intelligent Chronic Disease Management

Notes: All Chinese terms have been translated into English for presentation.

The construction of our Healthcare AI keyword lexicon follows three core principles to ensure data integrity and methodological rigor:

(1) *Broad inclusion with fine-grained filtering.* To mitigate the risk of omitting potential AI-related projects, we initially adopted a broad inclusion strategy. This involved incorporating terms with wider semantic extensions, such as “informatization,” “electronic medical records,” and “assisted diagnosis”, to serve as a preliminary boundary. While these terms may not be exclusively confined to the AI domain, they provided a comprehensive baseline. In subsequent stages, we used semantic disambiguation and manual textual identification to strictly isolate and confirm entries that represent actual AI.

(2) *Minimum identifiable semantic units*. To enhance retrieval comprehensiveness, our taxonomy is structured around “minimum identifiable semantic units.” For instance, by using “supervised learning” as a root unit, we systematically captured all associated derivatives including “unsupervised,” “self-supervised,” and “semi-supervised learning.” This approach ensures the consistency of our keyword identification across diverse procurement contexts.

(3) *Normalization of synonyms and syntactic variations.* To ensure conceptual consistency, we prioritized the normalization of synonymous expressions and syntactic variations. Although terms such as “intelligent health insurance” and “health insurance intelligence” were identified as separate strings in the raw data, they were consolidated into a single standardized construct during the final data synthesis.

A.3.2. Validation Procedure

Keyword-based classification of unstructured text is a well-established approach in economics and finance for identifying targeted content in large text corpora (Loughran & McDonald, 2011; Hassan et al., 2019), but it is subject to three known limitations. *False positives* arise when procurement contracts adopt prevailing technological buzzwords that do not correspond to genuine AI deployment. *False negatives* arise when AI-enabled medical systems are described under broader labels such as medical imaging or general healthcare service procurement. *Temporal mismatch* arises because the keyword list is partially derived from recent terminology, whereas our procurement sample extends back to 2012. To assess the magnitude of these errors, we conducted a manual validation exercise following the procedure of Hassan et al. (2019).

Sample. We drew a stratified random sample of 600 contracts from the universe of healthcare-related hospital procurement contracts. Stratification was based on two criteria: (i) sample period: three subperiods of 2012-2015, 2016-2019, and 2020-2023; and (ii) candidate classification status: contracts initially identified as AI-related by the keyword filter versus those not identified. This yielded six strata, with 100 contracts drawn per stratum, enabling estimation of both precision (from the keyword-identified strata) and recall (from the non-identified strata), and assessment of temporal stability.

Coding rules. Two research-team members independently coded each contract as 1 (AI-related) or 0 (non-AI) based on the contract's title and full text. A contract was coded as 1 if its text explicitly described an AI-enabled function (e.g., machine learning, computer vision, natural language processing) or named an AI application module (e.g., intelligent diagnosis system, clinical decision support, AI-assisted image reading). Contracts describing only generic IT infrastructure, basic information systems, or hardware without algorithmic content were coded as 0. The two coders worked blind to each other's labels.

Reconciliation. We compared the two coders' labels and resolved disagreements through joint discussion; any remaining unresolved cases were adjudicated by a third senior research-team member. The reconciled labels serve as the ground-truth benchmark.

Validation results. We report two sets of validation metrics. First, the two human coders reached a Cohen's Kappa of 0.85 on the 600 validation contracts, indicating high inter-coder reliability. Second, against the reconciled human-coded benchmark, the keyword-based classifier achieves precision of 91.5%, recall of 88.7%, and an F1 score of 89.3%. Subperiod analysis confirms stable performance across the sample window: precision and recall both fall within a narrow 86%-94% range across the three subperiods, indicating that the lexicon performs comparably for contracts written in early-, mid-, and late-period vocabularies, and that temporal mismatch does not materially degrade classification accuracy.

Supplementary LLM-based check. As an additional check, we generated independent labels for the 600 validation contracts using a large language model with a structured prompt instructing the model to classify each contract using the same coding rules described above. The LLM-generated labels agreed with the reconciled human labels on 95.6% of contracts, providing further evidence that the human-coded benchmark is reliable and that the keyword-based classifier is well calibrated.

A.3.3. Classification into Diagnostic and Operational Categories

After identifying 44,786 AI-related contracts through the keyword approach, we further classify each contract as diagnostic or operational: (i) a contract is classified as diagnostic if its primary application-scenario match is a clinical/diagnostic term (e.g., intelligent diagnosis, clinical decision support system); (ii) a contract is classified as operational if its primary application-scenario match is a hospital-operations or service-extension term (e.g., intelligent management, intelligent health insurance). This procedure yields 17,897 diagnostic contracts (40.0%) and 26,889 operational contracts (60.0%), summing to the 44,786 AI-related contracts identified above.

We aggregate the validated contracts to the city-quarter level based on the procuring hospital's registered city and the contract's announcement date, yielding three dependent variables: *Total Healthcare AI Investment*, *Diagnostic Healthcare AI Investment*, and *Operational Healthcare AI Investment*. The final balanced panel contains 12,384 city-quarter observations across 258 prefecture-level cities and 48 quarters (2012Q1–2023Q4).

Appendix B: Robustness Checks for the Baseline Estimate

B.1. Altering Covariate Sets

To examine whether our estimates are sensitive to the inclusion of different sets of covariates, we conduct a series of robustness checks based on alternative model specifications. While our baseline specification already incorporates a comprehensive set of city-level covariates and year-quarter fixed effects, we further assess the stability of our findings through the following variations:

(1) drop all time-varying covariates while keeping the city and quarter fixed effects;

(2) maintain all covariates in the baseline model and add interactions between all covariates and the year-quarter fixed effects. This specification allows the effects of city-level covariates to vary flexibly over time, thereby absorbing potential differential time trends correlated with OGDP launch;

(3) maintain all covariates in the baseline model and additionally absorb city-specific linear and quadratic time trends;

(4) additionally control for the number of participants in the basic medical insurance scheme for employees. This variable serves as a proxy for local healthcare purchasing power and demand, further mitigating omitted-variable bias related to the healthcare financing environment.

Table B-1 reports the regression results from these alternative specifications, and Figure B-1 visualizes the corresponding event study estimates. Across all variations, the estimated coefficients remain statistically significant and quantitatively stable, exhibiting minimal deviation from the baseline results. This consistency across progressively stringent specifications provides evidence that our findings are not sensitive to covariate selection or specific modeling assumptions.

Table B-1. Replication of Table 2 with Different Covariates

DV:	Total Healthcare AI Investment			
	(1)	(2)	(3)	(4)
OGDP	0.380*** (0.059)	0.198*** (0.041)	0.172*** (0.042)	0.294*** (0.052)
Controls	No	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
# Observations	12,384	12,384	12,384	12,384
Adjusted R^2	0.672	0.738	0.758	0.688

Notes: All columns include city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

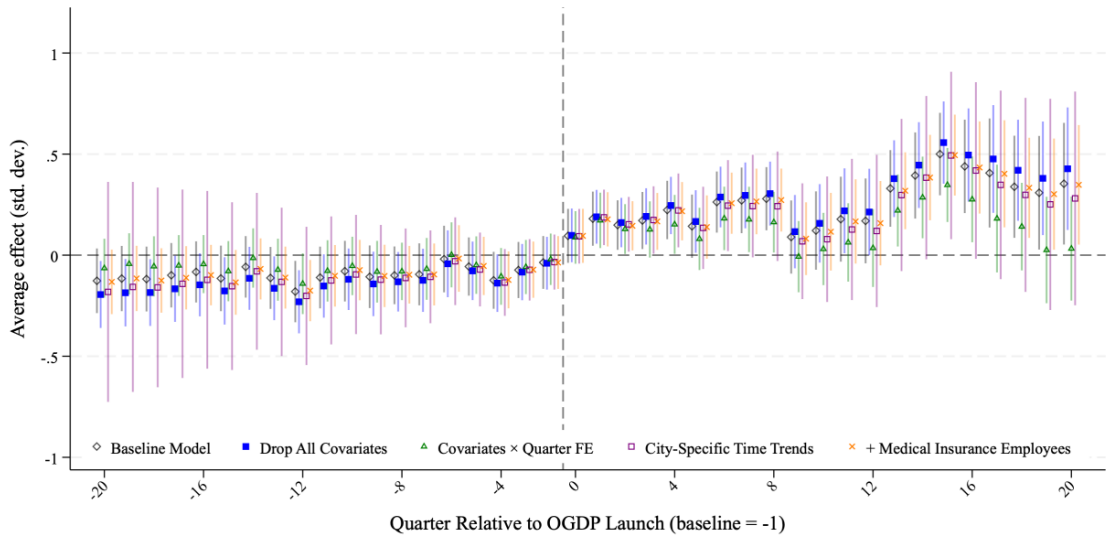


Figure B-1. Event Study Estimates from Specifications with Different Covariates

B.2. Propensity Score Matching

To assess the comparability of treated and control cities, we followed established practices in information systems and economics (Rosenbaum & Rubin, 1983; Heckman et al., 1997; Mithas & Krishnan, 2009) and employed propensity score matching (PSM) with a kernel matching algorithm. This approach mitigates concerns regarding potential unobserved heterogeneity between treated and control cities.

In the matching process, we incorporated a comprehensive set of city-level characteristics: the natural logarithms of resident population, GDP per capita, the number of higher-education institutions, the number of digital patents, the number of doctors, government R&D spending, and the number of hospitals. GDP per capita proxies for the overall economic environment, capturing local wealth and market sophistication; wealthier cities are better positioned to bear the research and implementation costs of advanced technologies such as healthcare AI. Resident population reflects the scale of the local labor market and potential healthcare demand, providing the human-capital pool and knowledge spillovers needed for AI adoption. The number of higher-education institutions measures the availability of skilled human capital, ensuring that investment disparities are not driven solely by the supply of STEM professionals. Digital patent counts capture existing technological capabilities and absorptive capacity, indicating a city's readiness for innovation commercialization. Government R&D spending signals institutional commitment to innovation, creating a favorable environment for AI projects. Finally, the number of hospitals serves as a proxy for healthcare infrastructure capacity; stronger institutional networks facilitate the clinical translation of AI investments.

We estimated propensity scores using a logit model based on these covariates and then applied kernel matching. To ensure match quality, we imposed a common-support condition, excluding any treated city without a suitable control. Post-matching diagnostics show that the standardized mean differences of the covariates between the treated and control groups decrease substantially, indicating successful covariate balance.

Figure B-2 illustrates the distribution of matching variables before and after PSM, demonstrating the improved balance achieved through the procedure. Table B-2 presents the re-estimated results of Table 2 using the matched sample, confirming the robustness of our baseline estimates.

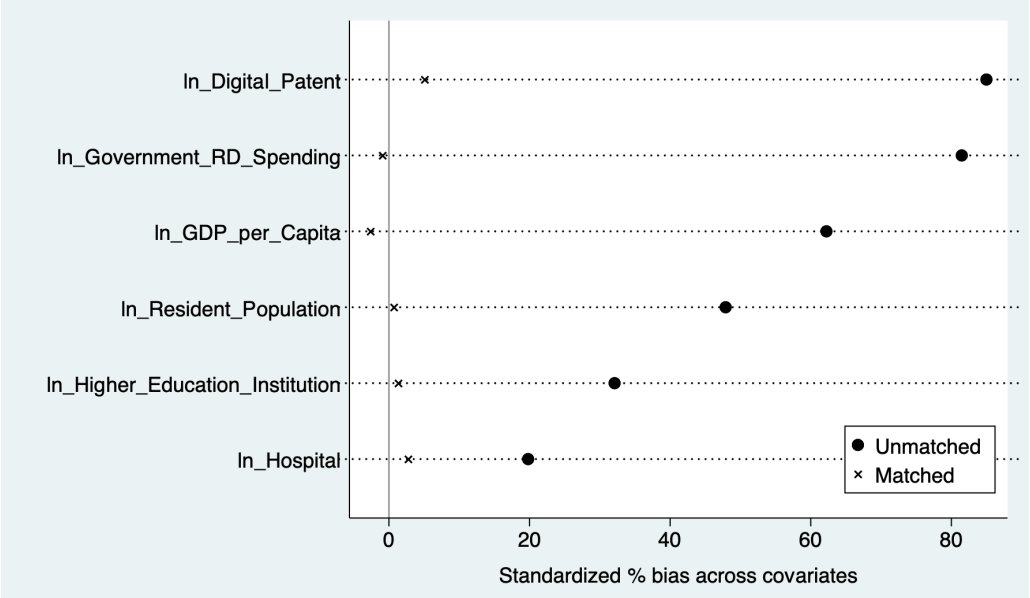


Figure B-2. Changes in Matching Variables Before and After PSM

Table B-2. Replication of Table 2 with Propensity Score Matching

DV:	Total Healthcare AI Investment
OGDP	0.303*** (0.054)
Controls	Yes
City FE	Yes
Quarter FE	Yes
# Observations	12,204
Adjusted R^2	0.666

Notes: The analysis includes covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.3. Coarsened Exact Matching

Given that cities were not randomly assigned to the treatment and control groups, unobserved heterogeneity between the two groups could bias our baseline estimates. To mitigate this selection concern and improve comparability between the groups, we employ Coarsened Exact Matching (CEM). Unlike PSM, which may require iterative manual inspection and re-estimation to achieve balance, CEM is a nonparametric approach that allows researchers to bound imbalance a priori. This method has been widely adopted in recent DiD analyses in IS and economics (e.g., Wang et al., 2022).

We matched the treatment and control groups using a comprehensive set of pre-treatment covariates: the natural logarithms of GDP per capita, the number of higher-education institutions, the number of digital patents, government R&D spending, the number of doctors, and the number of hospitals. The pre-matching multivariate imbalance measure (multivariate L1 statistic) was 0.974, indicating a substantial distributional divergence between the treatment and control groups. To define the matching strata, we use the Shimazaki-Shinomoto rule (autocuts(ss)) to coarsen these continuous covariates into bins. This procedure discretizes the variables and performs exact matching within these coarsened strata. Matching substantially improved comparability between the two groups: the L1 statistic decreased significantly from 0.974 to 0.469, representing a 51.8% reduction in overall imbalance. The resulting matched sample retains 384 out of 2,372 treated city-quarters and 416 out of 10,012 control city-quarters.

Finally, we re-estimated our baseline regression on this matched sample using the CEM-generated weights. As reported in Table B-3, the coefficient for OGDP remains positive and statistically significant. This consistency demonstrates that our main findings are robust to distributional differences between the treatment and control groups.

Table B-3. Replication of Table 2 with Coarsened Exact Matching

DV:	Total Healthcare AI Investment
OGDP	0.249*** (0.085)
Controls	Yes
City FE	Yes
Quarter FE	Yes
# Observations	792
Adjusted R^2	0.727

Notes: The analysis includes covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.4. Adjusted by Weights

To address potential concerns arising from regional heterogeneity in economic development and healthcare capacity, we re-estimated our baseline models using weighted least squares (WLS) instead of the unweighted specifications used in the main analyses. This approach assigns proportionally greater influence to cities with larger economies or more extensive healthcare resources. Specifically, we applied two distinct weighting schemes: (1) weights based on city-level GDP per capita, accounting for varying degrees of economic density; and (2) weights based on the total number of hospitals, adjusting for the concentration of medical resources.

Table B-4 reports the results. The estimated coefficients for OGDP remain positive and statistically significant across both specifications. The stability of these estimates across alternative weighting schemes suggests that our baseline findings are not driven by particular subsets of cities.

Table B-4. Replication of Table 2 with GDP- and Hospital-Based Weights

DV:	Total Healthcare AI Investment	
	Weights by GDP per capita	Weights by the number of hospitals
	(1)	(2)
OGDP	0.297*** (0.053)	0.297*** (0.055)
Controls	Yes	Yes
City FE	Yes	Yes
Quarter FE	Yes	Yes
# Observations	12,384	12,384
Adjusted R^2	0.690	0.704

Notes: All columns include covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.5. Controlling for Outlier Cities

To address concerns that our results may be driven by the distinctive characteristics of China's most developed metropolises, we excluded the six largest and structurally distinct cities: the four municipalities directly under the central government (Beijing, Shanghai, Tianjin, and Chongqing) plus the two remaining first-tier cities (Guangzhou and Shenzhen). These cities differ from the rest of the sample in economic agglomeration, digital infrastructure, policy capacity, and innovation ecosystems, could disproportionately influence the estimated effects.

As shown in Table B-5, re-estimating the model on this restricted sample yields results consistent with our baseline findings. The persistence of these estimates suggests that our conclusions are not an artifact of top-tier city influence, but reflect the broader urban landscape in China.

Table B-5. Replication of Table 2 After Excluding Outlier Cities

DV:	Total Healthcare AI Investment
OGDP	0.290*** (0.053)
Controls	Yes
City FE	Yes
Quarter FE	Yes
# Observations	12,096
Adjusted R^2	0.665

Notes: The analysis includes covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.6. Bacon Decomposition

A concern with TWFE estimation in staggered DiD designs is that treatment effects may vary over time. When treatment effects are heterogeneous, the use of already-treated units as controls for later-treated units, or so-called forbidden comparisons can introduce negative weighting bias, potentially yielding estimates that diverge substantially from the true average treatment effect on the treated (ATT). To diagnose this concern, we employ the decomposition method of Goodman-Bacon (2021). This procedure decomposes the aggregate TWFE coefficient into its constituent 2 × 2 DiD comparisons, enabling examination of the distribution of weights across different control groups. Because the Bacon decomposition applies only to TWFE specifications without time-varying covariates, we implement it on the no-covariate version of our baseline model (Column 1 of Table B-1, aggregate estimate = 0.380).

Figure B-3 presents the decomposition results. Our overall TWFE estimate is predominantly driven by the "clean" comparison between treated and never-treated units, which accounts for the dominant share of the total weight and yields an estimate aligned closely with our baseline findings in both magnitude and direction. The potentially problematic comparisons involving variation in treatment timing (e.g., later-treated units against earlier-treated units) carry negligible weight. The within-group component (represented by green crosses in the figure) contributes to the overall variation but does not drive the primary results. We therefore find no evidence that our estimates are artifacts of time-varying effects or negative weighting bias, confirming the robustness of the OGDP effect.

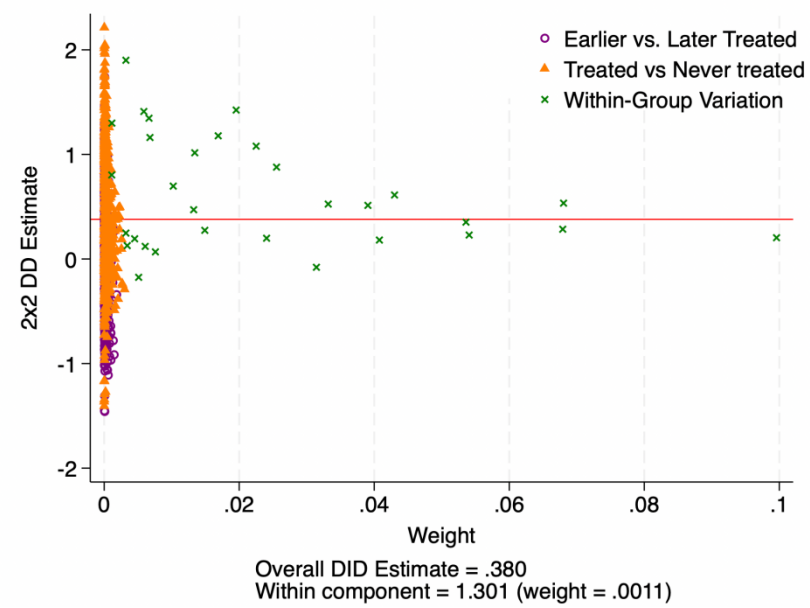


Figure B-3. Bacon Decomposition of the TWFE Estimator

B.7. Alternative Operationalizations of Independent Variable

As an additional robustness check, we construct an alternative measure of OGDP launch timing to account for within-quarter implementation lags. Specifically, for cases where the official launch of the OGDP occurred in the last month of a quarter, we reclassified the effective start date to the subsequent quarter, on the grounds that policies promulgated late in a quarter often lack sufficient time to generate observable effects within the same period.

We re-estimated our baseline models using this adjusted timing. As shown in Table B-6, the results remain qualitatively unchanged, consistent with our baseline estimates in terms of both coefficient magnitude and statistical significance.

Table B-6. Replication of Table 2 with Alternative Operationalization of IV

DV:	Total Healthcare AI Investment	
	TWFE-DID	CS-DID
	(1)	(2)
OGDP	0.296*** (0.053)	0.216*** (0.076)
Controls	Yes	Yes
City FE	Yes	Yes
Quarter FE	Yes	Yes
# Observations	12,384	12,384

Notes: All columns include covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. The TWFE-DiD model yields an adjusted R^2 of 0.685.We first implement a TWFE-DiD estimation, followed by a robustness check using Callaway and Sant'Anna (2021), with both approaches yielding consistent results. *** p<0.01, ** p<0.05, * p<0.1.

B.8. Alternative Operationalizations of Dependent Variable

To assess the robustness of our findings to alternative measures of the dependent variable, we adopt two conservative operationalizations.

The first targets the identification of AI-related contracts. Our baseline measure captures a broad spectrum of AI-related contracts, including those where AI integration is functional but not explicitly labeled. For this robustness check, we restrict the sample to contracts that explicitly mention AI terminology, excluding projects where AI functions are likely present but embedded within broader IT infrastructure or maintenance descriptions. For instance, contracts related to electronic medical records (EMRs) frequently mandate advanced data-processing tasks embedded within maintenance clauses. Specific examples include "data quality control scoring subsystems" (which automate the monitoring of data consistency and accuracy), "automated error detection" for system stability, and "interoperability data cleaning" (e.g., standardizing and converting millions of patient records for national platform integration). While domain experts recognize that these modules typically employ algorithmic or AI-driven logic to process large-scale datasets and identify anomalies, they are frequently labeled generically without explicit AI tags. We exclude such ambiguous cases, reducing the number of AI-related contracts from 44,786 to 39,340.

The second operationalization replaces the count-based dependent variable with the log monetary value of awarded healthcare AI procurement contracts in each city-quarter.

We re-estimate our models under both operationalizations. As shown in Table B-7, the OGDP coefficient retains its sign and statistical significance in both columns.

Table B-7. Replication of Table 2 with Alternative Operationalization of DV

DV:	Total Healthcare AI Investment	
	Number of AI contracts from the explicitly AI-labeled sample	Monetary value of AI contracts from the baseline sample
	(1)	(2)
OGDP	0.353*** (0.051)	0.528*** (0.140)
Controls	Yes	Yes
City FE	Yes	Yes
Quarter FE	Yes	Yes
# Observations	12,384	11,720
Adjusted R^2	0.688	0.604

Notes: All columns include covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.9. Inverse Hyperbolic Sine Transformation

In our baseline analysis, the dependent variable is logarithmically transformed to address skewness. A potential concern, however, is the handling of zero values in healthcare AI investment data (i.e., city-quarters with no procurement contracts). Standard logarithmic transformations require dropping observations with zero values or adding an arbitrary constant (e.g., ln(Y+1)), which may introduce bias or reduce statistical power.

To address this, we employ the inverse hyperbolic sine (IHS) transformation as a robustness check. The IHS transformation is defined as $asinh(y) = ln(y + \sqrt{y^2 + 1})$. This transformation approximates the natural logarithm for large values while remaining defined at zero (IHS(0) = 0), allowing us to retain the full sample including observations with zero investment (Burbidge et al., 1988; Bellemare & Wichman, 2020).

We re-estimate the baseline two-way fixed effects (TWFE) model using the IHS-transformed volume of healthcare AI investment as the dependent variable. Table B-8 presents the results. The estimated coefficient for OGDP is 0.345 and remains statistically significant at the 1% level. This suggests that the implementation of OGDP is associated with a significant increase in healthcare AI investment. This result indicates that the OGDP effect is not an artifact of the log(Y+1) transformation or the handling of zero-valued observations.

Table B-8. Replication of Table 2 with IHS Transformation

DV:	Total Healthcare AI Investment
OGDP	0.345*** (0.063)
Controls	Yes
City FE	Yes
Quarter FE	Yes
# Observation	12,384
Adjusted R^2	0.683

Notes: The analysis includes covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.10. Alternative Model Specifications

We further assess the sensitivity of our estimates to model specification choices. Because OLS may yield biased estimates in the presence of heteroscedasticity or when the dependent variable contains a substantial number of zeros, we employ the Poisson pseudo-maximum likelihood (PPML) as a robust alternative. Table B-9 reports the results. The estimates remain consistent in both magnitude and significance, providing further evidence that our main conclusions are robust to alternative distributional assumptions.

Table B-9. Replication of Table 2 with Poisson Pseudo-Maximum Likelihood

DV:	Total Healthcare AI Investment
OGDP	0.199*** (0.076)
Controls	Yes
City FE	Yes
Quarter FE	Yes
# Observations	12,336
Log pseudolikelihood	-20587.67415

Notes: The analysis includes covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.11. Alternative Estimation of Standard Errors (Adjusted by Clustering)

To account for potential spatial correlation and interdependence among error terms within provinces, we re-estimated our baseline models with standard errors clustered at the province level (rather than the city level). This specification relaxes the assumption of independence across cities within the same province, allowing for arbitrary correlation of residuals among them.

As shown in Table B-10, the statistical significance of our key estimate remains unaffected. The consistency confirms that our statistical inference is robust to alternative assumptions regarding the structure of the error terms.

Table B-10. Replication of Table 2 with Standard Errors Clustered at the Province Level

DV:	Total Healthcare AI Investment
OGDP	0.297*** (0.079)
Controls	Yes
City FE	Yes
Quarter FE	Yes
# Observations	12,384
Adjusted R^2	0.686

Notes: The analysis includes covariates as well as city and year-quarter fixed effects. Robust standard errors (clustered at the province level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.12. Controlling for Other Shocks

We further examine the robustness of our results by excluding the COVID-19 pandemic period. The pandemic represented a severe exogenous shock that disrupted the healthcare sector, redirecting financial resources toward emergency response rather than strategic investment. Such systemic irregularities could introduce structural breaks or estimation bias if included in the long-run analysis.

To address this, we restrict our sample window to the pre-pandemic period (2012–2019). The results, reported in Table B-11, remain consistent with our baseline estimates in both magnitude and statistical significance. This confirms that our primary findings are not driven by the anomalous healthcare-spending dynamics induced by the global health crisis.

Table B-11. Replication of Table 2 Controlling for COVID-19

DV:	Total Healthcare AI Investment	
	TWFE-DID	CS-DID
	(1)	(2)
OGDP	0.292*** (0.068)	0.300*** (0.109)
Controls	Yes	Yes
City FE	Yes	Yes
Quarter FE	Yes	Yes
# Observations	8,256	8,256

Notes: All columns include covariates as well as city and quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. The TWFE-DiD model yields an adjusted R^2 of 0.541.We first implement a TWFE-DiD estimation, followed by a robustness check using Callaway and Sant'Anna (2021); both approaches yield consistent results. *** p<0.01, ** p<0.05, * p<0.1.

B.13. Placebo Test

To rule out the possibility that our findings are driven by spurious correlations or unobserved time-varying confounders, we conduct a placebo test following the random-permutation procedure widely adopted in the information systems and economics literatures (e.g., Chetty et al., 2009; La Ferrara et al., 2012).

We construct a counterfactual treatment variable, OGDP_placebo, by randomly assigning treatment status and launch timing to cities. In each iteration, we randomly select 134 cities from the full sample to serve as the pseudo-treatment group and randomly assign a hypothetical launch date to each (ensuring dates fall within the sample period, excluding the pre-treatment initialization phase). OGDP_placebo takes a value of 1 for the selected cities after their hypothetical launch dates, and 0 otherwise. We then re-estimate the baseline specification using OGDP_placebo in place of the actual policy variable. This process is repeated 500 times to generate a distribution of placebo coefficients. The logic of the test is straightforward: if the baseline OGDP effect reflects a genuine causal mechanism, the coefficients on OGDP_placebo should be centered around zero and statistically insignificant, because the pseudo-treatment has no actual policy impact on healthcare AI investment. Conversely, if the baseline result were driven by chance or unobserved time-varying confounders, the distribution of OGDP_placebo coefficients would resemble the baseline estimate.

Figure B-4 illustrates the distribution of the estimated coefficients and their corresponding *p*-values from the 500 simulations. The kernel density curve depicts the probability density of the placebo coefficients, while the red dots represent the *p*-values. The vertical solid line indicates the magnitude of our baseline coefficient (0.297), and the horizontal dashed line marks the conventional significance threshold at $p = 0.1$. The distribution of OGDP_placebo coefficients centers tightly around zero, and the vast majority of *p*-values lie above the 0.1 threshold. Our baseline estimate lies far in the right tail of the placebo distribution, clearly distinct from the placebo coefficients, indicating that the main results are not driven by chance or spurious correlation.

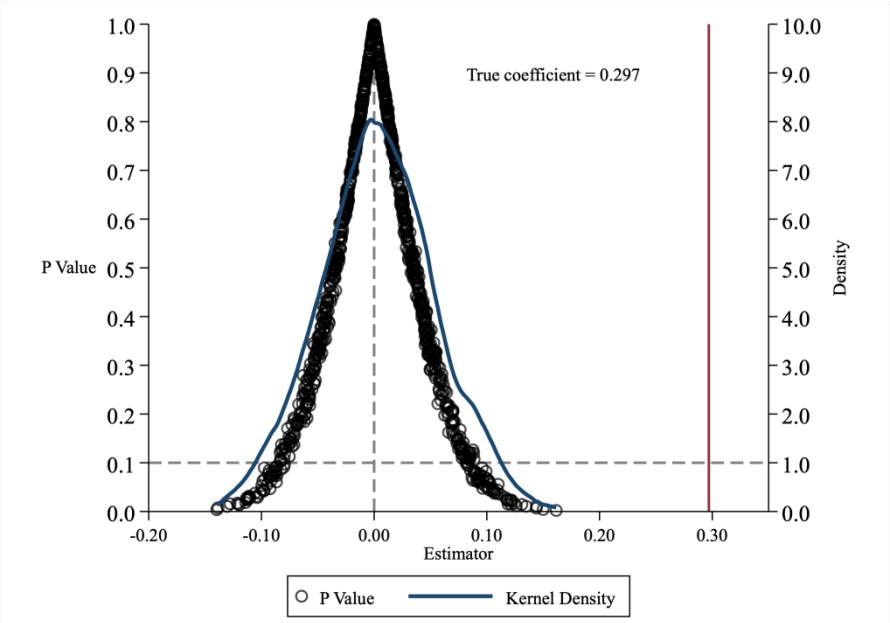


Figure B-4. Distribution of Regression Coefficients and p-Values from Placebo Test

Notes: The figure plots the kernel density of the coefficients obtained from 500 simulations where the treatment status and timing were randomly permuted. The red dots indicate the p-values for each simulation. The vertical solid line represents the benchmark OGDP coefficient (0.297), and the horizontal dashed line indicates the 10% significance level.

B.14. Generalized Synthetic Control

When time-varying confounders are present, estimates from a standard difference-in-differences (DID) approach can be biased. To address this concern, we employ the synthetic control framework, which constructs a counterfactual outcome trajectory for each treated city by reweighting a set of untreated control cities to match both pre-treatment covariates and outcome trends. This framework uses the pre-treatment period as the criterion for evaluating match quality, making the comparison between treated cities and their counterfactual trajectories transparent (Abadie et al., 2010).

To accommodate city- and time fixed effects and to handle the unbalanced structure of our panel data, we adopt the generalized synthetic control (GSC) framework proposed by Xu (2017). Compared with the conventional synthetic control method, the GSC framework allows for multiple treated units and staggered treatment timing, incorporates a two-way fixed-effects structure, and is well-suited for panel settings with heterogeneous policy adoption. In our context, this approach is appropriate because OGDP was implemented at different times across cities. We control for multiple time-varying covariates at the city level, including log-transformed GDP per capita, resident population, number of higher-education institutions, number of digital patents, government R&D spending, number of doctors, and number of hospitals.

Within the GSC framework, we estimate the model using the matrix completion (MC) estimator developed by Athey et al. (2021), which extends the synthetic control logic by directly imputing counterfactual outcomes through a nuclear-norm-regularized matrix completion procedure. Compared with the interactive fixed-effects (IFE) estimator in Xu (2017) , the MC estimator is more robust to high-dimensional unobserved heterogeneity and imposes weaker functional-form assumptions on the latent factor structure. The optimal regularization hyperparameter (λ) is selected via cross-validation from the set {0.001, 0.002, 0.003, 0.004, 0.005}. Standard errors and 95% confidence intervals are constructed using a nonparametric bootstrap procedure with 500 replications.

The results are presented in Table B-12 and Figure B-5. The estimated average treatment effect on the treated (ATT) is 0.369 and is statistically significant at the 1% level ($p < 0.01$). This indicates that, even after constructing a counterfactual that flexibly accounts for time-varying unobserved factors, the OGDP effect on healthcare AI investment remains positive and robust, consistent with our baseline findings.

Table B-12. Replication of Table 2 with Synthetic Control Estimation

DV	Total Healthcare AI Investment
ATT	0.369*** (0.070)
Controls	Yes
City FE	Yes
Quarter FE	Yes
# Observations	12,384

Notes: ATT is estimated using the Matrix Completion (MC) estimator (Athey et al., 2021) within the Generalized Synthetic Control framework (Xu, 2017). The regularization parameter λ is selected via cross-validation from {0.001, 0.002, 0.003, 0.004, 0.005}. Standard errors (in parentheses) and confidence intervals are constructed using a non-parametric bootstrap procedure with 500 replications. *** p<0.01, ** p<0.05, * p<0.1.

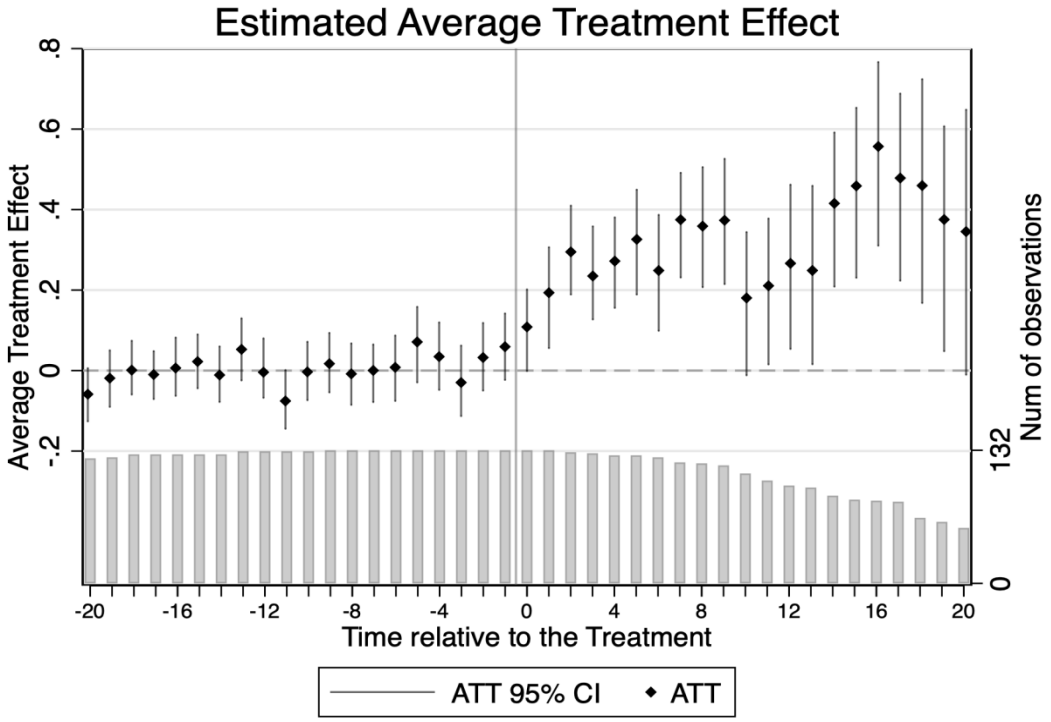


Figure B-5. Estimated Dynamic Effects of OGDP on Healthcare AI Investment Using the Generalized Synthetic Control Method

Notes: The figure plots the estimated dynamic average treatment effects on the treated (ATT) of the OGDP policy over time, based on the matrix completion (MC) estimator (Athey et al., 2021) within the Generalized Synthetic Control framework (Xu, 2017). The solid line connects the point estimates of the ATT for each period. Vertical bars represent the 95% confidence intervals constructed using a non-parametric bootstrap procedure with 500 replications. The trajectory shows that the difference between treated cities and their counterfactual trajectories is close to zero prior to the OGDP launch (indicating a good pre-treatment fit) and diverges positively afterward.

Appendix C. Alternative Explanations

C.1. Concurrent AI and Data Policies

A potential concern is that OGDP launches may coincide with broader provincial policy initiatives, such as Big Data or AI development plans, and that the estimated OGDP effect may capture these contemporaneous interventions. To isolate the effect of OGDP from such broader policy initiatives, we construct a dataset of provincial-level Data Policies and AI Policies. We manually collected policy documents from the PKULaw (Beida Fabao) Database (https://www.pkulaw.com/), a comprehensive repository of Chinese legal and regulatory documents widely used in Chinese legal and economics research.

To ensure that our measure captures substantive strategic shifts rather than minor administrative adjustments, we applied two filtering criteria. First, we restricted our search to high-level documents issued directly by the Provincial People's Congress or the Provincial People's Government, excluding lower-level departmental notices that lack significant administrative force. Second, we focused exclusively on comprehensive policy frameworks (e.g., *Development Plan for Big Data Industry or Action Plan for New Generation Artificial Intelligence*), excluding narrow sector-specific notices that do not reflect a region-wide strategic shift. For each province, we identified the implementation date of the first major policy in each category (regardless of subsequent revisions) to capture the timing of the initial policy announcement. We then constructed two time-varying indicators defined at the province-quarter level: $Data\ Policy_{p,t}$ equals 1 if province p has implemented a provincial-level data policy by quarter t, and 0 otherwise; $AI\ Policy_{p,t}$ is defined analogously. Selected examples are reported in Table C-1.

Table C-1. Summary of Provincial Data and AI Policies (Selected Examples)

Province	Policy Category	Document Title	Issuing Authority	Implementation Date	Current Status
Beijing	Data Policy	Notice of the Beijing Municipal People's Government on Issuing the "Beijing Action Plan for Big Data and Cloud Computing Development (2016–2020)"	People's Government of Beijing Municipality	August 3, 2016	Valid
Jiangsu	Data Policy	Notice of the Jiangsu Provincial People's Government on Issuing the "Jiangsu Action Plan for Big Data Development"	People's Government of Jiangsu Province	August 19, 2016	Repealed / Expired
Zhejiang	AI Policy	Notice of the People's Government of Zhejiang Province on Issuing the "Zhejiang Development Plan for New Generation Artificial Intelligence"	People's Government of Zhejiang Province	December 4, 2017	Valid
Anhui	AI Policy	Notice of the People's Government of Anhui Province on Issuing the "Anhui Development Plan for New Generation Artificial Intelligence Industry (2018–2030)"	People's Government of Anhui Province	May 11, 2018	Valid

Table C-2 reports the estimates from three specifications: augmenting Equation (1) with Data Policy alone (Column 1), AI Policy alone (Column 2), and both indicators jointly (Column 3). The AI Policy coefficient is positive and significant (0.149, $p < 0.01$), implying a roughly 16% increase in healthcare AI investment associated with the introduction of a provincial AI strategy. This is consistent with the interpretation that explicit AI strategies provide additional regulatory clarity. The Data Policy coefficient is small and statistically insignificant. Importantly, the OGDP coefficient remains positive, stable, and highly significant across all three specifications, indicating that the OGDP effect is not absorbed by the most salient concurrent provincial policy initiatives.

Table C-2. Robustness of the OGDP Effect to Controlling for Concurrent AI and Data Policies

DV:	Total Healthcare AI Investment		
	Control Data Policy	Control AI Policy	Control Data & AI Policy
	(1)	(2)	(3)
OGDP	0.297*** (0.053)	0.292*** (0.053)	0.292*** (0.053)
Data Policy	-0.014 (0.038)		0.002 (0.038)
AI Policy		0.149*** (0.051)	0.149*** (0.052)
Controls	Yes	Yes	Yes
City FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
# Observations	12,384	12,384	12,384
Adjusted R^2	0.686	0.687	0.687

Notes: All columns include covariates as well as city and quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

C.2. Local Digitalization Trajectories

A potential concern is that OGDP launches may correlate with a city's broader digitalization trajectory, and that the estimated OGDP effect partly reflects this pre-existing digital capacity rather than the OGDP itself. To examine this concern, we use two proxies for city digitalization and conduct subsample analyses comparing high- and low-digitalization cities.

The first proxy is the city's annual count of digital patents, obtained from the China National Intellectual Property Administration (CNIPA), China's authoritative registry for innovation output in China. We identify digital patents based on the official *Key Digital Technology Patent Classification System* issued by CNIPA,[1] which provides a standardized list of International Patent Classification (IPC) codes corresponding to digital technologies. The second proxy is the city's annual telecommunications revenue, obtained from the *China City Statistical Yearbook*, which reports the scale of local information infrastructure across Chinese cities. For each proxy, we classify cities into high- and low-digitalization groups based on the sample-period median of the proxy.

Table C-3 reports the subsample estimates. Columns 1 and 2 present results using digital patents, and Columns 3 and 4 present results using telecommunications revenue. Chi-square tests for the equality of OGDP coefficients between the high- and low-digitalization groups yield 1.08 ($p > 0.10$) for digital patents and 0.03 ($p > 0.10$) for telecommunications revenue. We therefore cannot reject the null hypothesis that the OGDP effect is uniform across digitalization levels. This pattern suggests that the OGDP effect is not driven by cities' pre-existing digital capacity, ruling out the alternative interpretation that the estimated effect simply reflects a more general local digitalization trajectory.

Table C-3. Subsample Analyses by Local Digitalization Trajectories

Alternative Explanations	Empirical Proxies	Type	TWFE-DID	# Observation	Adjusted R^2
Local Digitalization Trajectories	digital patents (year-level)	High	0.186*** (0.061)	6,192	0.710
		Low	0.285*** (0.073)	6,192	0.519
		Chi-Square	1.08		
	telecommunications revenue (year-level)	High	0.228*** (0.068)	6,192	0.744
		Low	0.212*** (0.060)	6,192	0.546
		Chi-Square	0.03		

Notes: Cities are split at the sample-period median of each proxy. All columns include covariates as well as city and quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

[1] https://www.cnipa.gov.cn/attach/0/d32119ae1faa4fbf9e308824862f87dd.pdf

C.3. Fiscal Capacity and Spending Priorities

A potential concern is that fiscally stronger cities may be both more likely to launch OGDPs and more capable of supporting healthcare AI investment through public funding or subsidies. If so, the estimated OGDP effect may partly reflect local fiscal capacity rather than the OGDP itself. To examine this concern, we construct two time-varying measures of local fiscal capacity.

The first is local public budget expenditure, defined as each city's total general public budget expenditure in a given year, obtained from the *China City Statistical Yearbook*. The second is the share of the city's annual budget allocated to science and technology, defined as science-and-technology fiscal expenditure divided by total general public budget expenditure, from the same source. The two measures capture different aspects of local fiscal conditions: total spending capacity and its science-and-technology orientation.

For each measure, we classify cities into high- and low-capacity groups based on the sample-period median of the proxy and estimate Equation (1) separately on the two subsamples, using chi-square tests to evaluate the equality of OGDP coefficients between the two groups. Table C-4 reports the results. The OGDP coefficient is positive and significant in both subsamples for both fiscal-capacity measures, and the chi-square tests fail to reject the null of equal OGDP coefficients between groups ($p > 0.10$ in both cases). This pattern suggests that the OGDP effect is not driven by local fiscal capacity, ruling out the alternative interpretation that the estimated effect simply reflects greater public spending power in cities that adopt OGDPs.

Table C-4. Subsample Analyses by Local Fiscal Capacity and Spending Priorities

Alternative Explanations	Empirical Proxies	Type	TWFE-DID	# Observations	Adjusted R^2
Fiscal Capacity and Spending Priorities	local general public budget expenditure (year-level)	High	0.201*** (0.057)	6,192	0.722
		Low	0.283*** (0.083)	6,192	0.520
		Chi-Square	0.66		
	the share of fiscal spending allocated to science and technology (year-level)	High	0.185*** (0.066)	6,192	0.737
		Low	0.286*** (0.067)	6,192	0.597
		Chi-Square	1.15		

Notes: All columns include covariates as well as city and quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix D. Economic Significance

To quantify the incremental healthcare AI investment induced by the launch of OGDPs, we examine the effect along two complementary dimensions (technical attributes and application scenarios) and benchmark the resulting magnitudes against public expenditures on OGDP construction and operation. The economic magnitudes reported below correspond to the procurement-value response in treated municipalities, not to the contract-count outcomes used in our main specification.

For a treated city i, the cumulative incremental AI investment induced by the OGDP from its launch up to the post-treatment period T is estimated as:

$$\Delta Y_{Total} = \sum_{i \in Treat} \sum_{t=T_{i0}}^{T} \left[\overline{Y}_{i,pre} \times \left(e^{\beta} - 1 \right) \right]$$

Where β is the OGDP coefficient from a value-based version of Equation (1) reported in Tables D-1 and D-2, and $\overline{Y}_{i,pre}$ is city *i*'s baseline average healthcare AI procurement value prior to the OGDP launch.

D.1. Technical Attributes Perspective

Following the technical-attribute classification described in Appendix A.3, we group healthcare AI procurement contracts into four categories based on their core technical function: infrastructure, models & algorithms, perception & cognition, and execution & interaction. Table D-1 reports the value-based OGDP effect for each category.

The overall coefficient of 0.895 (Column 1) corresponds to a semi-elasticity of 144.7% (= $e^{0.895}$-1) relative to the pre-treatment baseline. Mapping this onto observed procurement values yields a cumulative incremental investment of approximately ¥5.67 billion (~$801 million; precise value ¥ 5,669,541,688) across treated municipalities.

The increase is concentrated in data-intensive analytical systems. Models & algorithms (Column 3, coefficient = 0.811) rises by 125.0% (= $e^{0.811}$-1), and perception & cognition (Column 4, coefficient = 0.694) rises by 100.2% (= $e^{0.694}$-1). Execution & interaction (Column 5, coefficient = 0.205) rises more modestly by 22.8% (= $e^{0.205}$-1), plausibly reflecting the greater organizational adjustments required to integrate AI outputs into clinical workflows and patient-facing applications. Infrastructure technologies (Column 2) show no statistically significant response, consistent with limited changes in hospital data throughput following OGDP launch.

Table D-1. OGDP Effect on Healthcare AI Procurement Value, by Technical Attribute

DV:	All	Infrastructure	Models & Algorithms	Perception & Cognition	Execution & Interaction
	(1)	(2)	(3)	(4)	(5)
OGDP	0.895*** (0.155)	0.052 (0.066)	0.811*** (0.153)	0.694*** (0.129)	0.205* (0.116)
Controls	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
# Observations	12,384	12,384	12,384	12,384	12,384
Adjusted R^2	0.536	0.173	0.427	0.348	0.284

Notes: All columns include covariates as well as city and quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

D.2. Application Scenarios Perspective

Following the application-scenario classification described in Appendix A.3, we group healthcare AI procurement contracts into three functional domains: clinical workflow & services, hospital operations & management, and extended services. Table D-2 reports the value-based OGDP effect for each.

The overall coefficient of 0.540 (Column 1) corresponds to a 71.6% (= $e^{0.540}$-1) increase in application-layer AI procurement value, equivalent to approximately ¥5.1 billion (~$720 million; precise value ¥5,098,464,140) across treated municipalities.

The effect concentrates in hospitals' core operational and clinical functions. Clinical workflow & services (Column 2, coefficient = 0.411) rises by 50.8% (= $e^{0.411}$-1), and hospital operations & management (Column 3, coefficient = 0.461) rises by 58.6% (= $e^{0.461}$-1). Extended services (Column 4, coefficient = 0.001) show no statistically significant response, indicating that OGDP-induced investment is directed toward hospitals' core production and coordination processes rather than peripheral or externally oriented services.

Table D-2. OGDP Effect on Healthcare AI Procurement Value, by Application Scenario

DV:	All	Clinical Workflow & Services	Hospital Operations & Management	Extended Services
	(1)	(2)	(3)	(4)
OGDP	0.540*** (0.137)	0.411*** (0.139)	0.461*** (0.127)	0.001 (0.027)
Controls	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
# Observations	12,384	12,384	12,384	12,384
Adjusted R^2	0.510	0.464	0.314	0.055

Notes: All columns include covariates as well as city and quarter fixed effects. Robust standard errors (clustered at the city level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

D.3. Economics Significance

To evaluate the fiscal leverage of OGDPs, we construct a cost-estimation framework encompassing both initial construction (fixed costs) and ongoing operations (variable costs). We stratify the 134 treated cities into three tiers based on *the Notice on Adjusting the Criteria for City Size Classification* issued by the State Council of China, combined with their administrative ranks and digital economy scales. Specifically, large metropolitan areas (approximately 15% of the sample), medium-sized regional hubs (35%), and smaller prefecture-level cities (50%) incur estimated baseline construction costs of ¥ 15 million (~$2.12 million), ¥8 million (~$1.13 million), and ¥ 3 million (~$0.42 million), respectively. Recognizing that digital infrastructure requires continuous maintenance, we add an annual operating expense equivalent to 10% of the initial construction cost,[2] accumulated quarterly post-launch. Aggregating these capital and operational expenditures across all treated municipalities yields a total estimated public investment of approximately ¥1.34 billion (~$189 million).

Building on the incremental investment increases calculated above, we compute the investment-leverage ratio associated with OGDP launch. The technical-attributes and application-scenarios represent two analytical lenses applied to the same underlying investment; their leverage ratios are therefore evaluated in parallel rather than additively. The leverage ratio is calculated by dividing the aggregate marginal effects of each dimension by the public expenditures incurred for platform construction and maintenance:

$$Leverage_Ratio = \frac{\Delta Y_{Total}}{\sum Cost_{OGDP}}$$

Under the technical-attributes lens, the ¥ 5.7 billion incremental investment yields a leverage ratio of approximately 4.23 (= 5.7/1.34). Under the application-scenarios lens, the ¥ 5.1 billion incremental investment yields a leverage ratio of approximately 3.80 (= 5.1/1.34). These leverage ratios should be interpreted as suggestive rather than causal: they do not account for OGDPs' broader social returns or full operational costs, nor do they capture spillovers beyond healthcare AI.

[2] This 10% rate falls within the 5–15% range documented in provincial government budget standards for IT operations and maintenance. See, e.g., Henan Provincial Department of Finance (2020), Provisions on Budget Standards for Provincial Information System Operations and Maintenance Projects (Trial), available at https://czt.henan.gov.cn/2020/08-03/1750092.html; Heilongjiang Provincial Department of Finance (2023), Budget Standards for Provincial Government Information System Operations and Maintenance Projects (Trial), available at https://czt.hlj.gov.cn/czt/c110776/202311/c00_31682398.shtml; Hunan Provincial Department of Finance (2024), Guidelines for Budget Preparation and Fiscal Review of Provincial Government-Invested Information Technology Projects (Trial), available at http://czt.hunan.gov.cn/czt/xxgk/zcfg/zcjd/202404/t20240415_33289667.html.

Appendix E. Exploring Underlying Mechanisms

E.1. Descriptive Statistics and Measurement of Mechanism Variables

Table E-1. Mechanism Variables, Measurement and Descriptive Statistics

Variable		Measurement	Mean	Std. Dev	Min	Max
			(1)	(2)	(3)	(4)
Governmental Signaling Mechanism Variables	Weibo posts	Number of OGDP-related posts disseminated through the government's official Weibo account in a city-quarter	0.109	0.546	0	17
	WeChat posts	Number of OGDP-related posts disseminated through the government's official WeChat account in a city-quarter	0.106	0.537	0	13
	Non-healthcare government departments	Number of non-healthcare government departments providing open data in a city-quarter	1.917	13.355	0	659
Data Provision Mechanism Variables	Healthcare data volume*	Number of healthcare data entries newly uploaded in a city-quarter	1.311	15.721	0	955
	Download per healthcare entry*	Download volume per healthcare data entry in a city-quarter	32.646	779.703	0	69,067

Notes: This table reports descriptive statistics for the mechanism variables used in the mechanism analyses. The variables are organized into two groups corresponding to the two mechanisms tested in the main text.

(1) *Governmental signaling variables*. The *non-healthcare government departments* enter regressions directly; WeChat and Weibo posts are standardized (z-score) and aggregated into a composite signaling-intensity index.

(2) *Data provision variables.* Healthcare data volume and Downloads per healthcare entry capture the volume and active use of healthcare data released through the OGDP. Both variables enter regressions in raw count form.

(3) Note on the "non-healthcare government departments" variable: as with the other mechanism variables in this table, the distribution is highly right-skewed because most city-quarters correspond to pre-OGDP periods or cities that never launch an OGDP, where the variable takes the value zero. In addition, this variable counts dataset-issuing units at different administrative levels (municipal bureaus, district-level branches, and sub-district offices) separately when each appears as a distinct data provider on the OGDP; the maximum value of 659 therefore reflects the cumulative participation of multi-tier administrative units in a single city-quarter rather than 659 independent municipal-level departments.

(4) Winsorization. Statistics above are based on raw data. In regressions, Non-healthcare government departments and the variables marked with * (whose right-tail values reach 659, 955, and 69,067, respectively) are winsorized at the 2nd and 98th percentiles. WeChat and Weibo posts are not winsorized, as their maxima (13 and 17) are too small to generate comparable leverage.

E.2. Robustness Check of Mechanisms Analysis

The median-split subsample tests reported in main-text Figures 6 and 7 impose a binary cutoff that can obscure variation near the median and discard information at the tails of each proxy's distribution. As a robustness check, we re-estimate each mechanism using continuous-interaction specifications that exploit cross-city variation in each proxy on a continuous scale. For each proxy, we standardize the variable (z-score) and interact it with OGDP, adding this interaction to Equation (1) while retaining the OGDP indicator and all baseline controls. Standardization places the four proxies on a common scale, allowing the magnitudes of the interaction coefficients to be directly compared across mechanisms. A positive and significant interaction coefficient indicates that cities scoring higher on the proxy experience a stronger OGDP effect, consistent with the corresponding mechanism.

Figure E-1 reports the four interaction coefficients, each estimated from a separate regression. The two governmental signaling interactions, corresponding to the cost dimension (number of non-healthcare government departments) and the observability dimension (intensity of governmental media communication), are positive and statistically significant. The two data-provision interactions, corresponding to the total volume of healthcare data and the downloads per healthcare data entry, are statistically indistinguishable from zero. These continuous-interaction estimates mirror the median-split evidence in main-text Figures 5 and 6, confirming that the asymmetry between the two mechanisms is not an artifact of the binary cutoff.

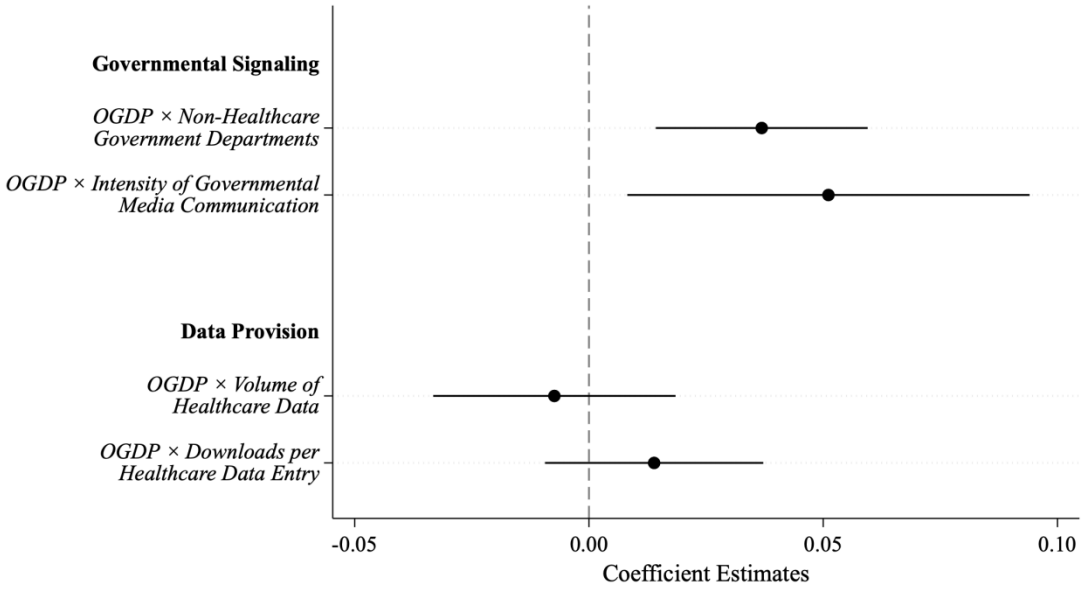


Figure E-1. Continuous-Interaction Estimates for the Two Mechanisms

Notes: The figure plots the OGDP×proxy interaction coefficients (with 95% confidence intervals) from four separate regressions, each adding one interaction term to Equation (1). Robust standard errors are clustered at the city level.

Appendix F. Heterogeneous Effects

We examine how city-level contextual factors moderate the OGDP effect on healthcare AI investment, organized around three core healthcare AI ecosystem participants whose responsiveness shapes whether OGDPs translate into investment outcomes.

As illustrated in Figure F-1, we propose a conceptual framework that depicts the interplay among local governments (the OGDP builder), hospitals (the demand-side investor), and AI vendors (the supply-side developer). The local conditions shaping each are, respectively, government digitalization priority, healthcare demand pressure, and local AI vendor supply. While Section 6 of the main text develops the theoretical motivation for these three moderators, this appendix details the measurement of each moderator, the subsample procedure, and the full regression estimates with formal tests of inter-group differences.

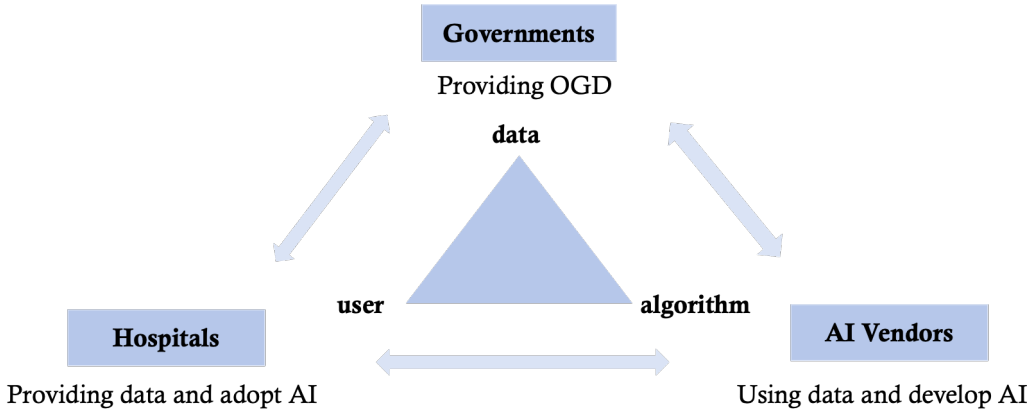


Figure F-1. Conceptual Framework: A Multi-Stakeholder Perspective

F.1. Measurement and Data Sources

Government digitalization priority. We construct this measure through textual analysis of annual local government work reports, which articulate municipal governments' policy priorities for the year. For each city–year, we count the frequency of digitalization-related keywords (e.g., "数字化"/digitalization, "数字经济"/digital economy, "数字政府"/digital government, "智慧城市"/intelligent city, "大数据"/big data) and normalize by the total word count of the report. This text-based measure of governmental policy attention follows an approach widely used in studies of policy commitment and signaling (Hassan et al., 2019). Government work reports are obtained from each municipality's official government website.

Healthcare demand pressure. We measure this as the public-hospital bed-occupancy rate at the province – year level, obtained from the *China Health Statistical Yearbook*. The bed occupancy rate is a widely used indicator of healthcare system strain in the health-economics literature (Berry Jaeker & Tucker, 2017). Higher occupancy rates indicate greater pressure on existing service capacity and stronger latent demand for efficiency-enhancing technologies.

Local AI vendor supply. We measure this as the share of healthcare AI procurement contracts in each city–quarter that are awarded to vendors registered in the same city as the procuring hospital. To identify vendor geography, we extract legal entity names and Unified Social Credit Codes (USCCs) from procurement contracts and match them to national and provincial firm-registration databases to determine each vendor's registered city. We then harmonize vendor and hospital locations using the National Bureau of Statistics' unified city-code system. Vendors registered in the same city as the procuring hospital are classified as local; others are classified as non-local. This local-share measure captures the depth of local market embedding between healthcare AI suppliers and hospital buyers, consistent with the geographic-clustering literature that emphasizes the role of local complementarity in enabling effective use of common infrastructure inputs (Porter, 1998).

F.2. Estimation and Inter-Group Tests

For each local condition, we split cities at the sample-period median and estimate Equation (1) separately on the high- and low-subsamples (each paired with the never-treated control cities). We then test the equality of the OGDP coefficients across the two subsamples using a chi-square statistic, which treats the subsample estimates as independent draws conditional on the city- and quarter-fixed effects already absorbed in each regression.

Table F-1 reports the results. Across all three conditions, the OGDP coefficient is positive and statistically significant in both the high and low subsamples, but the magnitude is substantially larger in the high subsample, with chi-square tests rejecting the equality of coefficients in each case.

Table F-1. Heterogeneous Effects of Contextual Characteristics

Contextual Characteristics	Grouping	Effect	Chi-square
Government digital priority	high	0.432***(0.078)	4.024**
	low	0.244***(0.052)	
Healthcare demand pressure	high	0.409***(0.058)	3.171*
	low	0.243***(0.073)	
Local AI vendor supply	high	0.707***(0.092)	19.343***
	low	0.240***(0.053)	

Notes: Robust standard errors (clustered at the city level) in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

References

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, 105(490), 493–505.

Athey, S., Bayati, M., Doudchenko, N., Imbens, G., & Khosravi, K. (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536), 1716-1730.

Bellemare, M. F., & Wichman, C. J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, 82(1), 50-61.

Beraja, M., Yang, D. Y., & Yuchtman, N. (2023). Data-intensive innovation and the state: Evidence from AI firms in China. *The Review of Economic Studies*, 90(4), 1701-1723.

Berry Jaeker, J. A., & Tucker, A. L. (2017). Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity. *Management Science*, 63(4), 1042-1062.

Burbidge, J. B., Magee, L., & Robb, A. L. (1988). Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association*, 83(401), 123-127.

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

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics,* 134(3), 1405-1454.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and Taxation: Theory and Evidence. The *American Economic Review*, 99(4), 1145–1177.

De Chaisemartin C, & d'Haultfoeuille X (2020) Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964-96.

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

Hassan, T. A., Hollander, S., van Lent, L., & Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, 134(4), 2135 - 2202.

Heckman, J. J., Ichimura, H., & Todd, P. E. (1997). Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme. *The Review of Economic Studies*, 64(4), 605–654.

La Ferrara, E., Chong, A., & Duryea, S. (2012). Soap operas and fertility: Evidence from Brazil. *American Economic Journal: Applied Economics*, 4(4), 1-31.

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

Mithas, S., & Krishnan, M. S. (2009). From association to causation via a potential outcomes approach. *Information Systems Research*, 20(2), 295–313.

Pham, P., Zhang, H., Gao, W., & Zhu, X. (2024). Determinants and performance outcomes of artificial intelligence adoption: Evidence from US Hospitals. *Journal of Business Research*, 172, 114402.

Rosenbaum, P. R., & Rubin, D. B. (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*, 70(1), 41–55.

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

Wang, J., Li, G., & Hui, K. L. (2022). Monetary incentives and knowledge spillover: Evidence from a natural experiment. Management Science, 68(5), 3549 - 3572.

Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1), 57-76.