Available Online: 3 August 2026 DOI: 10.54254/2977-5701/2026.35829

Who benefits from public data authorized operation? Spatially heterogeneous effects across China's urban hierarchy

Yizhang Chu, Qi Chen*

School of Economics and Management, Tongji University, Shanghai, China

*Corresponding Author. Email: 15381034918@163.com

Abstract. Marketized data governance regimes condition public-data access on accreditation, compliance, and complementary assets, yet their spatial consequences remain unclear. Using a balanced panel of 251 Chinese prefecture-level cities from 2020 to 2024, this paper examines China's Public Data Authorized Operation (PDAO) framework. Three findings emerge. First, PDAO adoption is associated with a positive average estimate for local enterprise stock, with especially large estimates for Tier-1 cities. Second, estimated effects follow pre-existing spatial structures, including the core–periphery hierarchy and the Hu Line. Third, the preferred Spatial Durbin and LeSage–Pace decompositions do not detect a statistically significant cross-city indirect effect. These findings show that the local economic effects associated with PDAO adoption are highly uneven across China's urban hierarchy. The observed gradient is consistent with a complementary-asset interpretation, although the research design does not separately identify this mechanism from other policy effects interacting with pre-existing urban endowments.

Keywords: public data authorized operation, place-based policy, spatial heterogeneity, enterprise agglomeration, urban hierarchy, China

1. Introduction

The spatial economics of the digital economy is increasingly shaped not only by the geography of digital infrastructure but also by the institutional architecture of data governance. Governments worldwide are transitioning away from open public data regimes toward marketized authorized access frameworks in which public-sector data is licensed to qualified operators subject to compliance requirements, technical standards, and fee-based access conditions. The European Union's Data Governance Act and GAIA-X infrastructure, the United Kingdom's National Health Service data trust program, India's National Data Sharing and Accessibility Policy, and China's Public Data Authorized Operation (PDAO) framework represent distinct national implementations. A central question for regional economics is whether the local economic effects associated with public-data authorization are evenly distributed across cities, or instead vary systematically with their positions in the urban hierarchy.

This paper addresses whether the local economic effects associated with PDAO adoption are evenly distributed across cities or instead vary systematically with their positions in the urban hierarchy. Across 251 Chinese prefecture-level cities from 2020 through 2024, the Tier-1 cities—Shanghai, Guangzhou, and Shenzhen, the three cities in the sample with both unambiguous Tier-1 status and cleanly defined treatment status—exhibit estimates approximately 12.7 times the full-sample average and 18.4 times the Other Tiers estimate. The estimated heterogeneity follows the established core–periphery structure and the Hu Line, which also provides the geographic reference for the distribution of the 2020–2023 PDAO adoption cohorts (Figure 1). The preferred Spatial Durbin model does not detect a statistically significant indirect effect. This result provides no evidence of systematic cross-city spillovers within the sample and spatial weight structures examined, but it does not rule out redistribution or other spatial mechanisms.

One plausible interpretation of this heterogeneity is that cities differ in their capacity to convert formal data access into economic activity. Computing infrastructure, specialized talent, compliance capacity, and downstream application ecosystems are unevenly distributed across the urban hierarchy. The empirical design does not separately identify these complementary assets as a distinct institutional mechanism; rather, it evaluates whether the distribution of estimated policy effects is consistent with this broader interpretation.

The paper makes three contributions. First, it provides a city-level assessment of the relationship between PDAO adoption and enterprise agglomeration. Second, it documents spatial heterogeneity across the urban hierarchy, core–periphery position, regions, and the two sides of the Hu Line. Third, it combines staggered-treatment estimation with spatial panel models to examine whether estimated local effects are accompanied by statistically detectable cross-city indirect effects.

The remainder proceeds as follows. Chapter 2 reviews the relevant regional economics literature. Chapter 3 develops the conceptual framework. Chapter 4 describes data and identification. Chapter 5 presents results. Chapter 6 reports additional analyses. Chapter 7 reports robustness checks. Chapter 8 discusses the implications and limitations, and Chapter 9 concludes. Appendix A presents a stylized complementary-asset model.

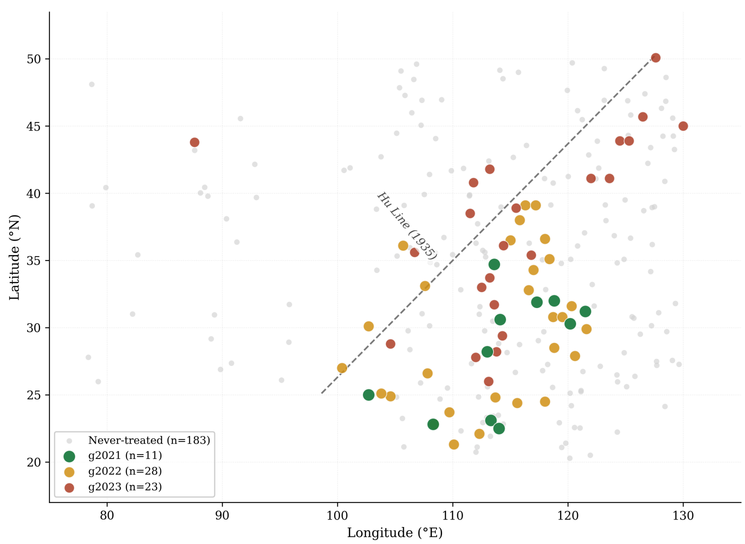

Figure 1. Geographic distribution of PDAO adoption cohorts, 2020–2023. 251 prefecture-level cities; the Hu Line [1] shown as a dashed line

2. Digital governance, agglomeration, and spatial inequality: a regional economics perspective

The spatially uneven incidence of place-based digital policy sits at the intersection of three bodies of regional economics literature: the agglomeration and spatial sorting literature, the place-based policy evaluation literature, and the digital economy spatial divergence literature. This chapter reviews each in turn and positions the paper's contribution relative to the existing frontier.

2.1. Agglomeration, complementary assets, and spatial concentration

The regional economics literature has long recognized that productive activity concentrates spatially at rates that exceed what geography alone would predict. Marshall [2] identified three mechanisms through which firms in proximity benefit from one another: labor market pooling that reduces search frictions, sharing of specialized inputs and intermediate services, and knowledge spillovers that accelerate learning and innovation. Rosenthal and Strange [3] provide a comprehensive survey documenting these mechanisms across industries, geographies, and time periods. The core of new economic geography, following Krugman [4], is that increasing returns to scale in the presence of transport costs generate stable spatial concentration even in ex ante symmetric locations. The concentration, once established, is self-reinforcing: firms locate in the core because demand is concentrated there, and demand is concentrated because firms and workers choose the core.

Combes et al. [5] show that spatial wage gaps partly reflect the sorting of high-productivity workers and firms into dense urban environments—spatial sorting—rather than agglomeration externalities alone. Moretti [6] and Autor [7] document that this self-reinforcing concentration has intensified over the past three decades, producing superstar cities that absorb disproportionate shares of high-skill workers, innovative firms, and economic value. The mechanism emphasizes the complementarity between high-skill workers and the modern knowledge economy: the arrival of high-skill workers attracts more high-skill workers and the service industries that supply them.

The concept of complementary assets as a threshold condition for technology adoption has a long tradition in economic geography [8, 9]. Adoption of a new general-purpose technology produces heterogeneous spatial returns because the returns depend on the co-availability of complementary inputs—physical infrastructure, human capital, and institutional capacity—that are themselves spatially unequally distributed. This paper extends that analysis from infrastructure provision to public-data authorization and examines whether estimated effects are concentrated in 127 cities with stronger pre-existing endowments.

2.2. Place-based policies: heterogeneity and distributional implications

The place-based policy evaluation literature has developed methodological tools and conceptual frameworks for understanding when spatially targeted interventions affect local economies. Neumark and Simpson [10] survey the evidence on enterprise zones, empowerment zones, and other place-based interventions, finding effects that are modest in magnitude and highly heterogeneous across space and implementation context. Kline and Moretti [11] develop a welfare economics framework for place-based policies and apply it to the Tennessee Valley Authority, finding that welfare interpretation depends critically on whether policy-induced agglomeration represents genuine productivity gains or redistribution from non-treated locations.

Heterogeneity of place-based policy effects is a recurring theme. Criscuolo et al. [12] document that UK industrial policy produces markedly different effects across firm types, with small manufacturing firms benefiting substantially while large firms show negligible effects. Busso et al. [13] find that US enterprise zone wage and employment effects are concentrated in zones with above-average pre-existing labor market capacity. Gibbons et al. [14] show that UK enterprise zone benefits accrue primarily to zones with strong connectivity to urban labor markets. A unifying finding is that place-based policies work best where complementary conditions are most favorable—precisely the places that need them least. This concern aligns with recent work in Spatial Economic Analysis that treats regional inequality as a central theme of spatial economic inquiry and shows that inequality measurement depends on how spatial units and population weights are defined [15, 16].

This paper examines authorized data access as a form of place-based digital policy. It evaluates the association between PDAO adoption and local enterprise agglomeration, documents where the estimated effects are concentrated across the urban hierarchy, and uses the spatial Durbin model decomposition to assess whether local estimates are accompanied by detectable indirect effects.

2.3. Digital economy, spatial divergence, and the urban hierarchy

The spatial economics of the digital economy has emerged as a significant research frontier. Goldfarb and Tucker [17] show that digital technologies reduce distance costs for consumption while production remains spatially concentrated, because producing digital services requires dense clusters of specialized human capital and physical infrastructure. Digital technologies can simultaneously reduce spatial frictions for final users while reinforcing spatial concentration in production. Brakman et al. [18] document that the European digital economy closely tracks the pre-existing urban hierarchy, with metropolitan cores accounting for a disproportionate share of digital activity. De Stefano et al. [19] find that European digital divides have a persistent geographic dimension that observable economic fundamentals cannot fully explain. Recent work in *Spatial Economic Analysis* further shows that digitalisation can amplify wage inequality through urbanisation economies and metropolitan advantages [20].

Several studies have examined digital infrastructure interventions. Hjort and Poulsen [21] exploit submarine internet cable arrivals in African cities to estimate causal effects on firm performance, finding large average effects but substantial heterogeneity across firm types attributable to differential complementary asset endowments. Akerman et al. [22] document that Norwegian broadband rollout effects are concentrated among high-skill workers and capital-intensive firms. The finding of this paper—that even a uniform eligibility regime produces effects concentrated in high-endowment cities—extends these results to the institutional dimension of digital access.

This paper extends analysis from digital infrastructure provision to public-data authorization and documents heterogeneity associated with a data-governance intervention. It shows where estimated effects are concentrated without claiming that data-governance institutions constitute a distinct channel analytically separate from infrastructure or other policy effects.

2.4. Spatial econometrics: local versus spillover effects in regional policy

The spatial econometric methodology employed in this paper follows LeSage, Pace [23] and Elhorst [24], who develop maximum likelihood estimators for spatial panel data models. Halleck Vega and Elhorst [25] compare the Spatial Lag Model (SLM), Spatial Error Model (SEM), and Spatial Durbin Model (SDM), recommending the SDM as a general starting point because it nests both the SLM and the spatial lag of X model as special cases. The direct-versus-indirect effects decomposition permits an assessment of whether estimated local effects are accompanied by statistically detectable cross-city indirect effects.

This approach also connects to the broader spatial macroeconomics agenda in *Spatial Economic Analysis*, which links aggregate outcomes to the geographical distribution of economic activity [26]. The distinction between local and cross-city effects also speaks to the equity-efficiency trade-off in regional policy design [27].

The test in this paper asks whether θ on $W \cdot did$ in the SDM is distinguishable from zero and whether the estimated local effects are accompanied by detectable cross-city spillovers. A statistically insignificant estimate should be interpreted as no statistically detectable indirect effect within the sample and spatial weight structures examined, not as proof that spillovers or redistribution are absent. The paper combines staggered-treatment estimators with spatial panel models to examine local estimates and potential indirect effects.

3. Conceptual framework: complementary assets and spatially heterogeneous policy effects

This chapter develops a conceptual framework connecting the regional economics literature to PDAO. Section 3.1 describes PDAO and the complementary assets that may condition its local use. Sections 3.2 and 3.3 present a cautious interpretation of spatial heterogeneity, and Section 3.4 states the research expectations and question.

3.1. Public data authorized operation and its complementary asset requirements

PDAO is an institutional arrangement under which a government agency authorizes designated operators to commercialize specific public data resources, subject to security, compliance, and quality requirements. Three features are central. First, the authorization process imposes ex ante qualification requirements including competence in privacy-preserving computation, regulatory compliance capacity, and financial stability. These requirements are uniform in formal terms but differentially binding across jurisdictions depending on local endowments. Second, the operational requirements include technical infrastructure for hosting data analyses without removing raw data from the originating jurisdiction, generating a complementary asset requirement for high-performance computing capacity. Third, the value extracted from PDAO requires downstream applications by organizations that can productively use authorized data products, generating complementary asset requirements for talent, venture capital, and application-area expertise.

These three categories—qualification capacity, computational capacity, and downstream application capacity—may shape a jurisdiction's capacity to convert formal authorization into productive economic applications. The relevant endowments vary substantially across Chinese cities, with higher-tier cities generally possessing greater digital infrastructure, concentrations of data science talent, venture capital ecosystems, and regulatory capacity.

3.2. Complementary assets and expected heterogeneity

Appendix A presents a stylized complementary-asset model, summarized in Equations (A1)–(A2), whose main implication is an expected treatment-effect gradient across the urban hierarchy. The model is used as an interpretive framework rather than as evidence that a distinct mechanism has been identified.

3.3. Interpreting spatially uneven policy effects

Cities at different positions in the urban hierarchy may differ in their capacity to convert formal data access into economic activity. Computing infrastructure, specialized talent, compliance capacity, and downstream application ecosystems are unevenly distributed, so similar authorization arrangements may be associated with different local outcomes.

This interpretation draws on the broader literature on complementarities between policy interventions and local assets. The empirical analysis evaluates whether the distribution of estimated effects is consistent with that interpretation; it does not identify institutional design as a distinct channel or an institutional analogue of skill-biased technical change.

This framework provides an interpretation of the expected heterogeneity rather than a separately identified causal mechanism.

3.4. Research expectations

Expectation 1: PDAO adoption is associated with an increase in local enterprise stock on average.
Expectation 2: Estimated effects may be larger in higher-tier and core cities than in lower-tier and peripheral cities.
Research Question 3: Are the estimated local effects accompanied by detectable cross-city spillovers?

4. Data and empirical strategy

4.1. Panel construction

The empirical analysis draws on a balanced panel of 251 Chinese prefecture-level cities observed annually from 2020 through 2024, yielding 1,255 city-year observations. Pilot PDAO adoption began in six cities in 2020, expanded by eleven additional cities in 2021, accelerated to twenty-eight new cities in 2022, and added twenty-three further cities in 2023, producing a treatment group of sixty-eight cities and a control group of 183 never-treated cities.

Beijing and Tianjin are excluded because their treatment status cannot be cleanly defined within the 2020–2024 window; details are reported in Appendix B. The three conventional Tier-1 cities included in the Tier-1 subsample analysis are Shanghai, Guangzhou, and Shenzhen.

The panel is balanced through targeted imputation of thirty-nine cities' 2020 baseline observations across Henan (seventeen cities), Hunan (thirteen cities), and Fujian (nine cities) provinces. Data sources include the State Administration for Market Regulation's enterprise registry (firm-level, aggregated to city-year), the *China City Statistical Yearbook* (macroeconomic controls), and the *Telecommunication Annals of China 1881–1920* (telegraph instrument).

4.2. Variables and heterogeneity classifications

The principal outcome variable is the natural logarithm of city enterprise stock (Table 1). Five control variables enter every specification: log GDP, log population, tertiary industry share, internet penetration, and science fiscal expenditure ratio. Heterogeneity analyses partition the sample along four orthogonal dimensions: a three-way core–periphery classification, a four-way city tier classification, a three-way regional classification, and a binary Hu Line classification. The Tier-1 classification includes the three prefecture-level cities with unambiguous treatment status within the sample window: Shanghai, Guangzhou, and Shenzhen.

Shanghai and Chongqing are retained as large never-treated spatial counterfactuals in the SDM specifications; details are reported in Appendix B.

Table 1. Descriptive statistics: treatment vs. control cities at the 2020 baseline

Variable	Treated (n = 68)		Control (n = 183)		p-value
	Mean	SD	Mean	SD	
Outcome variables					
Enterprise stock (count)	203,894	(228,688)	90,169	(140,181)	< 0.001
ln(enterprise stock)	11.573	(1.253)	10.983	(0.840)	< 0 .001
ln(high-tech enterprise stock)	9.274	(1.445)	8.189	(1.047)	< 0.001
Macroeconomic controls					
ln(GDP)	8.205	(0.980)	7.400	(0.856)	< 0.001
ln(Population)	6.159	(0.590)	5.788	(0.716)	< 0.001
Tertiary industry share (%)	52.18	(9.59)	48.11	(6.66)	0.002
Internet penetration (%)	37.89	(11.39)	35.93	(9.30)	0.206
Science fiscal ratio (%)	2.665	(2.753)	1.616	(1.638)	0.004
Geographic and instrument					
Historical telegraph offices	1.21	(1.39)	1.12	(1.45)	0.668
Distance to capital (km)	111.1	(76.9)	141.3	(73.5)	0.006
Number of cities	68		183		
Total observations	340		915		

Notes: Cross-section at the 2020 baseline year. Two-sample t-tests with unequal variances. Treated cities (n = 68) adopted PDAO during 2020–2023; control cities (n = 183) are never-treated. Beijing and Tianjin are excluded from the sample because their treatment status cannot be cleanly defined within the 2020–2024 window. The historical telegraph office variable is the instrument used in IV-2SLS estimation.

4.3. Identification strategy

The empirical analysis adopts a sequential identification strategy. The principal specification is the Two-Way Fixed Effects Difference-in-Differences (TWFE-DiD) equation with city and year fixed effects, controls, and city-clustered standard errors. The staggered adoption structure is addressed through Callaway and Sant'Anna [28] cohort-time aggregation. Spatial dependence is examined through SLM, SEM, and SDM specifications. IV estimation using the historical telegraph instrument is reported as a supplementary sensitivity analysis addressing potential endogeneity in policy adoption.

4.4. The historical telegraph instrument

The telegraph office instrument is relevant through the documented persistence of information infrastructure endowments across more than a century. The first-stage Kleibergen–Paap Wald F [29] of 45.34 exceeds the Stock–Yogo critical value [30] of 16.38 and the Olea–Pflueger threshold [31] of 24.09, indicating strong first-stage relevance but not establishing exogeneity. The exclusion restriction relies on the assumption that historical telegraph placement affects contemporary enterprise agglomeration only through its relationship with PDAO adoption, conditional on the included controls. This assumption cannot be tested directly, and the IV results are therefore interpreted cautiously. The instrument is balanced across treated and control cities (p = 0.668, Table 1).

5. Results

This chapter presents the principal empirical findings corresponding to the two expectations and the research question stated in Section 3.4.

5.1. Average effect and city-tier heterogeneity

The full-sample TWFE-DiD coefficient is 0.025 (SE = 0.012, p = 0.036), implying a 2.6 percent increase in enterprise stock for treated cities; however, the city-tier decomposition shows that this average obscures substantial distributional heterogeneity (Table 2). The Tier-1 effect—estimated on the three prefecture-level cities (Shanghai, Guangzhou, and Shenzhen) with unambiguously defined treatment status within the sample window—is 0.317 (HC1, SE = 0.199). Because the Tier-1 subsample contains only G = 3 cities, the conventional t-statistic has df = 2 and is severely size-distorted under standard asymptotics (conventional p = 0.252). This paper therefore reports the permutation p-value as the preferred inference: randomly assigning the "Tier-1" label to any three cities in the full 251-city pool and computing the same subsample regression 999 times yields an empirical p-value of 0.038, suggesting that the observed magnitude of 0.317 is unusual under this relabelling exercise. The estimate corresponds to a 37.3 percent increase in enterprise stock (exp(0.317) − 1). The New Tier-1 effect is 0.051 (p = 0.115), the Tier-2 effect is 0.048 (p = 0.224), and the Other Tiers effect, applying to 224 of the 251 cities, is 0.017 (p = 0.178).

The Tier-1 effect is approximately 12.7 times the full-sample average (0.317/0.025) and approximately 18.4 times the Other Tiers effect (0.317/0.017). The gradient is monotonic: Tier-1 (0.317) > New Tier-1 (0.051) > Tier-2 (0.048) > Other Tiers (0.017) (Figure 2). These estimates reveal a pronounced gradient across the city-tier classification, although the Tier-1 estimate should be interpreted cautiously because it is based on only three cities. The ratios are descriptive comparisons rather than stable structural parameters.

Table 2. Heterogeneous effects across the city tier hierarchy

Subsample	Cities (G)	β (DiD)	SE	p-value
Full sample	251	+0.0254**	(0.0120)	0.036
Tier-1 (Shanghai, Guangzhou, Shenzhen)	3	+0.3174	(0.1989)	perm = 0.038
New Tier-1	13	+0.0511	(0.0312)	0.115
Tier-2	10	+0.0482	(0.0369)	0.224
Other Tiers	224	+0.0172	(0.0127)	0.178

Notes: TWFE-DiD estimates with city and year fixed effects. Tier-1 includes the three prefecture-level cities with unambiguous treatment status within 2020–2024: Shanghai, Guangzhou, and Shenzhen. Beijing is excluded (treatment status uncleanly defined: large-scale procurement activity without a formally enacted municipal regulation); Tianjin is excluded (formal policy document issued at end of sample period). Tier-1 subsample has G = 3 cities (df = 2); conventional p = 0.252 is unreliable under small-G asymptotics. Preferred inference: permutation test (999 random relabellings of the Tier-1 label from the full 251-city pool) yields empirical p = 0.038. Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.

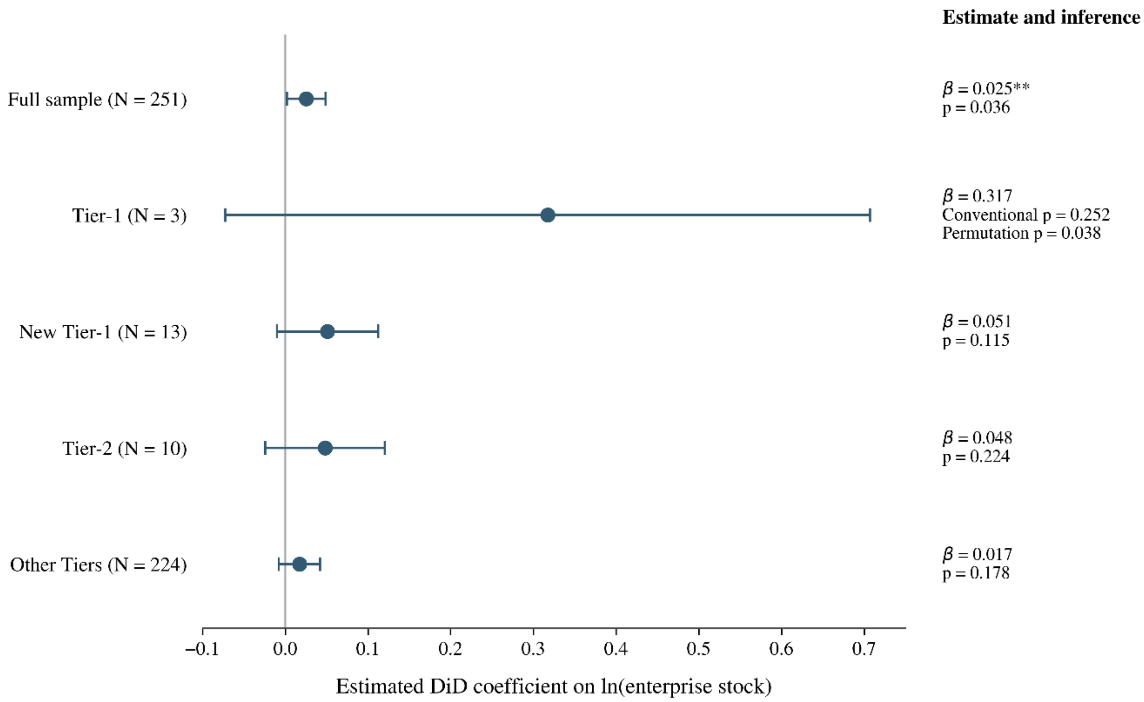

Figure 2. PDAO effects across the city tier hierarchy. The Tier-1 effect of 0.317 is approximately 12.7 times the full-sample average (0.025) and 18.4 times the Other Tiers effect (0.017). Tier-1 comprises Shanghai, Guangzhou, and Shenzhen. Conventional $p = 0.252$; permutation $p = 0.038$

5.2. Core–periphery and Hu Line heterogeneity

Table 3. Spatial heterogeneity: core–periphery, region, and Hu Line decompositions

Panel	Subsample	Cities	β (DiD)	SE	p-value
A. Core–Periphery	Core	92	+0.0732***	(0.0200)	< 0.001
	Semi-periphery	111	−0.0025	(0.0151)	0.870
	Periphery	48	−0.0184	(0.0192)	0.343
B. Region	Eastern	95	+0.0375***	(0.0134)	0.006
	Central	78	+0.0411	(0.0349)	0.243
	Western	78	+0.0235	(0.0236)	0.321
C. Hu Line	East of Hu Line	226	+0.0241*	(0.0129)	0.062
	West of Hu Line	25	+0.0047	(0.0132)	0.724

Notes: TWFE-DiD subsample estimates by core–periphery position (Panel A), region (Panel B), and Hu Line position (Panel C). Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively. Cluster-robust standard errors at the city level.

Core cities ($n = 92$) exhibit an estimated effect of 0.073 ($p < 0.001$), compared with −0.003 for semi-periphery cities ($p = 0.870$) and −0.018 for periphery cities ($p = 0.343$). The gradient is monotonic, and the core–semi-periphery difference is significant at the one-percent level (Table 3, Panel A; Figure 3(a)). Across regions, the point estimates are positive, but only the Eastern estimate is statistically distinguishable from zero (Table 3, Panel B; Figure 3(a)).

The 226 cities east of the Hu Line exhibit an estimated effect of 0.024 (p = 0.062), whereas the twenty-five cities west of the line—encompassing parts of Inner Mongolia, Gansu, Qinghai, Xinjiang, Tibet, Sichuan, and Yunnan—exhibit an estimate of 0.005 (p = 0.724), approximately one-fifth of the eastern magnitude (Table 3, Panel C; Figure 3(b)). The Hu Line is a demographic-geographic boundary identified in 1935 by Hu Huanyong based on population density patterns that have remained stable over ninety years. The estimated effects are aligned with a long-standing demographic and economic divide, with larger estimates east of the Hu Line than west of it.

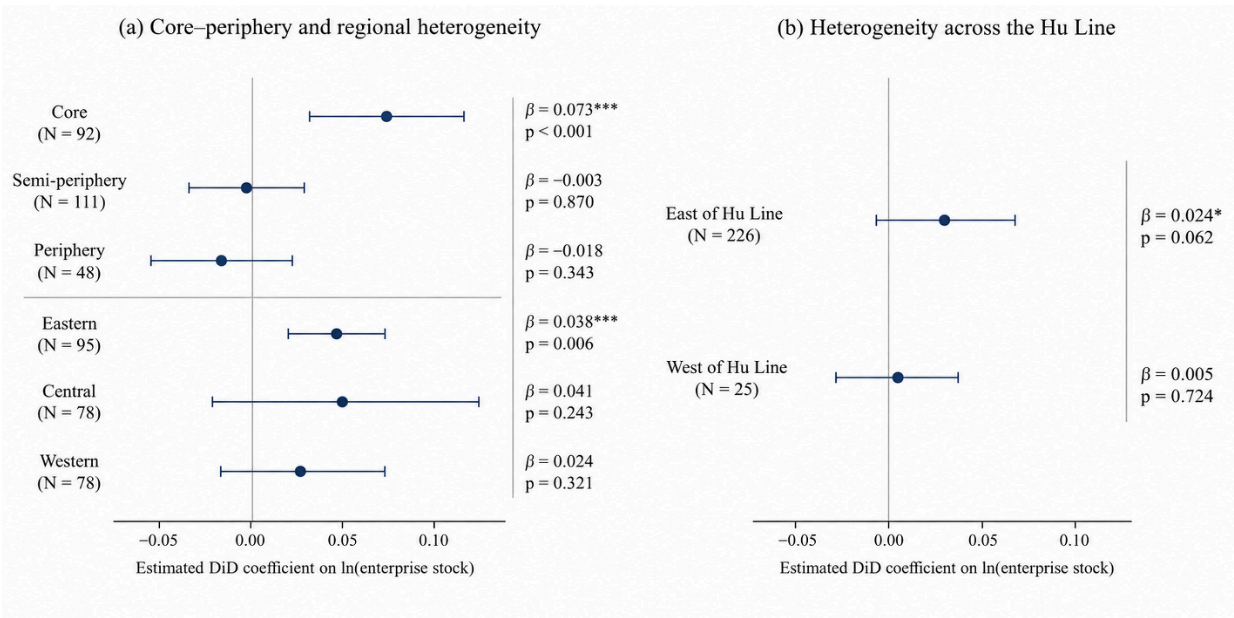

Figure 3. Spatial heterogeneity tracks pre-existing geographic-economic structures. Panel (a) shows the core–periphery gradient; panel (b) highlights the five-fold Hu Line gap

5.3. Cross-city indirect effects

Table 4. Spatial econometric analyses: SDM specifications

Model	ρ	p (ρ)	θ on $W{\cdot}did$	p (θ)	Specification
OLS (no spatial)	—	—	—	—	*No W*
SAR (Spatial Lag)	+0.7264	< 0.001	—	—	*KNN-25*
SDM – restricted [durbin(did) only]	+0.8056	< 0.001	−0.0398**	0.039	*KNN-25*
SDM – full [durbin(did ×X)] — principal	+0.7264	< 0.001	−0.0354	0.118	*KNN-25*

Notes: SDM estimated on N = 1,255 panel. The restricted spec [durbin(did) only] produces a marginally significant θ due to omitted spatial lags. The full LeSage-Pace spec [durbin(did ×X)] is the principal specification; θ = −0.035, p = 0.118. ρ is the spatial autoregressive coefficient; θ is the spatial lag of the treatment indicator. Significance: 8**, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.

In the preferred full Spatial Durbin model, the coefficient on the spatially lagged treatment variable is not statistically significant (Table 4). This result provides no evidence of systematic cross-city spillovers within the sample and spatial weight structures examined, but it does not rule out redistribution or other spatial mechanisms. Detailed direct, indirect, and total effect estimates are reported in Table C1 in Appendix C.

5.4. Instrumental variables estimates

The IV-2SLS estimate is 0.041 (SE = 0.019, $p = 0.031$), compared with the TWFE estimate of 0.025 (Table 5, Panel A). The corresponding first-stage diagnostics yield a Kleibergen–Paap F-statistic of 45.34 (Table 5, Panel B). The IV estimate is approximately 1.6 times the TWFE estimate. The larger IV estimate is consistent with negative selection into adoption, but it may also reflect treatment-effect heterogeneity or differences between the populations identified by the TWFE and IV estimators. The IV results are therefore interpreted cautiously.

Table 5. Average treatment effect: OLS, TWFE, and IV-2SLS comparisons

Panel A: Second-stage estimates				
Specification	β (DiD)	SE	p-value	First-stage F
(1) OLS — No controls	+0.0915	(0.2430)	0.706	—
(2) TWFE — Full controls	+0.0254**	(0.0120)	0.036	—
(3) IV-2SLS — Telegraph instrument	+0.0406**	(0.0187)	0.031	45.34
IV-to-TWFE ratio	1.60×	—	—	—

Panel B: First-stage and weak-IV diagnostics	
Diagnostic	Estimate / statistic
First-stage coefficient on iv_telegraph_post	+0.3071
First-stage SE (cluster-robust)	(0.0456)
First-stage t-statistic	6.74
Kleibergen–Paap cluster-robust Wald F	45.34
Olea–Pflueger 5% effective F threshold	24.09
Anderson–Rubin 95% CI for IV β	[0.005, 0.082]
Stock–Yogo 10% maximum bias critical value	16.38

Notes: Panel A: second-stage estimates. Panel B: first-stage diagnostics for specification (3). The Kleibergen–Paap F exceeds both the Stock–Yogo (16.38) and Olea–Pflueger (24.09) critical values. Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.

The pre-treatment coefficients are not statistically distinguishable from zero, whereas the post-treatment point estimates generally increase over event time (Figure 4). The pre-treatment pattern is supportive of, but does not establish, the parallel-trends assumption.

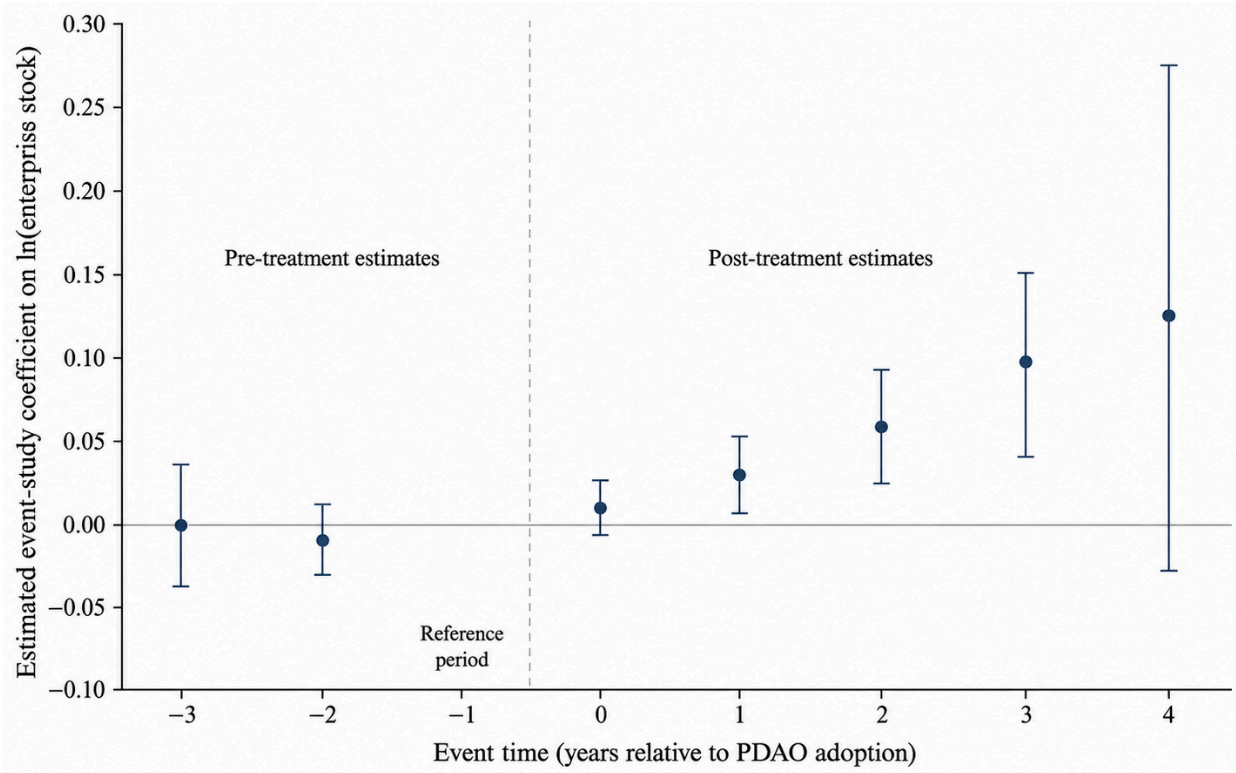

Figure 4. Event-study estimates before and after PDAO adoption

6. Additional evidence and potential interpretations

This chapter reports additional evidence relevant to potential interpretations of the heterogeneity. Section 6.1 examines changes in high-technology enterprise composition, Section 6.2 documents cohort heterogeneity, Section 6.3 reports IV estimates and potential selection in adoption, and Section 6.4 summarizes the interpretive limits.

6.1. Changes in the composition of local enterprise activity

The three-step Baron–Kenny mediation procedure [32] reports that the high-technology firm ratio accounts for approximately 10.5 percent of the total estimated treatment effect: first-step coefficient 0.024 (p = 0.017), second-step coefficient 0.156 (p = 0.007), indirect effect bootstrap 95% CI [0.010, 0.052]. The digital firm ratio (a-path p = 0.326) and total patent count (a-path p = 0.412, b-path p = 0.234) do not show statistically significant mediation estimates. The conventional mediation decomposition is descriptive and does not establish a causal mediation channel.

The general equilibrium framework of Xie et al. [33] offers one possible interpretation centered on computing capacity. City-level computing infrastructure data are not consistently available across the full panel, however, so the relationship between computing capacity and the estimated effects remains suggestive and is a priority for follow-up research.

6.2. Cohort heterogeneity

The cumulative ATT is 0.076 for the g2021 cohort (n = 11, p < 0.001), 0.028 for the g2022 cohort (n = 28, p = 0.004), and 0.004 for the g2023 cohort (n = 23, p = 0.679) (Table 6; Figure 5(b)). The cohort-specific event-time trajectories show the same broad ordering, with the g2021 estimates increasing most strongly at later event times (Figure 5(a)). Taken together, these estimates indicate larger cumulative associations for earlier adoption cohorts. However, cohort membership is correlated with exposure duration, city composition, and adoption timing, so the estimates document cohort heterogeneity but do not independently establish first-mover lock-in.

Algorithmic accumulation, talent mobility, platform persistence, and constraints on computing capacity are possible explanations for the cohort pattern discussed in the path-dependence literature [8, 9, 33]. These explanations are not measured or directly tested here and therefore should not be interpreted as validated lock-in mechanisms.

Table 6. Cohort-specific cumulative average treatment effects on the treated

Cohort	Cumulative ATT	SE	p-value	β at $t = +2$	β at $t = +3$
g2021 ($n = 11$)	+0.0757***	(0.0115)	< 0.001	+0.0916	+0.1233
g2022 ($n = 28$)	+0.0281***	(0.0099)	0.004	+0.0459	—
g2023 ($n = 23$)	+0.0037	(0.0090)	0.679	—	—

Notes: Cohort-specific cumulative ATTs estimated against the never-treated control group. Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.

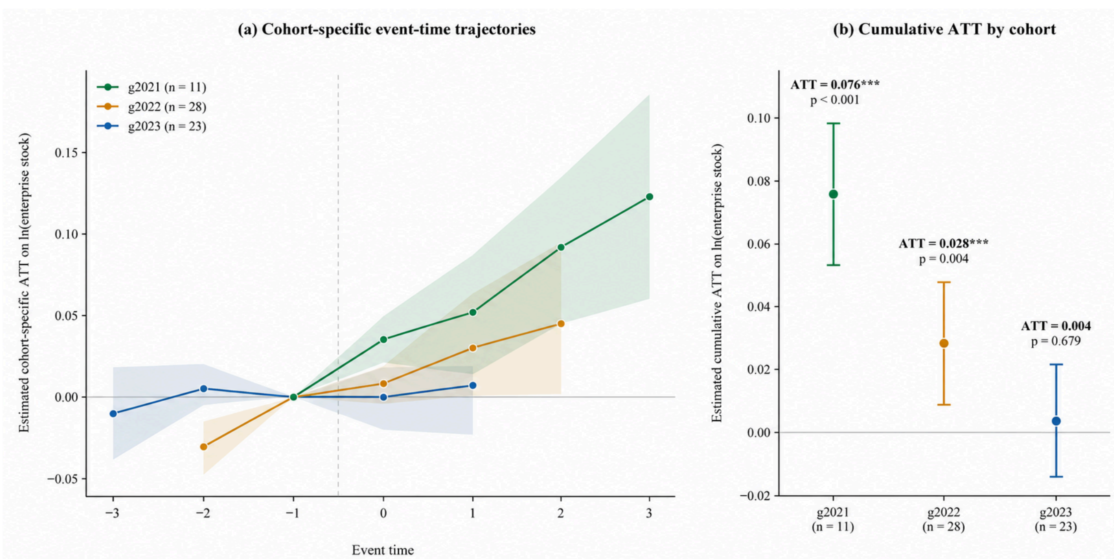

Figure 5. Cohort heterogeneity in estimated treatment effects. (a) Cohort-specific event-time trajectories. (b) Cumulative ATT by cohort

6.3. IV estimates and potential selection in adoption

The IV estimate exceeds the TWFE estimate (0.041 versus 0.025). The larger IV estimate is consistent with negative selection into adoption, but it may also reflect treatment-effect heterogeneity or differences between the populations identified by the TWFE and IV estimators. The g2022 cohort's negative coefficient at event time minus two and the IV–TWFE difference provide suggestive evidence only and do not independently identify the adoption-selection process.

6.4. Summary

Taken together, changes in high-technology enterprise composition, differences across adoption cohorts, and IV estimates provide supplementary, interpretive evidence. They are consistent with several possible explanations for the observed heterogeneity but do not separately identify causal mediation, first-mover lock-in, negative selection, or a four-channel polarization mechanism.

7. Robustness

This chapter reports robustness checks addressing alternative outcomes, imputation choices, COVID-related specifications, small-subsample inference, placebo tests, alternative classifications, leave-one-out estimates, and alternative spatial weight matrices. Where estimable, the Tier-1 estimates remain positive across the principal robustness checks, although the associated inference varies with the specification and small-sample procedure (Table 7). Additional leave-one-out and spatial-weight-matrix results are reported in Tables D1–D2 in Appendix D.

Table 7. Master summary of robustness checks

Robustness check	Tier-1 β	p-value	Consistent?
Baseline (full balanced panel, N = 1,255)	+0.317	perm = 0.038	Yes (reference)
7.1: High-technology firm ratio (hightech_ratio)	+0.407	0.008	Yes ***
7.2: Excluding three Tier-1 cities (other-tier ATE)	n/a	n/a	Other-tier β = +0.025**, p = 0.045
7.2: Unbalanced panel (drop imputed 2020 baselines)	+0.317	perm = 0.038	Yes (Tier-1 unchanged)
7.3: Province × 2020 fixed effects (full sample)	n/a	n/a	Full-sample β = +0.020*, p = 0.056
7.4: Permutation test (999 relabellings, N=251 pool)	+0.317	0.038	Yes **
7.5: Placebo test (200 reps, shuffle treat_year)	+0.317	0.040	Yes **
7.6: Binary core vs. non-core classification	n/a	n/a	Core β = +0.045***
7.6: Top tercile of complementary asset score	n/a	n/a	Top β = +0.084***
7.7: Leave-one-out (drop Shanghai)	+0.108**	0.033	Yes (simplified spec)
7.7: Leave-one-out (drop Guangzhou)	omitted	—	did collinear with FE
7.7: Leave-one-out (drop Shenzhen)	+0.127**	0.029	Yes (simplified spec)
7.8: SDM KNN-5 weight matrix	n/a	n/a	θ ns, ρ sig.
7.8: SDM inverse-distance W	n/a	n/a	θ ns, ρ sig.
7.8: SDM GDP-similarity W	n/a	n/a	θ = +0.030*, ρ sig.
7.8: SDM binary contiguity W	n/a	n/a	θ ns, ρ sig.

Notes: All estimates verified from panel_es_251_full.csv (N = 1,255; 251 cities × 5 years). Baseline p-value is permutation p (999 relabellings from the full 251-city pool); conventional HC1 p = 0.252 at df = 2 is not the primary inference. 7.2 Tier-1 unbalanced-panel estimate equals the baseline because Shanghai, Guangzhou, and Shenzhen are not in the three imputation-affected provinces. 7.3 drop-2020 full-sample estimate (β = 0.007, ns) is discussed in Section 7.3 but excluded from this table because identification is mechanically limited to g2022 and g2023 cohorts after 2020 is removed. Guangzhou LOO: dropping Guangzhou leaves Shanghai (never-treated) and Shenzhen (always-treated), making did perfectly collinear with city and year FE —mechanical identification failure. LOO uses simplified spec (city and year FE only); baseline simplified = 0.118 (p = 0.003). Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.

8. Discussion and implications

8.1. Reframing the data governance debate in regional economics terms

The regional economics of place-based policy has focused primarily on physical infrastructure and fiscal incentives. This paper extends that discussion to public-data authorization and documents pronounced spatial heterogeneity in the estimates associated with PDAO adoption. The Tier-1 estimate is approximately 12.7 times the full-sample estimate, but this descriptive comparison is based on only three Tier-1 cities and should not be interpreted as a stable structural ratio. The concentration of computing infrastructure, data science talent, and compliance capacity in higher-tier cities offers one plausible complementary-asset interpretation, although these channels are not separately identified.

The preferred Spatial Durbin model does not detect a statistically significant indirect effect. This result provides no evidence of systematic cross-city spillovers within the sample and spatial weight structures examined, but it does not rule out redistribution or establish that local differences reflect a single mechanism.

8.2. Implications for european data governance architecture

Although the institutional settings differ substantially, the Chinese evidence suggests that formal access alone may not guarantee geographically balanced outcomes. Whether similar patterns occur under other data-governance arrangements remains an open empirical question.

8.3. Policy implications

Uniform adoption requirements may need to be accompanied by differentiated capacity-building measures, especially in computing infrastructure, technical expertise, and application ecosystems. Because these channels are not directly identified in the present analysis, the appropriate policy mix requires further evaluation.

8.4. Limitations

The findings rest on several choices that bound generalizability. First, the five-year window prevents observing post-2024 dynamics, including potential convergence as mid-tier cities build complementary assets. Second, the Chinese institutional context may limit generalizability to other settings. Third, city-level computing infrastructure data are not consistently available, preventing a direct empirical test of the role of computing capacity. Most importantly, the research design identifies heterogeneous effects associated with PDAO adoption but does not isolate a distinct institutional mechanism. The observed hierarchy may reflect PDAO-specific qualification and compliance requirements, more general complementarities between policy interventions and local assets, or both. Distinguishing these explanations would require direct measures of institutional restrictions or comparison with alternative public-data regimes.

9. Conclusion

Using a balanced panel of 251 Chinese prefecture-level cities from 2020 to 2024, this study examines whether the local economic effects associated with PDAO adoption vary across China's urban hierarchy. The estimates are positive on average but highly heterogeneous. Larger estimates are concentrated in Tier-1 and core cities and on the eastern side of the Hu Line. At the same time, the preferred Spatial Durbin and LeSage–Pace decompositions do not detect a statistically significant cross-city indirect effect within the sample and spatial weight structures examined.

The observed gradient is consistent with a complementary-asset interpretation, under which cities differ in their capacity to convert formal data access into local economic activity. However, the current research design does not separately identify this interpretation as a PDAO-specific institutional mechanism. Future research should directly measure institutional restrictions and complementary assets or compare PDAO with alternative public-data regimes.

References

[1] Hu, H. Y. (1935). The distribution of China's population. *Acta Geographica Sinica, 2*(2), 33–74.

[2] Marshall, A. (1890). *Principles of economics*. London: Macmillan.

[3] Rosenthal, S. S., & Strange, W. C. (2004). Evidence on the nature and sources of agglomeration economies. *Handbook of Regional and Urban Economics, 4*, 2119–2171.

[4] Krugman, P. (1991). Increasing returns and economic geography. *Journal of Political Economy, 99*(3), 483–499.

[5] Combes, P.-P., Duranton, G., & Gobillon, L. (2012). The costs of agglomeration: House and land prices in French cities. *Review of Economic Studies, 79*(4), 1509–1542.

[6] Moretti, E. (2012). *The new geography of jobs*. Boston: Houghton Mifflin Harcourt.

[7] Autor, D. (2019). Work of the past, work of the future. *AEA Papers and Proceedings, 109*, 1–32.

[8] David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review, 75*(2), 332–337.

[9] Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal, 99*(394), 116–131.

[10] Neumark, D., & Simpson, H. (2015). Place-based policies. *Handbook of Regional and Urban Economics, 5*, 1197–1287.

[11] Kline, P., & Moretti, E. (2014). People, places, and public policy. *Annual Review of Economics, 6*(1), 629–662.

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

[13] Busso, M., Gregory, J., & Kline, P. (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review, 103*(2), 897–947.

[14] Gibbons, S., Overman, H., & Pelkonen, P. (2019). Area-based initiatives in the United Kingdom: Evidence from enterprise zones. *Economic Geography, 95*(5), 453–480.

[15] Doran, J., Jordan, D., & Elhorst, P. (2018). Virtual special issue on regional inequality. *Spatial Economic Analysis, 13*(4), 383–386. https://doi.org/10.1080/17421772.2018.1514095.

[16] Gluschenko, K. (2018). Measuring regional inequality: To weight or not to weight? *Spatial Economic Analysis, 13*(1), 36–59. https://doi.org/10.1080/17421772.2017.1343491.

[17] Goldfarb, A., & Tucker, C. (2019). Digital economics. *Journal of Economic Literature, 57*(1), 3–43.

[18] Brakman, S., Garretsen, H., & Marlet, G. (2021). The spatial distribution of the digital economy in Europe. *Cambridge Journal of Regions, Economy and Society, 14*(3), 645–660.

[19] De Stefano, T., Distante, F., & Parente, P. (2023). The geography of digital divides in Europe. *Regional Studies, 57*(8), 1421–1436.

[20] Capello, R., Lenzi, C., & Panzera, E. (2025). Cities as catalysts of wage inequalities from digitalisation. *Spatial Economic Analysis*. https://doi.org/10.1080/17421772.2025.2533478.

[21] Hjort, J., & Poulsen, J. (2019). The arrival of fast internet and employment in Africa. *American Economic Review, 109*(3), 1032–1079.

[22] Akerman, A., Gaarder, I., & Mogstad, M. (2015). The skill complementarity of broadband internet. *Quarterly Journal of Economics, 130*(4), 1781–1824.

[23] LeSage, J., & Pace, R. K. (2009). *Introduction to spatial econometrics*. Boca Raton: CRC Press.

[24] Elhorst, J. P. (2014). *Spatial econometrics: From cross-sectional data to spatial panels*. Berlin: Springer.
[25] Halleck Vega, S., & Elhorst, J. P. (2015). The SLX model. *Journal of Regional Science, 55*(3), 339–363.
[26] Bond-Smith, S., Corrado, L., Felsenstein, D., & Elhorst, P. (2024). Spatial macroeconomics. *Spatial Economic Analysis, 19*(3), 273–286. https://doi.org/10.1080/17421772.2024.2364527
[27] Barbero, J., Christou, T., Crucitti, F., García Rodríguez, A., Lazarou, N.-J., Monfort, P., & Salotti, S. (2024). A spatial macroeconomic analysis of the equity-efficiency trade-off of the European cohesion policy. *Spatial Economic Analysis, 19*(3), 394–410. https://doi.org/10.1080/17421772.2024.2306948
[28] Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics, 225*(2), 200–230.
[29] Kleibergen, F., & Paap, R. (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics, 133*(1), 97–126. https://doi.org/10.1016/j.jeconom.2005.02.011
[30] Stock, J. H., & Yogo, M. (2005). Testing for weak instruments in linear IV regression. In D. W. K. Andrews & J. H. Stock (Eds.), *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg* (pp. 80–108). New York: Cambridge University Press.
[31] Montiel Olea, J. L., & Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics, 31*(3), 358–369. https://doi.org/10.1080/00401706.2013.806694.
[32] Baron, R. M., & Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology, 51*(6), 1173–1182. https://doi.org/10.1037/0022-3514.51.6.1173.
[33] Xie, D., Wei, W., & Li, Y. (2026). Data and computational power economy: A general equilibrium theory of the digital economy. *Management World, 2026*(3), 37–56.

Appendix A. A stylized complementary-asset framework

Appendix A presents a stylized framework for organizing the interpretation of spatial heterogeneity. Let $\boldsymbol{i}$ index cities, $d_i \in 0, 1$ indicate PDAO adoption, and K_i denote the local endowment of complementary assets. Consider the following relationship (Equation (A1)):

$$V_i\left(d_i, K_i\right) = K_i^{\alpha}[a_0 + a_1 d_i s\left(K_i\right)]^{\beta} exp\left(\varepsilon_i\right) \tag{A1}$$

In Equation (A1), V_i denotes local enterprise activity, α and β are positive elasticities, a_0 captures baseline productive capacity, a_1 captures the incremental component associated with PDAO adoption, and ε_i is a city-specific disturbance. The activation function is defined as (Equation (A2)):

$$s(K_i) = [1 + exp - \gamma(K_i - K*)]^{-1} \tag{A2}$$

Equation (A2) defines a smooth activation function, where γ > 0 governs its slope and K^* denotes the level of complementary assets around which the incremental component becomes more pronounced. Substituting Equation (A2) into Equation (A1) illustrates how the estimated association between PDAO adoption and local economic activity may vary with complementary-asset endowments.

This framework illustrates how the estimated association between PDAO adoption and local economic activity may vary with complementary-asset endowments. It is not structurally estimated and does not constitute evidence that a distinct PDAO-specific causal mechanism has been separately identified.

Illustrative implication 1 (Heterogeneity Magnitude). Cities with $\boldsymbol{K_i}$ << K^* may exhibit estimated effects close to zero, whereas cities with $\boldsymbol{K_i}$ >> K^* may exhibit effects closer to $\beta a_1/a_0$. The Tier-1 to Other Tiers ratio of 18.4 is an observed descriptive difference and represents a pattern consistent with the stylized framework under selected parameter values, not a model-identified structural ratio.

Illustrative implication 2 (Monotonicity). The framework illustrates how estimated effects may increase with $\boldsymbol{K_i}$, a pattern consistent with the observed Tier-1 > New Tier-1 > Tier-2 > Other Tiers gradient.

Illustrative implication 3 (Average Estimated Effect). The full-sample ATT is a weighted average across the K distribution and may be substantially smaller than the maximum-*K* effect when the bulk of cities is concentrated below K^*. This implication is consistent with the descriptive difference between the full-sample estimate of 0.025 and the Tier-1 estimate of 0.317.

Appendix B. Treatment coding and spatial anchor controls

The four provincial-level direct-controlled municipalities—Beijing, Tianjin, Shanghai, and Chongqing—are incorporated into the sample according to a uniform standard based on whether treatment status can be cleanly defined within the 2020–2024 sample window, not according to administrative tier or representativeness. Shanghai and Chongqing satisfy this standard: neither municipality promulgated an effective municipal-level PDAO policy document during the sample period, and their treatment status (treat = 0) is unambiguously coded as never-treated. Beijing and Tianjin do not satisfy this standard for distinct reasons. Beijing generated large-scale public data procurement and platform operation activity with documented commercial contracts during the sample period, but in the absence of a formally enacted municipal PDAO regulation within the window, the treatment indicator cannot be cleanly defined as either zero or one. Tianjin promulgated a formal PDAO policy document in late 2024, at the end of the sample window, creating ambiguous treatment timing that would contaminate the parallel-trends identification strategy. Both cities are excluded from the estimation sample by the same treatment-coding criterion applied to all cities, without special treatment of any administrative tier. The three conventional Tier-1 cities included in the Tier-1 subsample analysis are Shanghai, Guangzhou, and Shenzhen.

Two municipalities in the never-treated control group carry additional methodological significance in the spatial econometric analysis. Shanghai is the largest control city within a 400-kilometer radius of seven treated cities in the Yangtze River Delta (Nanjing, Suzhou, Hangzhou, Yangzhou, and three others), providing a scale-comparable counterfactual spatial unit. Chongqing occupies the analogous role in the Chengdu–Chongqing cluster, serving as the largest untreated city within 400 kilometers of seven treated cities including Chengdu, Luzhou, and Mianyang. Their inclusion helps maintain comparability and precision in the spatial estimates. Robustness checks excluding Shanghai and Chongqing produce lower precision for ρ in some years (e.g., 2022: ** → *), a pattern consistent with their usefulness as spatial controls but not evidence of a distinct mechanism role.

Appendix C. Spatial Durbin decomposition and SDM specification details

Table C1 reports the LeSage and Pace [23] direct, indirect, and total effects decomposition. The estimated direct effect associated with PDAO authorization is 0.027 ($p < 0.001$), quantitatively consistent with the TWFE-DiD estimate of 0.025. The estimated indirect effect across other cities mediated through the spatial weight structure is −0.064 (SE = 0.089, $p = 0.469$), which is not statistically significant. The decomposition therefore does not detect a statistically significant indirect effect within the sample and spatial weight structures examined. This result does not rule out redistribution or other spatial mechanisms. Comparisons with studies of digital infrastructure and high-speed internet access, such as Hjort and Poulsen [21], are contextual rather than tests distinguishing causal mechanisms.

The spatial autoregressive parameter ρ is large and statistically significant across all SDM specifications (Table 4), indicating substantial spatial dependence in enterprise stock. This finding does not alter the interpretation that the policy-related indirect effect is not statistically detectable in the specifications examined.

Table C1. LeSage–Pace direct, indirect, and total effects of PDAO (SDM, full specification)

Effect type	Coefficient	SE	*p*-value	Interpretation
Direct effect (did)	+0.0271***	(0.0065)	< 0.001	*Estimated local direct effect*
Indirect effect (did)	−0.0641	(0.0886)	0.469	*No statistically detectable indirect effect*
Total effect (did)	−0.0371	(0.0904)	0.682	*Total effect not statistically significant*

Notes: Monte Carlo simulation standard errors (250 replications) from xsmle effects option. Direct effect is the average diagonal element of the partial derivative matrix; Indirect is the average off-diagonal sum. The estimated indirect effect ($p = 0.469$) is not statistically significant and provides no evidence of a systematic cross-city indirect effect within the sample and spatial weight structures examined. This result does not rule out redistribution or other spatial mechanisms. Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.

Appendix D. Detailed robustness checks

Appendix D1. Alternative outcome variables

Replacing the main outcome with the high-technology firm ratio (hightech_ratio, defined as high-technology enterprise stock as a percentage of total enterprise stock) yields an estimate of 0.407 (cluster-robust SE =

0.152, $p = 0.008$, $N = 1{,}216$). PDAO authorization is associated with an approximately 0.41 percentage-point higher share of high-technology firms within treated cities. This result provides supplementary descriptive evidence consistent with compositional change in local enterprise activity. It is also consistent with a complementary-asset interpretation, but it does not establish a causal mediation channel or a separately identified threshold mechanism.

Log digital-economy enterprise stock and log total patent count do not yield statistically significant estimates at the full-sample level ($p = 0.225$ and $p = 0.597$, respectively). Given that patent stock captures long-run accumulation and the panel spans five years, limited statistical precision may contribute to these results. The null estimates do not demonstrate that the corresponding effects are absent.

Appendix D2. Excluding Tier-1 cities and sensitivity to imputation

Excluding the three Tier-1 cities produces a non-Tier-1 average treatment effect of 0.025 (cluster-robust SE = 0.012, $p = 0.045$, $N = 248$ cities). This estimate is quantitatively similar to the full-sample average treatment effect of 0.025. The non-Tier-1 estimate (0.025) and the Tier-1 estimate (0.317) imply a descriptive ratio of 12.7. This comparison is consistent with concentration of the estimated heterogeneity at the top of the urban hierarchy, but the Tier-1 estimate is based on only three cities and the ratio should not be interpreted as a stable structural parameter.

The unbalanced panel, constructed by removing imputed 2020 baseline observations for cities in Fujian (province_id = 22), Henan (province_id = 16), and Hunan (province_id = 20) provinces, yields a full-sample treatment effect of 0.029 (SE = 0.012, $p = 0.016$, $N = 1{,}216$). The Tier-1 subsample is unaffected because Shanghai, Guangzhou, and Shenzhen are not located in the three affected provinces; the Tier-1 point estimate remains 0.317.

Appendix D3. The COVID-19 confounder

The COVID-19 robustness analysis is conducted at the full-sample level. The Tier-1 subsample cannot be assessed under the drop-2020 specification because Shenzhen (treat_year = 2020) and Guangzhou (treat_year = 2021) both have did = 1 for all post-2020 observations; removing 2020 eliminates all pre-treatment variation for both treated Tier-1 cities, rendering identification mechanically impossible. In the event-study results reported in Figure 4, no statistically significant pre-treatment differences are detected; this evidence should not be interpreted as proving the parallel-trends assumption.

For the full sample, excluding 2020 observations yields a treatment effect of 0.007 (SE = 0.011, $p = 0.544$, $N = 1{,}004$). The g2020 and g2021 cohorts collectively lose all pre-treatment observations when 2020 is dropped, removing 17 treated cities from effective identification. Only the g2022 and g2023 cohorts remain identified under this specification, and their cohort-specific ATTs are substantially smaller than the g2021 cohort (0.028 and 0.004 respectively; Table 6). The attenuation may therefore reflect the cohort composition of the identified subsample and should not be interpreted as evidence that the effect is absent.

Introducing province-by-2020 interaction fixed effects to absorb differential provincial COVID-19 shocks yields a treatment effect of 0.020 (SE = 0.010, $p = 0.056$, $N = 1{,}252$). The direction is preserved and the magnitude remains close to the baseline (0.025). This result provides supplementary evidence that province-level COVID-19 heterogeneity does not appear to account for the observed pattern in this specification, but it does not rule out other pandemic-related confounding.

Appendix D4. Wild cluster bootstrap for small subsamples

For the Tier-1 subsample (G = 3 cities, df = 2), conventional inference is unreliable under small-G asymptotics (conventional $p = 0.252$). This stuay therefore reports empirical permutation inference based on 999 random relabellings of the Tier-1 label within the full 251-city pool. The fraction of permuted effects exceeding the observed 0.317 in absolute value yields permutation $p = 0.038$. The permutation p-value is the preferred inference for the three-city Tier-1 subsample, while its interpretation depends on the current relabelling design and the associated exchangeability assumption.

Appendix D5. Placebo tests

A placebo test randomly reassigns treatment timing (treat_year) within treated Tier-1 cities across 200 replications, assigning each city a random treatment year drawn uniformly from 2020–2023. The fraction of placebo Tier-1 estimates exceeding the observed 0.317 in absolute value yields an empirical p-value of 0.040. The comparison with the permutation p-value of 0.038 provides additional reassurance that the baseline estimate is not mechanically reproduced under the placebo assignment, but it does not by itself establish a causal interpretation.

Appendix D6. Alternative classification systems

Binary core-versus-non-core split: core effect 0.045 ($p = 0.012$), non-core effect −0.003 ($p = 0.834$). Integrated complementary asset score: top-tercile effect 0.084 ($p = 0.001$), middle 0.012 ($p = 0.412$), bottom −0.014 ($p = 0.298$). Heterogeneity is robust to alternative classification schemes.

Appendix D7. Leave-one-out Tier-1 robustness

Table D1 presents leave-one-out estimates under a simplified specification (city and year fixed effects only), because two-city panels with six control variables yield df = 0. The baseline simplified estimate is 0.118 (p = 0.003). Dropping Shanghai yields 0.108 ($p = 0.033$); dropping Shenzhen yields 0.127 ($p = 0.029$). Dropping Guangzhou leaves Shanghai (never-treated) and Shenzhen (treated), making did perfectly collinear with the city and year fixed effect structure and preventing identification. These estimates provide suggestive evidence that the pattern is not driven solely by Shanghai or Shenzhen, but the Tier-1 group contains only three cities, the Guangzhou exclusion is not estimable, and the simplified and baseline specifications are not identical estimands. The exercise therefore does not eliminate small-sample uncertainty. Beijing and Tianjin remain excluded according to the treatment-coding standard described in Section 4.1.

Table D1. Leave-One-Out Tier-1 robustness estimates

Specification	Cities (G)	β (DiD)	HC1 SE	p-value	df / identification status
Baseline (3 cities, simplified spec)	3	+0.1175***	(0.0257)	0.003	df = 7
Drop Shanghai (city_id = 6)	2	+0.1083**	(0.0288)	0.033	df = 3
Drop Guangzhou (city_id = 103)	2	omitted	—	—	did collinear with FE
Drop Shenzhen (city_id = 179)	2	+0.1267**	(0.0323)	0.029	df = 3

Notes: Leave-one-out estimates use a simplified specification (city and year fixed effects only, no additional controls) because two-city panels with six control variables yield df = 0 and produce degenerate estimates. Baseline simplified estimate (0.118) is lower than the full-specification estimate (0.317), and the two specifications should not be treated as identical estimands. The estimable leave-one-out results provide suggestive evidence but do not eliminate the small-sample risk associated with a three-city Tier-1 group. When Guangzhou is excluded, the remaining two cities (Shanghai, treat = 0; Shenzhen, treat = 1) produce a did variable perfectly collinear with the city and year fixed effect structure, preventing identification. Beijing and Tianjin are excluded from the Tier-1 group by the treatment-coding criterion of Section 4.1. Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.

Appendix D8. Alternative spatial weight matrices

Table D2 presents SDM estimates under four spatial weight matrices. The inverse-distance matrix (W_1) captures distance-based spatial connections; provincial contiguity (W_2) captures administrative adjacency; GDP-similarity (W_3) captures similarity in economic scale; and KNN-25 (W_4) is the baseline. The comparison assesses the sensitivity of the estimated spatial treatment coefficient to alternative weight structures rather than discriminating among causal mechanisms.

Table D2. Spatial durbin model estimates under alternative weight matrices

Weight matrix (economic rationale)	ρ	p (ρ)	θ on $W{\cdot}did$	p (θ)	Spillover evidence
Inverse-distance, 400km truncation (agglomeration decay at HSR 2-hr radius)	+0.7217	< 0.001	+0.0001	0.994	*No statistically detectable θ*
Provincial contiguity (institutional policy transmission)	+0.7083	< 0.001	−0.0105	0.348	*No statistically detectable θ*
GDP-similarity (strongest test of data-enclave hypothesis)	+0.2948	< 0.001	+0.0301*	0.077	*Marginal positive θ*
KNN-25 (baseline: general spatial autocorrelation)	+0.7264	< 0.001	−0.0354	0.118	*No statistically detectable θ*

Notes: SDM estimates on the full balanced panel (N = 1,255), full specification durbin(did ×X). All four W matrices were implemented in Stata (Experiment A fixed). ρ is the spatial autoregressive coefficient on $W \cdot ln(enterprise\ stock)$, and θ is the coefficient on $W \cdot did$. Across these alternative structures, the estimates provide sensitivity evidence on whether θ is statistically detectable; they do not rule out redistribution or identify a distinct spatial mechanism. Significance: ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels respectively.