

Artificial Proximity: Policy Diffusion in China's East–West Pairings

Juncheng Lyu

King's College London

Abstract

China's East–West Pairing programme has centrally assigned more developed eastern regions to support poorer western regions since 1996. This paper asks whether pairing transfers the richer partner's policies and practices to the poorer partner. Combining difference-in-differences analysis of a policy-text similarity index from 2006 to 2022 with a fieldwork case study, we find no detectable average convergence in binding-norm policy text after designation despite the channels for exchange created by the programme. The case study shows that these channels transfer working practices and create personal ties, while partner dissimilarity and accountability pressures may help explain why these exchanges do not appear as convergence in codified policy.

1. Introduction

Policy diffusion research has largely studied ties that governments form for themselves, between neighbours and peers (F. S. Berry & Berry, 2018; Shipan & Volden, 2008), and between the jurisdictions officials move between (Yi et al., 2018). Recent large-scale studies show how policy travels through China's own system (Luo et al., 2025; S. Wang & Yang, 2025). China's East–West Pairing programme (EWP) creates the connection administratively. Launched in 1996, it pairs more developed eastern regions with provinces in the western interior to support poverty reduction.

The centre assigns prefecture-level partnerships across provinces and levels of development. Cadre exchanges, joint offices, and routine contact coordinate the partners' poverty-reduction work rather than promote policy adoption. These arrangements nonetheless create the conditions through which policies and practices might spread between partners. Learning and imitation are established mechanisms of diffusion (Shipan & Volden, 2008), and the networks formed by transferred and seconded cadres provide a third channel here (Yi et al., 2018).

Studies of China's pairing mechanisms show that central mandates can create connections between local governments, but they do not empirically test whether those connections produce policy diffusion. Designed networks can also generate self-organising learning ties between paired cities (Yi et al., 2024), while work on the East–West programme documents the structure of the assigned partnership without evaluating whether the partners' policies converge (Z. Wang, 2025). Q. Wei and Yang (2025)'s qualitative study of the Shenzhen–Ganzhou partnership, which belongs to the same programme family, leaves the evaluation of outcomes to later research. X. Li et al. (2025) evaluate Zhejiang's within-province pairing and find effects on investment flows and fiscal outcomes, but not on household income.

Evidence from Chinese policy text also shows that borrowing is more common between similar jurisdictions, although this work examines diffusion among cities without imposed pairings (Luo et al., 2025). Research on China's policy process further shows that local models spread when the centre recognises and promotes them (Heilmann, 2008; J. Wang & Yu, 2026). What remains unclear is whether the centre can also produce horizontal diffusion simply by assigning a connection between local governments, and how the cadre evaluation system shapes that process.

This paper asks whether a connection created to coordinate poverty reduction also moves policy between the partners. It estimates whether the recipient's binding policy text moves towards its assigned partner's and uses fieldwork to examine what passes through the programme's transfer channels. The paper applies a difference-in-differences design to a policy similarity index, comparing each assisted prefecture's binding-norm corpus with its assigned partner's earlier corpus from 2006 to 2022. The design accounts for pre-existing resemblance between the partners and for the possibility that text converges through declared adoption rather than actual borrowing. The quantitative analysis is combined with a case study of the Shenzhen–Guangxi partnership based on two phases of fieldwork and the programme's administrative record.

The designation produced no detectable average convergence across different outcome measures and estimation assumptions. The assisted prefectures' policy texts did not move toward their partners' in the six years after 2016. However, the case study shows that the channels did produce meaningful transfer in the form of more advanced working practices and personal ties that outlast the formal pairing. Two barriers may explain why this transfer does not appear as policy convergence. The partners differ substantially, making direct policy adoption less likely. Local officials are also assessed against centrally defined targets that do not reward learning from the distinctive policies of their assigned partner. The two barriers reinforce each other as dissimilarity makes borrowing riskier, while the evaluation system provides little incentive to take that risk.

The remainder of the paper is structured as follows. Section 2 reviews the relevant literature and conceptual framework. Section 3 specifies the East–West Pairing programme and the channels it establishes. Section 4 describes the data. Section 5 constructs the policy similarity index. Section 6 presents the empirical model and identification strategy. Section 7 reports the estimates, their pre-trend diagnostics, and the descriptive comparisons. Section 8 presents the qualitative evidence from the Shenzhen–Guangxi case study. Section 9 concludes and discusses the policy implications.

2. Literature review and conceptual framework

Policy diffusion occurs when one government's choices are shaped by those made elsewhere (F. S. Berry & Berry, 2018). The literature distinguishes horizontal diffusion between jurisdictions at the same level from vertical diffusion through mandates or coercion (Shipan & Volden, 2008). A designed pairing creates three possible channels of horizontal diffusion. First, one government may learn from another's policy record. Its choices may then move towards those of its partner (Shipan & Volden, 2008), although this is more likely when the two jurisdictions face similar conditions. Second, a jurisdiction may imitate a prominent peer in order to resemble it, with less concern for whether the policy fits local conditions. The model is often a more developed

jurisdiction (Lim & Horesh, 2016; Shipan & Volden, 2008). Third, diffusion may occur through networks of officials. Officials who move between jurisdictions can carry policy ideas and practices with them. The resulting tie may produce convergence, but this again depends on institutional similarity between the origin and destination (F. S. Berry, 2023; Yi et al., 2018). Normative pressure and coercion move jurisdictions towards a central or field-wide template rather than towards a particular partner. The research design absorbs convergence of this kind (F. S. Berry & Berry, 2018; Shipan & Volden, 2008). Competition operates through rivalry for mobile residents, firms, and revenue, but the EWP's paired prefectures do not compete within the same market (F. S. Berry & Berry, 2018; Shipan & Volden, 2008).

The three partner channels all require a connection between the jurisdictions. Top-down mandated ties, including counterpart assistance and cadre transfers, can produce self-organising interlocal learning networks among Chinese cities (Yi et al., 2024). However, Yi et al. (2024) do not examine whether these networks lead connected cities to adopt more similar policies. Policy diffusion in China more often operates vertically than horizontally. The centre selects local models and mandates their expansion, producing experimentation under hierarchy rather than spontaneous diffusion among similar jurisdictions (Heilmann, 2008). A centrally mandated pairing such as the EWP therefore provides an appropriate case for examining whether an artificial horizontal connection produces diffusion.

An installed channel does not by itself produce convergence. This paper examines two barriers that can limit it in the channels the EWP installs. The first is dissimilarity, since the effect of a connection depends on the similarity between the jurisdictions it links. Glick (2014) argues that a government would rather learn from another that is similar to itself, and applies this reasoning to policy content. Similar jurisdictions may mimic one another by copying policies closely. Dissimilar jurisdictions are more likely to modify a borrowed policy to fit their own circumstances. Mimicking produces near-identical text, whereas modification produces partial rather than complete textual similarity. Policy transfer also changes policies as they move. Borrowing jurisdictions select and recombine particular elements rather than reproducing the original policy intact (McCann & Ward, 2013). Volden (2016) finds that learning from other US states' experiments disappears once the two governments become sufficiently distant ideologically. Learning is also weaker when the borrowing government lacks the capacity to examine other jurisdictions' experience.

Empirical studies of borrowing between jurisdictions provide similar evidence. Among Chinese prefectures, imitation is strongest between jurisdictions of the same administrative rank (Xiong et al., 2025). Economically similar jurisdictions also tend to have more similar policy portfolios where diffusion occurs without an imposed connection (Luo et al., 2025). However, these studies do not examine whether an imposed connection between dissimilar jurisdictions produces convergence in policy text.

A second barrier comes from the evaluation system to which officials are accountable. Local officials concentrate their effort on targets that evaluations rank and reward, while retaining wider discretion over tasks that receive less attention (O'Brien & Li, 1999). Borrowing a policy designed for another jurisdiction may therefore carry administrative and political risk when the evaluation system provides no corresponding reward. Where the centre signals that a reform is wanted, adoption accelerates even without fiscal inducement (Zhu & Zhang, 2016). J. Wang and Yu (2026) find two mechanisms operating together in China's provinces. Central recognition rewards generating what the centre will endorse and discourages borrowing what it has not. Peer recognition works alongside it, as governments borrow policies they judge exem-

plary and useful for local needs. Borrowing of that second kind has declined as central authority has strengthened. The two barriers compound as dissimilarity raises the risk that the cautious official then refuses. Whether convergence occurs depends on the relative strength of the barriers and the diffusion channel.

3. Institutional setting: the East–West Pairing programme and its installed channels

China’s East–West Pairing programme, established in 1996, is one of the central government’s main instruments for narrowing the gap between the richer eastern and poorer western regions. The 2016 directive places it within the national campaign to eliminate extreme rural poverty and frames it as cooperation on terms of complementary advantage and mutual benefit.¹ The channels through which policy might diffuse are a byproduct of building the cooperation rather than a deliberate aim. Partnerships are therefore established between more developed and less developed regions. The paired regions have very different socio-economic characteristics and are at different levels of development. Across the 40 prefecture pairs designated in 2016, the median donor-to-recipient ratio in GDP per capita was 4.8 in 2006, 3.8 in 2015, and 3.0 in 2022. The primary sector accounted for 17.4% of output in the median recipient in 2015, compared with 1.7% in the median donor (Table S16). The centre names specific prefecture-to-prefecture pairs in the directive itself, many of them ethnic minority autonomous prefectures. It also pairs each eastern province with a western one, while allowing the eastern province to select the remaining prefecture-level pairings. The arrangement is therefore a centrally engineered, top-down pairing rather than an organic horizontal tie between similar jurisdictions (Z. Wang, 2025).

The policy similarity index is constructed from each prefecture’s full corpus of binding norms rather than from poverty-related documents alone, for three reasons. First, poverty policy follows a central template passed down a fixed responsibility chain (Liu, Xin, & Yi, 2024), and the cooperation plans that run the programme are issued jointly by the two provinces. Convergence in poverty-related policy text would therefore confound joint issuance and the central template with horizontal diffusion. Second, the assisted prefectures issue too few documents in that domain alone to sustain the measure. Third, the channels themselves are not poverty-specific, since personnel support is mandated across education, health, science and technology, culture, and social work.

Fieldwork in the Shenzhen–Guangxi partnership shows how these channels operate. Interview evidence is cited by code, with the descriptor, location, and date for each code given in the online supplemental data (Table S15).

The eastern donor seconded a cadre to serve for one to three years as deputy secretary of the recipient county’s Party committee, with a seat on its standing committee. A coordinating office in the donor district acts as a node across its own functional bureaus and as a bridge to the recipient. On the recipient side, a standing Guangdong–Guangxi cooperation office runs the partnership at the provincial, prefecture, and county levels. It is jointly formed by the seconded donor work team and the recipient authorities (J.

¹General Office of the Central Committee of the Communist Party of China and General Office of the State Council, *Guanyu jingyibu jiaqiang dongxibu fupin xiezuo gongzuo de zhidao yijian* [Guiding opinions on further strengthening East–West poverty alleviation cooperation work] (Zhongbanfa [2016] No. 69), 7 December 2016, https://www.gov.cn/gongbao/content/2017/content_5156730.htm.

Wei, 2021).² Leading officials and delegations from the donor visit each year, and the secondment of teachers, doctors, and technical staff is mandatory and routine. Each district commits a fixed annual fiscal transfer to its partner county,³ released to the county treasury against project milestones.

4. Data and sample construction

The policy document corpus is drawn from PKU Law (Beida Fabao), a comprehensive legal database commonly used in studies of Chinese legal text (W. Chen et al., 2025; Luo et al., 2025; S. Wang & Yang, 2025). Three document types are excluded from its local administrative regulations. The first is administrative work documents (*difang gongzuo wenjian*), which make up 65% of them but are administrative circulars rather than policy documents. The second is administrative licensing replies (*xingzheng xuke pifu*, 14%). The third is local judicial documents (*difang sifa wenjian*, under 1%). The 935,641 records that remain are the eligible set, and the analysis draws on those issued by the prefectures in the sample. The eligible records fall into three types. Each carries an effectiveness weight based on its position in China's legislative hierarchy, following the three-tier classification in J. Chen et al. (2026). Local regulations and ethnic autonomous regulations receive weight three, local government rules weight two, and local normative documents (*guifanxing wenjian*) weight one. The corpus construction and document-type filtering, from the raw database to the analysis sample, are detailed in the online supplemental data (Table S1).

The socioeconomic controls cover 179 prefectures over 2000–2022. Most of the data come from the China City Statistical Yearbook, supplemented by the EPS China Regional Economic Database. The China Ethnic Statistical Yearbook adds coverage for the 12 autonomous prefectures in the sample. The panel's construction is further detailed in the online supplemental data (Table S2). Of the 179 prefectures, 94 are the assisted prefectures used in the main difference-in-differences, 43 are eastern partner prefectures supplying the partner-side covariates, and 42 belong to the Northeast and Guizhou case studies. Five time-varying controls enter the preferred specification, marked in Table 3. GDP per capita captures economic development, which also shapes legislative capacity. Population captures scale, since larger prefectures issue more documents. Secondary-industry share captures economic structure. Government expenditure as a share of GDP captures the intensity of government intervention. The fifth is the number of eligible documents the assisted prefecture issues in year t . A prefecture issuing more documents has mechanically more opportunities to match, independent of policy learning. GDP per capita, population and the document count enter in logs. The count of eligible documents its partner issued in the five years before t does not enter the preferred specification. It measures the size of the benchmark pool, which the partner-by-year fixed effects absorb.

The sample is drawn from the 10 assisted provinces that received new prefecture-to-prefecture EWP pairings in the 2016 directive. Within each province, prefectures fall into three groups. The treated units are those assigned a new specific prefecture-to-prefecture partner in 2016. Prefectures that remained under an existing 1996 arrangement with the same partner are excluded, since ongoing treatment could affect

²These offices were formally constituted and named *Yue-Gui xiezuo bangongshi* in 2021, near the end of the study window; equivalent frontline cooperation work groups had performed the same role since the 2016 designation (Yuan, 2021).

³The exact figure for the partnership studied is withheld at the request of fieldwork interviewees.

their trajectory. The controls are those that never received a specific prefecture-to-prefecture partner in any round. Three provinces are excluded entirely. Ningxia and Shaanxi had been covered by province-wide partnerships since 1996 and did not change in 2016. Xinjiang operates under a separate programme.

The resulting sample has three categories (Table 1). Category A provinces had no EWP history before 2016, providing the cleanest identification of whether a new partner designation produces diffusion in the absence of a pre-existing EWP framework. The trade-off is lower statistical power. Category B adds provinces where the 2016 directive restructured prefecture-to-prefecture pairings within province-level partnerships that had existed since 1996. Because the fixed effects absorb the shared province-level framework, this category identifies whether the new prefecture-to-prefecture designation produces additional diffusion within an existing long-term partnership. Category C adds the four provinces with too few within-province controls to sustain a within-province comparison (Table 2). The full sample includes all designated partners and assisted prefectures and has the most statistical power.

The panel spans 2006–2022, which gives 10 pre-treatment years for the event-study diagnostics. Through 2005 the median assisted prefecture issued fewer than 3 documents a year, rising to 18 in 2006, so the analysis begins in 2006. The panel ends in 2022, the last year with complete control variables. Although the pairings were restructured again in 2021, the single post-restructuring year is too short for that change to be reflected in binding policy text.

Table 1. EWP treatment assignments, 2016

Province	Treated units	Eastern partner(s)	Controls	Notes
<i>Category A: clean onset</i>				
Hubei	Enshi	Hangzhou	12	
Hunan	Xiangxi	Jinan	13	
Hebei	Zhangjiakou, Baoding, Chengde	Beijing, Tianjin	8	BTH overlap
<i>Category B: intensification</i>				
Gansu	8 prefectures	Tianjin, Fuzhou, Xiamen, Qingdao	4	1996: Tianjin
Sichuan	7 prefectures	Quzhou, Jiaxing, Zhoushan, Huzhou, Lishui, Foshan	9	1996: Zhejiang
Guangxi	8 prefectures	Maoming, Zhanjiang, Zhaoqing, Jiangmen, Shenzhen	6	1996: Guangdong
<i>Category C: pooled only</i>				
Jilin	Yanbian	Ningbo	0	Northeast (2017)
Guizhou	Bijie	Guangzhou	1	Switch 2016
Yunnan	8 prefectures	Shanghai, Zhuhai	1	Shanghai since 1996
Qinghai	Xining, Haidong	Nanjing, Wuxi	0	Partner switch

Notes. The 1996 blanket partner is shown for Category B. Controls are within-province prefectures that never received a prefecture-to-prefecture partner in any round, and the Controls column counts prefectures. Hebei overlaps the Beijing–Tianjin–Hebei (BTH) coordination.

Table 2. Sample composition

Sample tier	Pairs	Treated	Controls	Clusters	T_{pre}	T_{post}	Balanced 2007–2022
Full sample	288	40	248	94	10	7	216/288 (75.0%)
Cat A pooled	49	5	44	38	10	7	47/49 (95.9%)
Hubei $\times$ Hangzhou	13	1	12	13	10	7	11/13 (84.6%)
Hunan $\times$ Jinan	14	1	13	14	10	7	14/14 (100%)
Hebei $\times$ Beijing	11	2	9	11	10	7	11/11 (100%)
Hebei $\times$ Tianjin	11	1	10	11	10	7	11/11 (100%)
Cat B + Pool	239	35	204	56	10	7	169/239 (70.7%)

Notes. Pairs, Treated, and Controls are (assisted prefecture $\times$ partner) pair counts; Clusters is the number of unique assisted prefectures. A control prefecture evaluated against several partners contributes several control pairs, so the prefecture-level control counts are 54, 33, and 21 (Full sample, Cat A, Cat B+Pool), matching Table 1. T_{pre} and T_{post} are pre- and post-2016 years on the 2006–2022 panel. The right-most column is the share of pairs surviving a balanced 2007–2022 filter, meaning a non-missing outcome in all 16 years, and is the sample used by the estimators reported in the online supplemental data.

Table 3. Descriptive statistics

	N	Mean	SD	Min	Max
<i>Outcomes — bag-of-words</i>					
log(sum)	4,717	3.71	1.28	−0.76	6.27
psi_all	4,717	0.0586	0.0115	0.0212	0.1215
max	4,717	0.3695	0.0713	0.0895	0.7077
topa01	4,717	0.2756	0.0475	0.0844	0.5805
topa10	4,717	0.1716	0.0286	0.0719	0.4145
<i>Outcomes — Qwen3</i>					
log(sum)	4,717	6.09	1.34	1.61	8.47
psi_all	4,717	0.6180	0.0198	0.5500	0.6949
max	4,717	0.8122	0.0276	0.6584	0.9126
topa01	4,717	0.7702	0.0207	0.6492	0.8554
topa10	4,717	0.7150	0.0181	0.6356	0.8222
<i>Assisted prefecture (panel-invariant across vectorisations)</i>					
GDP per capita (yuan) [†]	4,896	34,434	22,883	2,767	145,545
Population (10,000) [†]	4,896	407.3	246.3	17.61	1,572
Primary industry share	4,896	0.167	0.0721	0.0117	0.417
Secondary industry share [†]	4,896	0.444	0.117	0.128	0.844
Tertiary industry share	4,896	0.387	0.103	0.114	0.727
Government expenditure / GDP [†]	4,896	0.230	0.138	0.0584	1.027
Fiscal revenue per capita (yuan)	4,896	2,187	2,079	80.78	21,052
Fiscal expenditure per capita (yuan)	4,896	6,402	3,845	570.6	26,281
Retail sales per capita (yuan)	4,896	11,871	10,015	619.6	85,811
Rural income per capita (yuan)	3,936	8,922	5,514	1,456	30,931
Number of documents issued [†]	4,896	62.13	114.9	0	1,191
<i>Partner prefecture</i>					
GDP per capita (yuan)	4,896	79,865	41,221	11,505	203,489
Population (10,000)	4,896	559.6	351.1	92.63	1,504
Primary industry share	4,896	0.0647	0.0602	0.0003	0.286
Secondary industry share	4,896	0.448	0.0916	0.158	0.656
Tertiary industry share	4,896	0.488	0.105	0.306	0.839
Government expenditure / GDP	4,896	0.142	0.0556	0.0554	0.356
Fiscal revenue per capita (yuan)	4,896	12,966	14,370	366.7	81,467
Fiscal expenditure per capita (yuan)	4,896	17,207	17,190	806.5	109,378
Retail sales per capita (yuan)	4,896	42,391	31,203	4,150	165,961
Rural income per capita (yuan)	3,010	13,474	6,871	3,868	41,934
Benchmark pool size (five-year window)	4,896	1,461	1,631	8	7,784

Notes. Statistics over the full estimation sample of 288 prefecture–partner pairs across 2006–2022 (4,896 pair-years; outcomes are missing for 179, leaving 4,717). The five reported continuous outcomes are shown under bag-of-words (primary) and Qwen3 (robustness). [†] Control in the preferred specification. GDP per capita enters as the log of GDP over registered population; the row reports the published per-capita series.

5. Measuring policy similarity

The outcome variable is the Policy Similarity Index (PSI), a clause-level text similarity measure between an assisted prefecture’s policy documents and its designated eastern partner’s documents. It is a continuous measure of policy content convergence beyond the binary “adopted or not” coding (W. Chen et al., 2025; Linder et al., 2020). Its construction follows W. Chen et al. (2025) with modifications for the EWP setting, where the unit of analysis is the prefecture-year rather than the document pair.

The index is built in three steps: clause preprocessing, document-pair best-match scoring, and legal-authority-weighted aggregation to the prefecture-year. Each document is decomposed into clauses. The clauses are then segmented into Chinese word tokens using the jieba library, supplemented by an EWP-specific user dictionary. A customised stopwords list removes Chinese function words and governance boilerplate. The remaining tokens form the bag-of-words (BoW) vector for each clause, and the online supplemental data (Section 2) present the preprocessing procedure.

Each outcome is computed under two vectorisations, BoW (primary) and Qwen3 dense embeddings (robustness). BoW represents lexical similarity, and Qwen3 represents dense-semantic similarity. Qwen3’s semantic cosine rewards shared genre more than shared content in a corpus that is uniform in bureaucratic tone and phrasing, so a high Qwen3 score alone is not evidence of convergence (Ethayarajh, 2019) (online supplemental data, Table S6).

For prefecture i in year t , let Z_{it} be the eligible documents it issued in year t and B_{jt} the eligible documents its designated partner j issued in the five years before t . The strict temporal ordering rules out contemporaneous learning. The assisted prefecture i is the focal side of the comparison. The index scores its documents and clauses against the partner’s earlier documents, which provide the benchmark. For each focal–benchmark document pair, the cosine is computed between every focal clause and every benchmark clause. Each focal clause then takes its highest cosine against any benchmark clause as its best available match. The document-pair score is the mean of these per-clause maxima,

$$s_m(z, b) = \frac{1}{|z|} \sum_{c_z \in z} \max_{c_b \in b} \cos(c_z, c_b), \quad (1)$$

where $|z|$ is the number of clauses in focal document z . This best-match aggregation follows a borrowing logic. A focal clause is assumed to draw on at most one source clause, so the signal is its closest benchmark clause rather than a flat mean over all clause pairs. The same max-then-aggregate construction is common in natural-language processing (T. Zhang et al., 2020). A threshold variant following W. Chen et al. (2025) scores the share of a document’s clauses whose best match exceeds a cosine threshold of 0.5. It records whether each clause has a close match rather than how close it is. Its construction is detailed in the online supplemental data (Table S4).

The document-pair scores are aggregated first across the partner’s benchmark pool B_{jt} into a document-level score $a_{\bullet}(z)$, and then across the focal prefecture’s documents into a legal-authority-weighted mean,

$$\text{PSI}_{ijt} = \frac{\sum_{z \in Z_{it}} w_z a_{\bullet}(z)}{\sum_{z \in Z_{it}} w_z}, \quad (2)$$

where $w_z \in \{1, 2, 3\}$ is the document's legal-authority weight. The benchmark pool is a rolling five-year window rather than a cumulative one, following Arts et al. (2021). The window keeps the benchmark to the partner's current policy stock, where a cumulative pool would treat a document issued a decade earlier as equally available to borrow. Prefecture-years with no eligible focal documents are coded as missing rather than zero, because a zero would imply policy was issued with zero similarity.

The PSI measures convergence between the focal prefecture's binding-norm corpus and the partner's prior corpus. This is also the tier at which the partnership itself produces formal cooperation documents, since stationed eastern cadres take part in drafting them (Y. Chen & Ling, 2023). The partnership therefore generates codified text, and the actions it coordinates are mainly shaped locally rather than copied from the partner (Z. Wang, 2025). The channel transfers content that is largely operational rather than codified policy text. Codified text captures intended policy, while implementation may diverge through bureaucratic discretion (Linder et al., 2020, pp. 549–550). The study therefore estimates formal, codified textual diffusion specifically.

When a government borrows from a named model, its policy documents usually make that reference explicit. The parliamentary report on India's zone legislation announced the Chinese model by name, yet what was built under the Act turned out fundamentally different. In such cases, convergence in policy documents records an attempt to adopt, rather than what the attempt actually delivers (Goodburn & Knoerich, 2022). Governments may write a model into official documents when adopting it is itself the aim, and when citing the model carries political or symbolic value (Lim & Horesh, 2016). The main purpose of the EWP is poverty reduction and development rather than policy diffusion. Adopting the partner's policies is therefore not rewarded in its evaluation system, and a declaration of adoption has limited benefit. The index captures codified adoption itself, and the estimated effect covers an attempt to adopt as well as the act, although highly modified or weakly expressed borrowing may be harder to detect. Where the centre runs a policy experiment across several localities, officials with stronger career incentives issue documents on that experiment more textually distinct from those of the other participants (S. Wang & Yang, 2025), so credit-seeking can push policy text apart even where the underlying policy is shared.

Transfers that do not enter codified policy fall outside the measure and are examined through the qualitative evidence. Over the longer term, such transfer may still enter binding policy text, but the analysis observes only six years after the designation. Formal policy text comes at the end of a multi-step process of codification. Project adaptation is followed by internal working documents, then binding prefecture-level norms, and finally linguistic convergence with the partner's corpus. The index observes only the binding norms at the end of that chain, since the working documents are excluded by design (Heffer & Schubert, 2023). Some operational transfer occurs below the prefecture level. It would move into prefecture-level binding norms through vertical personnel and performance channels rather than directly between partners (M. Li & Wang, 2025). That upward route is the vertical pressure the partner-by-year fixed effect absorbs. The prefecture-tier estimate is therefore an estimate on horizontal convergence toward the partner's corpus.

Five outcome-variable categories share the clause-level best-match score. They differ only in how widely they let the focal prefecture's attention range across the partner's policy archive when judging similarity (Table 4). A positive estimate for the maximum score, alongside a null result for "psi_all", would point to targeted copying of a single document rather than broad convergence. A null result across the full set of measures rules out textual convergence in any of these forms. The main outcome is "log(sum)",

where the log transform addresses the linear growth of the summation category and absorbs the proportionality into the fixed-effect structure.

For a focal document z , let $a_{\bullet}(z)$ denote its score aggregated across the partner's benchmark pool B_{jt} (of size $|B_{jt}|$), using the document-pair best-match score $s_m(z, b)$ of Equation (1). The reported categories differ only in this aggregation:

$$a_{\text{sum}}(z) = \sum_{b \in B_{jt}} s_m(z, b), \quad a_{\text{mean}}(z) = \frac{1}{|B_{jt}|} \sum_{b \in B_{jt}} s_m(z, b), \quad (3)$$

$$a_{\text{max}}(z) = \max_{b \in B_{jt}} s_m(z, b), \quad a_{\text{top}\alpha}(z) = \frac{1}{k} \sum_{b \in T_{\alpha}} s_m(z, b), \quad (4)$$

where T_{α} collects the $k = \lceil \alpha |B_{jt}| \rceil$ benchmarks with the highest $s_m(z, b)$. The pool mean a_{mean} is “psi_all”, a_{max} is “max”, and the top- α category with $\alpha \in \{0.01, 0.10\}$ gives “topa01” and “topa10”. Each outcome applies the prefecture-year aggregation of Equation (2) to its category aggregator $a_{\bullet}(z)$. The s_t threshold calculation, which produces the θ -thresholded category, together with the full outcome variants and selection diagnostics, is documented in Section 2 of the online supplemental data.

Table 4. Reported outcome variants

Variant	Family	Aggregation across the pool	Mechanism captured
log(sum)	Summation	sum, log-transformed	Accumulated textual overlap
psi_all	Pool mean	mean over all benchmarks	Broad convergence to the partner portfolio
max	Maximum	single best-matching document	Targeted copying from one document
topa01	Top- α	mean of the top 1% of benchmarks	Extreme-tail borrowing
topa10	Top- α	mean of the top 10% of benchmarks	Decile-level concentrated borrowing

Notes. The five reported scope-of-attention outcomes. Each is computed under BoW (primary) and Qwen3 (robustness). The threshold variants are defined and characterised in the online supplemental data (Tables S4 and S5).

6. Identification and estimation

The primary estimator is two-way fixed-effects (TWFE) difference-in-differences applied to a prefecture-partner-year panel. Each observation is characterised by focal prefecture i , partner prefecture j , and year t . In Category A provinces with a single eastern partner, this collapses to the standard prefecture-year structure. In Category B provinces, where multiple partners are assigned to different prefectures within the same province, each focal prefecture appears once per partner per year, so treated and control prefectures are evaluated against the same benchmark. Standard errors are clustered at the prefecture level. In the small single-province samples, joint pre-trend inference uses a wild-cluster bootstrap with Webb weights (Roodman et al., 2019).

Two specifications are used. T3 follows the bilateral-fixed-effects TWFE of Baier and Bergstrand (2007, p. 87). T6 is anchored to their theoretically motivated equation under an asymmetric adaptation (Baier & Bergstrand, 2007, pp. 88–89). Proximity, administrative custom, and leadership networks leave some prefecture pairs more alike from the start, and pair fixed effects absorb these baseline differences. The partner-by-year effect then absorbs the outside factors that move any pair's similarity. T3

specifies pair and year fixed effects with both focal-side and partner-side time-varying covariates,

$$\text{PSI}_{ijt} = \alpha_{ij} + \lambda_t + \delta (\text{Assigned}_{ij} \times \text{Post}_t) + X_{it}^{f'} \beta_f + X_{jt}^{p'} \beta_p + \varepsilon_{ijt}, \quad (5)$$

where α_{ij} absorbs prefecture-partner pair fixed effects, λ_t year fixed effects, and X^f , X^p are focal- and partner-side covariates. T6 replaces the year fixed effect with a partner-by-year fixed effect and drops the partner-side covariates it absorbs,

$$\text{PSI}_{ijt} = \alpha_{ij} + \lambda_{jt} + \delta (\text{Assigned}_{ij} \times \text{Post}_t) + X_{it}^{f'} \beta_f + \varepsilon_{ijt}. \quad (6)$$

A focal-by-year fixed effect would absorb the focal-prefecture time variation that identifies the treatment. T6 keeps only the partner side, λ_{jt} , and retains the focal-side covariates no longer absorbed. The partner-by-year fixed effect absorbs the full time path of partner-specific heterogeneity (pool growth, documentation shocks, and central templates) at any functional form. Variation across focal prefectures within the same partner-year cell then identifies the treatment effect δ . The treated prefecture is the one designated to that partner. Other prefectures in the same province serve as controls. They are scored against the same partner corpus but were never designated to that partner. The remaining specifications (T0, T1, T2, T4 and T5) are reported in the online supplemental data (Table S7), where controls are added in steps as a diagnostic.

These fixed effects also separate the diffusion mechanisms discussed earlier. Coercion and normative pressure act on every prefecture in a given year, so the year and partner-by-year effects absorb them. The coefficient δ captures whether a focal prefecture's policy text moves closer to that of its assigned partner. Two further factors that could lead to convergence without the installed channels are absent by construction. First, regions that resemble one another may converge in policy text over time because similarity creates the conditions for ties to develop. In the EWP setting, however, partners are deliberately assigned across unequal regions for poverty reduction purposes. The treated and partner prefectures are therefore not selected because they already resemble one another, and the pair fixed effect absorbs any time-invariant, unobserved factors that may have shaped the original pairing. Second, a government that is rewarded for citing a model can change its text through declaration alone. The EWP pairing does not reward such declarations of adoption. The treatment effect therefore captures the isolated effect of the channels installed by the EWP programme itself, namely learning, imitation, and agent-network diffusion.

TWFE estimates Equations (5) and (6) directly by fixed-effects OLS where the coefficient δ on $\text{Assigned} \times \text{Post}$ is the ATT under unconditional parallel trends.

The second main estimator is the fixed-effects counterfactual estimator (FEct) (Liu, Wang, & Xu, 2024). It treats treated units' post-treatment observations as missing and imputes their untreated counterfactual from a model fitted on untreated observations only,

$$\text{PSI}_{ijt}(0) = \alpha_{ij} + \lambda_{jt} + X_{it}^{f'} \beta_f + \varepsilon_{ijt}, \quad (7)$$

where $\text{PSI}_{ijt}(0)$ is the untreated potential outcome and α_{ij} , λ_{jt} , and X^f are as in Equation (6). The ATT is the observed outcome minus the imputed one, averaged over treated post-periods. Its identifying assumption is therefore the correctness of

the imputation model, not the unconditional parallel trends TWFE requires. FEct runs on the same full 2006–2022 panel as TWFE. Its factor-augmented extensions select no latent factors at any sample under the preferred specification (online supplemental data, Table S8). Consistent estimates from TWFE and FEct therefore show the null under two estimation strategies applied to the same sample. The result is additionally checked on a balanced 2007–2022 sub-panel against a doubly robust estimator and synthetic difference-in-differences, both under distinct assumptions (online supplemental data, Tables S9–S10) (Arkhangelsky et al., 2021; Sant’Anna & Zhao, 2020).

7. Quantitative results

The main outcome is the log of accumulated overlap with the partner’s prior corpus. It is reported alongside the rest of the outcome spectrum for two-way fixed effects and the counterfactual estimator across the three pooled samples (Tables 5 and 6). Under the preferred specification the BoW estimate at the full sample of newly paired prefectures is +0.001 (0.010) by two-way fixed effects and -0.016 (0.034) by the counterfactual estimator, neither significant. The 2016 designation did not move policy-text similarity. The same null holds at the other two categories under the BoW measure. In Category A, the absence of convergence after a fresh designation cannot be explained by the programme being too new or untested. In Category B, where the 2016 directive changed only the assigned partner, the estimates show no convergence toward the new partner. The two-way fixed-effects confidence interval rules out average shifts in textual similarity above roughly 2%.

The concentrated measures show small positive coefficients under the baseline specification (BoW “max” at +0.012, $p = 0.009$, full sample), but these follow the size of the partner’s archive rather than convergence. A partner with a larger archive offers more documents to match against, so its single best match is higher even when the typical document is no more similar. The positive coefficient falls monotonically as the measure widens and disappears in the archive-invariant pool mean. The partner-by-year fixed effect reduces it to a marginal positive on the single-best-match measure alone, and the counterfactual estimator removes even that (Tables 5 and 6). Overall, the BoW measure consistently produces a null estimate, and the dense embedding provides similar estimates apart from Category A. The pre-treatment evidence helps distinguish whether the Category A difference results from designation or reflects broader underlying differences.

Pre-trend tests can be underpowered, and conditioning on a passed test can introduce bias (Roth et al., 2023). Two diagnostics are therefore reported across the full outcome spectrum under each estimator. They are the close-leads pre-trend test and the HonestDiD smoothness breakdown $\hat{M}^*$ (Rambachan & Roth, 2023) (Tables 7 and 8). $\hat{M}^*$ records how far the smoothness restriction must be relaxed before the estimate no longer rules out a zero effect. A null estimate never rules out a zero effect to begin with, so a breakdown of zero is expected rather than evidence that the null is robust to larger violations. The pre-treatment leads are flat across the outcome spectrum at the full sample and Category B (Figures 1, 2 and 4).

The one apparent exception is Category A under Qwen3. The two-way fixed-effects estimate is significant for the broad aggregation outcomes. Under the counterfactual estimator the main outcome is insignificant, but the pooled mean and top-decile measures remain significant. The Qwen3 leads start well below the reference period

and rise smoothly through 2016, with no visible break at designation. This pattern is consistent with the continuation of a pre-existing trend (Figure 3). The rise has two pre-2016 sources. One is that Category A pools three provinces whose archives grow at different rates, which the pair fixed effect does not remove. The other is that the treated prefectures, including the autonomous prefectures of Xiangxi and Enshi, converge upward from a low base (the province-by-province estimates show the same pattern and are reported as descriptive; online supplemental data, Tables S11–S14). Neither source is a result of the 2016 designation.

The trend is visible under Qwen3 and not under BoW. The dense embedding compresses the series into a narrow range, so the small pre-existing trend is captured as a sharp significant one. By contrast, the Category A estimate under the BoW measure is null. Under the preferred specification the close-leads test and the smoothness breakdown flag only these Category A estimates (Tables 7 and 8).

The counterfactual estimator produces three small negative Qwen3 estimates at the Category B and pooled levels, for “psi_all”, “topa01”, and “topa10”. Each is no larger than 0.005 on the compressed Qwen3 scale (Table 6). These estimates appear only when three conditions hold simultaneously: the T6 specification, the counterfactual estimator, and the dense-embedding measure. If any one of these conditions changes, the estimate becomes null (Table 5). The placebo tests also detect pre-period movement in two of the three cases (Table 8).

Two further centrally assigned pairing programmes are studied as comparisons. The first is the 2013 restructuring of Guizhou’s partners,⁴ distinct from Bijie’s 2016 switch. Bijie is the province’s only prefecture in the main sample. The 2013 directive matched eight of Guizhou’s nine prefecture-level units, all but the provincial capital Guiyang, each with its own eastern partner city. The second is the 2017 Northeast counterpart-cooperation programme.⁵ The programme pairs each north-eastern province with an eastern one, alongside four designated prefecture pairs. Unlike the EWP, its main aim is market-oriented industrial and institutional cooperation rather than poverty alleviation.

Because these comparisons are restricted to paired prefectures, neither contains an untreated prefecture group. The comparison therefore comes from within the pool of assigned eastern partners. The study measures each focal prefecture’s overlap with the eastern partner it was assigned and compares that against its overlap with the programme’s other eastern partners. Each of those other partners is some other prefecture’s assignment, so together they serve as the placebo. If designation produced partner-specific diffusion, the assigned line would rise above the rest at the designation year. In Guizhou the assigned line and the placebo are nearly coincident across the observed period. In the Northeast the assigned line follows the placebo pool closely (Figure 5). Neither programme displays visible changes at the designation year. With no clean within-province untreated control, the case-study evidence is descriptive rather than causal.

Both trajectories rise from 2006 onward as the prefectures issue more documents. The gap between the assigned and placebo lines is in place well before either designation, and does not move post-treatment. The Northeast programme falls outside the

⁴General Office of the State Council, *Guanyu kaizhan duikou bangfu Guizhou gongzuo de zhidao yijian* [Guiding opinions on carrying out counterpart assistance work for Guizhou] (Guobanfa [2013] No. 11), 4 February 2013, https://www.gov.cn/zhengce/content/2013-02/07/content_6084.htm.

⁵General Office of the State Council, *Dongbei diqu yu dongbu diqu bufen shengshi duikou hezuo gongzuo fang'an* [Work plan for counterpart cooperation between Northeast China and selected eastern provinces and cities] (Guobanfa [2017] No. 22), 7 March 2017, https://www.gov.cn/gongbao/content/2017/content_5181083.htm.

poverty domain, while Guizhou, designated in 2013, is observed for almost a decade after designation, which is longer than both the Northeast programme and the main EWP sample.

Treated prefectures in the newly incorporated EWP provinces follow trajectories that are indistinguishable from those of never-treated prefectures, with no break in 2016. The trajectories' ordering predates the treatment year as early as 2006, in the main panel and in both comparison programmes (Figures 6, 7, 8 and 9). A partner's position relative to a prefecture's corpus is therefore shaped by the underlying characteristics of the two prefectures, not by the act of pairing itself. Arranged cooperation of this kind follows a political geography that precedes it (Jiang et al., 2023).

Table 5. ATT across the outcome spectrum (two-way fixed effects)

	BoW		Qwen3	
	T3	T6	T3	T6
<i>Full sample</i>				
log(sum)	+0.022 (0.017)	+0.001 (0.010)	+0.002 (0.002)	−0.001 (0.001)
psi_all	+0.001 (0.001)	+0.000 (0.001)	+0.001 (0.001)	−0.000 (0.001)
max	+0.012*** (0.004)	+0.008* (0.004)	+0.002 (0.002)	+0.001 (0.001)
topa01	+0.007** (0.004)	+0.003 (0.003)	+0.001 (0.002)	−0.000 (0.001)
topa10	+0.003 (0.002)	+0.001 (0.002)	+0.001 (0.001)	−0.000 (0.001)
<i>Cat A pooled</i>				
log(sum)	+0.012 (0.036)	+0.002 (0.037)	+0.010*** (0.003)	+0.009*** (0.003)
psi_all	+0.001 (0.002)	+0.000 (0.002)	+0.006*** (0.002)	+0.005*** (0.002)
max	+0.004 (0.014)	+0.003 (0.015)	+0.001 (0.003)	+0.000 (0.003)
topa01	+0.007 (0.012)	+0.006 (0.012)	+0.004 (0.003)	+0.004 (0.003)
topa10	+0.004 (0.006)	+0.003 (0.006)	+0.006*** (0.002)	+0.006** (0.002)
<i>Cat B+Pool</i>				
log(sum)	+0.020 (0.018)	−0.003 (0.008)	+0.000 (0.002)	−0.001 (0.001)
psi_all	+0.001 (0.001)	−0.000 (0.000)	+0.000 (0.001)	−0.001 (0.001)
max	+0.013*** (0.005)	+0.007* (0.004)	+0.002 (0.002)	+0.000 (0.001)
topa01	+0.008** (0.004)	+0.002 (0.003)	+0.001 (0.002)	−0.001 (0.001)
topa10	+0.003 (0.002)	+0.001 (0.002)	−0.000 (0.001)	−0.001 (0.001)

Notes. Single-coefficient ATT (cluster-robust SE in parentheses). Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6. ATT across the outcome spectrum (counterfactual estimator)

	BoW		Qwen3	
	T3	T6	T3	T6
<i>Full sample</i>				
log(sum)	+0.030 (0.020)	−0.016 (0.034)	+0.001 (0.002)	−0.014 (0.032)
psi_all	+0.002 (0.001)	−0.001 (0.001)	+0.001 (0.001)	−0.002 (0.001)
max	+0.013** (0.005)	+0.003 (0.006)	+0.002 (0.002)	−0.001 (0.002)
topa01	+0.008* (0.004)	+0.001 (0.004)	+0.001 (0.002)	−0.003 (0.002)
topa10	+0.004 (0.002)	+0.000 (0.002)	+0.000 (0.001)	−0.001 (0.001)
<i>Cat A pooled</i>				
log(sum)	+0.036 (0.053)	+0.031 (0.071)	+0.014** (0.007)	+0.021 (0.050)
psi_all	+0.002 (0.003)	+0.002 (0.003)	+0.008** (0.004)	+0.008** (0.004)
max	+0.007 (0.022)	+0.005 (0.023)	+0.003 (0.006)	+0.003 (0.006)
topa01	+0.010 (0.017)	+0.008 (0.017)	+0.006 (0.005)	+0.006 (0.004)
topa10	+0.007 (0.008)	+0.006 (0.009)	+0.008* (0.005)	+0.008** (0.004)
<i>Cat B+Pool</i>				
log(sum)	+0.025 (0.021)	−0.039 (0.037)	−0.001 (0.003)	−0.001 (0.043)
psi_all	+0.001 (0.001)	−0.002 (0.001)	−0.000 (0.001)	−0.003* (0.001)
max	+0.015*** (0.005)	+0.001 (0.006)	+0.002 (0.002)	−0.003 (0.002)
topa01	+0.008* (0.005)	−0.003 (0.004)	+0.000 (0.002)	−0.005** (0.002)
topa10	+0.003 (0.002)	−0.001 (0.003)	−0.001 (0.001)	−0.003** (0.001)

Notes. Counterfactual-imputation ATT (cluster bootstrap, 500 reps). Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Figure 1. Full-sample event studies, BoW

Notes. T3 baseline and T6 preferred overlaid; leads $\tau = -10, \dots, -2$ and lags $\tau = 0, \dots, +6$ with reference $\tau = -1$; 95% confidence intervals clustered by focal prefecture; dashed line at the 2016 designation.




Figure 2. Full-sample event studies, Qwen3

Notes. Specification and conventions as in Figure 1.




Figure 3. Category A event studies, Qwen3

Notes. Specification and conventions as in Figure 1, restricted to the Category A sample. The BoW counterpart is in the online supplemental data (Figure S3).




Figure 4. Category B and pooled event studies, Qwen3

Notes. Specification and conventions as in Figure 1, for the combined Category B and pooled sample. The BoW counterpart is in the online supplemental data (Figure S4).

Table 7. Pre-trend and sensitivity across the outcome spectrum (two-way fixed effects)

	Close-leads p				HonestDiD $\hat{M}^*$			
	BoW		Qwen3		BoW		Qwen3	
	T3	T6	T3	T6	T3	T6	T3	T6
<i>Full sample</i>								
log(sum)	0.598	0.696	0.519	0.152	0	0	0	0
psi_all	0.645	0.672	0.524	0.152	0	0	0	0
max	0.856	0.976	0.516	0.396	0	0	0	0
topa01	0.608	0.835	0.589	0.306	0	0	0	0
topa10	0.777	0.753	0.466	0.142	0	0	0	0
<i>Cat A pooled</i>								
log(sum)	0.064	0.008	0.242	0.199	0	0	.05	.05
psi_all	0.114	0.014	0.246	0.206	0	0	.05	.05
max	0.909	0.942	0.385	0.483	.05	.05	.05	.05
topa01	0.043	0.126	0.141	0.104	.05	.05	.05	.05
topa10	0.070	0.102	0.079	0.020	0	0	.05	0
<i>Cat B+Pool</i>								
log(sum)	0.878	0.922	0.995	0.617	0	0	0	0
psi_all	0.877	0.865	0.995	0.596	0	0	0	0
max	0.722	0.930	0.715	0.495	.05	0	0	0
topa01	0.419	0.766	0.643	0.504	0	0	0	0
topa10	0.910	0.790	0.752	0.498	0	0	0	0

Notes. Close-leads p : wild-cluster-bootstrap p of the joint test on the three immediate pre-treatment leads $\tau \in \{-2, -3, -4\}$ (Borusyak et al., 2024) (Webb weights, 9,999 reps; Roodman et al., 2019); bold marks $p < 0.10$. $\hat{M}^*$: HonestDiD smoothness breakdown, the first grid value ($M \in \{0, 0.05, \dots, 0.25\}$) at which the robust confidence set admits zero (0 = admits zero at the strictest grid value).

Table 8. Pre-trend and sensitivity across the outcome spectrum (counterfactual estimator)

	Placebo p				HonestDiD $\hat{M}^*$			
	BoW		Qwen3		BoW		Qwen3	
	T3	T6	T3	T6	T3	T6	T3	T6
<i>Full sample</i>								
log(sum)	0.475	0.735	0.390	0.479	0	0	0	0
psi_all	0.695	0.576	0.368	0.318	0	0	0	0
max	0.817	0.624	0.954	0.720	0	0	0	0
topa01	0.781	0.410	0.474	0.386	0	0	0	0
topa10	0.907	0.743	0.245	0.238	0	0	0	0
<i>Cat A pooled</i>								
log(sum)	0.273	0.334	0.323	0.699	.05	0	.05	0
psi_all	0.341	0.238	0.334	0.288	.05	0	.05	.05
max	0.925	0.985	0.398	0.474	0	0	0	0
topa01	0.752	0.792	0.220	0.250	0	0	.05	0
topa10	0.513	0.486	0.320	0.233	.05	0	.05	.05
<i>Cat B+Pool</i>								
log(sum)	0.931	0.326	0.104	0.807	0	0	0	0
psi_all	0.798	0.260	0.109	0.049	0	0	0	0
max	0.688	0.572	0.750	0.921	0	0	0	0
topa01	0.748	0.328	0.258	0.204	0	0	0	0
topa10	0.624	0.423	0.057	0.071	0	0	0	0

Notes. *Placebo p* : FEct placebo-test p (test that the average placebo-period effect is zero), treatment shifted three periods earlier (Liu, Wang, & Xu, 2024); bold marks $p < 0.10$. $\hat{M}^*$: the smoothness breakdown as in Table 7, computed under the HonestDiDFEct adaptation, which applies the sensitivity analysis to the estimator's placebo-period dynamic effects.




Figure 5. Assigned and placebo trajectories, Guizhou 2013 and Northeast 2017

Notes. Calendar-year mean log overlap (BoW) of each programme's focal prefectures with their assigned partners, against the within-pool placebo mean, with the Guizhou prefectures on the left and the 35 Northeast prefectures on the right. The Guizhou comparison uses the province's seven paired prefectures. Bijie is excluded because it enters the main sample, and Guiyang because the 2013 directive did not pair it. Yanbian is excluded because it is also paired under the main East–West programme. The Guizhou panel draws two placebos, the primary excluding the three switched prefectures' former partners and a robustness variant including them. Red dashed lines mark each programme's designation year, here and in Figures 6–9; the grey line at 2020 marks the COVID onset.




Figure 6. Treated prefectures and their placebo clouds, newly incorporated provinces

Notes. Log accumulated overlap (BoW) of every prefecture in each newly incorporated province with the province's assigned eastern partner, one panel per province and partner, 2006–2022. The prefectures designated to the panel's partner are navy and the province's never-treated prefectures grey; common scale across panels.




Figure 7. Guangxi against its five 2016 partners

Notes. Log accumulated overlap (BoW), 2006–2022, one panel per eastern partner. The prefectures assigned to the panel's partner are navy (Hechi and Baise for Shenzhen; Guilin and Hezhou for Zhaoqing; Nanning and Laibin for Maoming; Chongzuo for Jiangmen; Liuzhou for Zhanjiang), treated prefectures scored against partners not assigned to them green, and never-treated prefectures grey. The five pair panels share a common scale. The dashed-border panel repeats the Shenzhen data at its native scale.




Figure 8. Partner-change prefectures, Guizhou 2013 restructuring

Notes. Log accumulated overlap (BoW) of each Guizhou prefecture whose 2013 restructuring changed its partner, shown against the newly assigned partner (navy), the old partner (green), and the five within-pool placebo partners (grey).




Figure 9. Designated prefecture pairs, 2017 Northeast programme

Notes. Log accumulated overlap (BoW) of each north-eastern prefecture with its 2017-designated partner (navy; Changchun carries a second assigned partner, Hangzhou, in red) against its within-block placebo partners (grey, one line per partner). Grey dash-dot line at the 2020 COVID-19 onset; common scale across panels.

8. Qualitative evidence: the Shenzhen–Guangxi case

The interviews discussed earlier cover Shenzhen–Guangxi pairings from two assignment rounds. The first began in 2016 and involved counties in Hechi and Baise. The second began in 2020 and involved counties in Nanning. Shenzhen is among the programme’s best-resourced donors. Its own development also involved substantial policy borrowing, as the city built market institutions partly by learning from Hong Kong within an experiment authorised by the centre (J. Zhang, 2012). These features make Shenzhen a useful case for examining whether a centrally assigned partnership can produce policy diffusion. Two of the pairings described in the interviews, Shenzhen–Hechi and Shenzhen–Baise, are also treated pairs in the quantitative analysis. The case study therefore observes the same units the estimates cover.

The two provincial governments set the framework through a joint opinion. A provincial work list then assigns tasks, quotas, and deadlines to each prefecture, and cadres on both sides draft within this administrative chain.⁶ In Nanning, the prefecture-level implementation plan for the 2021–2025 cooperation programme was drafted in the city’s Guangdong–Guangxi cooperation office, where Shenzhen’s seconded cadre serves as deputy director (Y. Chen & Ling, 2023). The channel therefore reaches the level captured by the index, but the resulting text is additive and partnership-specific. The seconded cadre acts as a broker of the partnership, not as the author of the county’s or prefecture’s general regulatory texts. This differs from the mechanism in Lu (2024)’s study, where a rotated provincial leader’s annual report is itself the text being measured. In that setting, similarity rises while the leader serves and may fade after the leader moves on. Where operational content enters policy text, the template is set by the centre. “Town-level policy innovation and flexibility are very limited,” one interviewee said. Decisions must align with higher-level criteria and funding directives (Interview 13).

The national evaluation rules assess the organisational, fiscal, and developmental performance of each partnership. They include leadership arrangements, cadre dispatch, financial contributions, enterprise investment, employment outcomes, and the coverage of county pairings.⁷ The pairs’ performance is assessed by a third party. Adopting the partner’s policies is not included among the indicators. Instead, the cooperation text produced by cadres is itself an audited and rewarded output. The quantitative analysis provides no evidence that this text spreads beyond the cooperation domain into other policy fields.

The dissimilarity between the partners is the first thing officials describe. A coordination official put the difficulty plainly: the counterparts in the assisted counties “have a very different mindset,” and “they want results without demonstrating the necessary patience” (Interview 1). Implementation is shared between county leaders and seconded cadres through a dual standing committee that is formally equal, and in practice the arrangement can slow decision-making (Interviews 2–3).

⁶People’s Government of Guangdong Province and People’s Government of Guangxi Zhuang Autonomous Region, *Guanyu jingyibu jiaqiang Yue-Gui fupin xiezuo gongzuo de yijian* [Opinions on further strengthening Guangdong–Guangxi poverty alleviation cooperation work] (Yuefuhuan [2017] No. 284), 15 October 2017, https://www.gd.gov.cn/gkmlpt/content/0/146/mpost_146414.html; General Office of the People’s Government of Guangdong Province and General Office of the People’s Government of Guangxi Zhuang Autonomous Region, *Yue-Gui fupin xiezuo he quyu hezuo gongzuo qingdan (shengji cengmian) (xiuding)* [Work list for Guangdong–Guangxi poverty alleviation cooperation and regional cooperation (provincial level, revised)] (Yuebanhan [2019] No. 198), 28 May 2019, https://www.gd.gov.cn/zwgk/gongbao/2019/22/content/post_3366424.html.

⁷State Council Leading Group for Poverty Alleviation and Development, *Dongxibu fupin xiezuo kaohe banfa (shixing)* [Evaluation measures for East–West poverty alleviation cooperation (trial)] (Guokaifa [2017] No. 6), 8 August 2017.

The local conditions the policies are expected to address also differ. “The local demand is more for basic infrastructure construction, but the central emphasis is more on industrial poverty alleviation” (Interview 13). Counties often prioritise roads, water, and housing, reflecting needs at a stage of development that the partner’s own policy agenda no longer addresses. Entrenched local interests also limit diffusion. One local official drew the boundary clearly in conversation with the Shenzhen stationed official: “This is my home. I can’t just dismiss anyone at my will. You can pat yourself on the back and leave, but I have to stay here.” Local officials carry career and political stakes in the county long after seconded cadres depart.

In education, visiting teachers were received warmly and given standing, often a vice-principal title. However, local teachers were often reluctant to adopt their methods, and the younger teachers from Shenzhen found their advice met with resistance or deference. A senior educator who led District F’s teaching delegation in County H proposed a structured scheme for cultivating young local teachers and received little support. She described her work at the county school as having “a sense of hanging in midair”, and later asked to be transferred to a village school (Interview 9).

The second barrier to diffusion comes from the evaluation system. Officials concentrate on what the evaluation rewards. Adopting a dissimilar partner’s distinctive policy is risky rather than rewarded. A Shenzhen stationed cadre explained, “Three years of investment is tied to a lifetime of responsibility. If the project collapses, it will follow the person who approved it.” “In the end,” the same cadre added, “what matters is that the project doesn’t fail. That is the priority” (Interview 2). Another cooperation office favoured co-investment in projects that county or central funding has already vetted, precisely to avoid errors that might trigger accountability investigations (Interview 3). These accounts are more directly linked with project selection and implementation that serve the goal of poverty reduction. However, they can also provide indirect evidence of the calculations a local official undertakes when deciding whether to adopt policies from Shenzhen. A partner’s policy is untested in the recipient prefecture’s local conditions. Its adoption, or even the declaration of adoption, is not rewarded by any evaluation indicator and is not required as a programme goal. By contrast, a failed adoption could expose officials to accountability risks.

The people who conduct the evaluation strengthen this caution. Each partnership is assessed by an inspection team drawn from a central agency and from officials of another eastern province in the programme. In 2023 the team evaluating District F was led by officials from the United Front Work Department of the CPC Central Committee and from Tianjin (Interview 1). Evaluators from another partnership have incentives to scrutinise closely and identify faults, since this can make their own partnership appear stronger by comparison. This pushes assessed officials further toward what the indicators measure and away from untested initiatives.

Rotation and the temporary nature of the stationing also work against diffusion. “Even if there is a good initiative, when the teams change, the connection is lost,” a Shenzhen stationed cadre said (Interview 3). In County H the incoming team dismissed its predecessor’s projects (Interview 5), and newly assigned teams seek to distinguish themselves. Incoming leaders tend to discard inherited policies for new ones of their own (J. Wang & Yu, 2026). Rotation removes the position, not the acquaintance. Donor-side officials may still contact officials in counties whose pairing ended years earlier. What leaves with the posting is the standing-committee seat, the budget line, and the mandate through which practice could be codified into policy text. A policy built for the partner’s conditions is riskier to carry out on the assisted side. The inspection system further increases the costs of that risk, and no indicator rewards

the adoption of the partner's policy. Declining to borrow therefore carries little cost, while borrowing can stake an official's career on a policy designed for somewhere else.

The channels seem to be more effective where what they carry is not policy. Most of the people the channels move are teachers, doctors, and technical staff. They carry the practice of their own work. In health, one seconded doctor introduced a standard ultrasound reporting template across the county hospital in County M. Another built the county's stroke centre around a workflow that delivers thrombolytic drugs within 45 minutes of arrival (Interviews 6–7). In County L, an orthopaedic specialist performed the hospital's first operation on a patient over 80 (Interview 8). In education, a visiting teacher built the first physical-education curriculum at his assigned school in County H and trained the county's PE teachers (Interview 10). In business, eastern firms brought new operating concepts into recipient counties. As the manager of a Shenzhen firm operating in County L observed, their presence also “pressured local companies to learn advanced practices” through competition (Interview 12).

Clinical workflow, teaching method, and business practice are not what a binding prefecture-level norm records. These practices are also unlikely to reach the policy text within the study's window. They are personal, reaching the staff, students, and colleagues who worked beside the person who carried them (Interviews 8, 10–11, 13). The programme does not actively choose or sustain the ties once a posting ends. What they carry therefore depends on the person who holds them (Interviews 6–9).

9. Conclusion and policy implications

This paper uses the East–West Pairing programme as a natural experiment to test whether an assigned tie between dissimilar prefectures moves one partner's binding policy text towards the other's. The estimates show no detectable average convergence after designation. Yet the programme does connect the partners through seconded cadres, mandatory professional exchanges, and routine contact between their offices. Working practices and personal ties pass through these channels without appearing in the recipient's binding policy. Partner dissimilarity makes direct borrowing less likely, while the accountability system gives local officials little incentive to adopt the partner's policies.

The null is limited to codified policy in the six years after the designation. Binding norms may follow a slow progression from projects to working documents and then to formal policy. In the longer Guizhou comparison the assigned and placebo trajectories stay nearly coincident for almost a decade after designation, although that evidence is descriptive rather than causal. The study also does not test whether a direct measure of the transferred practices would reveal convergence outside binding policy text. However, the current identification strategy does not show whether partner dissimilarity or the accountability regime is more important in producing the null. The fieldwork provides evidence on how both barriers operate, but it does not identify their separate effects. Furthermore, the fieldwork comes from one of the programme's best-resourced donors, where the conditions for transfer are relatively favourable. The case study is therefore used to illustrate how the proposed mechanisms operate rather than to identify their effects causally.

Previous studies have described the channels created by pairing schemes (Turok & Habiyaremye, 2020; Q. Wei & Yang, 2025), while the available causal evidence examines investment flows and fiscal outcomes rather than policy change (X. Li et al., 2025). This article extends that evaluation to policy by combining quantitative

and qualitative evidence. The fieldwork also reveals forms of transfer that the policy measure does not capture, and these may still contribute to regional development.

A prosperous eastern prefecture may have administrative practices that could benefit a poorer western partner. The EWP, however, rewards poverty reduction and regional development rather than adoption of the donor's policies. Policy borrowing in China is more likely when the centre recognises a local model and gives officials an incentive to adopt it (J. Wang & Yu, 2026). A government seeking policy convergence through assigned partnerships would therefore need to reward that outcome directly. However, large differences between partners would still make direct borrowing difficult.

What lasts after the pairing is the working relationship, and it can outlive the formal tie. To draw lasting value from a pairing, a government should preserve the relationships it has already built rather than start fresh at each round of postings (Xing & Yang, 2025). Formal targets are unlikely to preserve personal relationships. Over-stringent monitoring and sanctions can crowd out the intrinsic motivation officials bring to a task (Tu & Gong, 2022), so a relationship converted into a reporting requirement risks eroding the very motivation it depends on. Future research could test whether partner dissimilarity and accountability pressures also limit policy diffusion in other administratively designed partnerships. It could also examine working documents below the binding-norm tier or settings in which officials are rewarded for citing and adopting a designated model.

References

- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12), 4088–4118. <https://doi.org/10.1257/aer.20190159>
- Arts, S., Hou, J., & Gomez, J. C. (2021). Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures. *Research Policy*, 50(2), 104144. <https://doi.org/10.1016/j.respol.2020.104144>
- Baier, S. L., & Bergstrand, J. H. (2007). Do free trade agreements actually increase members' international trade? *Journal of International Economics*, 71(1), 72–95. <https://doi.org/10.1016/j.jinteco.2006.02.005>
- Berry, F. S. (2023). Finding common ground: Innovation and diffusion across political science and public management research. *PS: Political Science & Politics*, 56(3), 329–338. <https://doi.org/10.1017/S1049096523000112>
- Berry, F. S., & Berry, W. D. (2018). Innovation and diffusion models in policy research. In C. M. Weible & P. A. Sabatier (Eds.), *Theories of the policy process* (4th ed., pp. 253–297). Routledge.
- Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6), 3253–3285. <https://doi.org/10.1093/restud/rdae007>
- Chen, J., Wen, J., Deng, Y., Chen, M., & Ma, F. (2026). Measuring policy diffusion intensity: A text-driven analysis of government documents. *Information Processing & Management*, 63(2), 104383. <https://doi.org/10.1016/j.ipm.2025.104383>
- Chen, W., Liao, L., & Yi, H. (2025). Multilevel policy textual learning in Chinese local environmental policies. *Journal of Public Policy*, 45(3), 500–521. <https://doi.org/10.1017/S0143814X25100652>
- Chen, Y., & Ling, P. (2023, October 20). *Yue-Gui xiezuo bangfu ganbu Chen Hao: Ba guazhi dangcheng renzhi, ba taxiang dangcheng guxiang [Guangdong–Guangxi cooperation cadre Chen Hao: Treating the secondment as a substantive post, the other's home as one's own]*. Nanning xinwen wang 南宁新闻网. <http://www.nnnews.net/p/3178252.html>
- Ethayarajh, K. (2019). How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 55–65.
- Glick, D. M. (2014). Learning by mimicking and modifying: A model of policy knowledge diffusion with evidence from legal implementation. *The Journal of Law, Economics, and Organization*, 30(2), 339–370. <https://doi.org/10.1093/jleo/ews041>
- Goodburn, C., & Knoerich, J. (2022). Importing export zones: Processes and impacts of replicating a Chinese model of urbanization in rural south India. *Urban Geography*, 43(10), 1496–1518. <https://doi.org/10.1080/02723638.2021.2014669>
- Heffer, A. S., & Schubert, G. (2023). Policy experimentation under pressure in contemporary China. *The China Quarterly*, 253, 35–56. <https://doi.org/10.1017/S0305741022001801>
- Heilmann, S. (2008). From local experiments to national policy: The origins of China's distinctive policy process. *The China Journal*, (59), 1–30.

- Jiang, X., Eaton, S., & Kostka, G. (2023). No county is an island: The rise of inter-jurisdictional cooperation among counties in China. *Regional Studies*, 57(10), 2021–2036. <https://doi.org/10.1080/00343404.2022.2154333>
- Li, M., & Wang, J. (2025). The impact of vertical and horizontal pressures on policy diffusion: Evidence from APLS adoption in china. *Journal of Chinese Governance*, 10(4), 558–582. <https://doi.org/10.1080/23812346.2024.2426769>
- Li, X., Wang, Y., Zhang, C., & Yang, H. (2025). Can territorial cooperation improve regional development? evidence from Zhejiang’s Mountain–Sea cooperation. *Regional Studies*, 59(1), 2549591. <https://doi.org/10.1080/00343404.2025.2549591>
- Lim, K. F., & Horesh, N. (2016). The “Singapore fever” in China: Policy mobility and mutation. *The China Quarterly*, 228, 992–1017. <https://doi.org/10.1017/S0305741016001120>
- Linder, F., Desmarais, B., Burgess, M., & Giraudy, E. (2020). Text as policy: Measuring policy similarity through bill text reuse. *Policy Studies Journal*, 48(2), 546–574. <https://doi.org/10.1111/psj.12257>
- Liu, L., Wang, Y., & Xu, Y. (2024). A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data. *American Journal of Political Science*, 68(1), 160–176. <https://doi.org/10.1111/ajps.12723>
- Liu, L., Xin, G., & Yi, H. (2024). Information asymmetry and vertical collective action dilemma: The case of targeted poverty alleviation in china. *Review of Policy Research*, 41(5), 790–814. <https://doi.org/10.1111/ropr.12575>
- Lu, J. (2024). Leadership rotations and the convergence of subnational economic policies in China: Evidence from provincial government work reports. *The Journal of Law, Economics, and Organization*, 40(2), 528–569. <https://doi.org/10.1093/jleo/ewac026>
- Luo, K., Wang, S., & Yang, D. Y. (2025). *Laboratories of autocracy: Landscape of central–local dynamics in China’s policy universe* (NBER Working Paper No. 34219). National Bureau of Economic Research. <https://doi.org/10.3386/w34219>
- McCann, E., & Ward, K. (2013). A multi-disciplinary approach to policy transfer research: Geographies, assemblages, mobilities and mutations. *Policy Studies*, 34(1), 2–18. <https://doi.org/10.1080/01442872.2012.748563>
- O’Brien, K. J., & Li, L. (1999). Selective policy implementation in rural China. *Comparative Politics*, 31(2), 167–186. <https://doi.org/10.2307/422143>
- Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5), 2555–2591. <https://doi.org/10.1093/restud/rdad018>
- Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., & Webb, M. D. (2019). Fast and wild: Bootstrap inference in Stata using boottest. *The Stata Journal*, 19(1), 4–60. <https://doi.org/10.1177/1536867X19830877>
- Roth, J., Sant’Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2), 2218–2244. <https://doi.org/10.1016/j.jeconom.2023.03.008>
- Sant’Anna, P. H. C., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1), 101–122. <https://doi.org/10.1016/j.jeconom.2020.06.003>

- Shipan, C. R., & Volden, C. (2008). The mechanisms of policy diffusion. *American Journal of Political Science*, 52(4), 840–857. <https://doi.org/10.1111/j.1540-5907.2008.00346.x>
- Tu, W., & Gong, T. (2022). Accountability intensity and bureaucrats' response to conflicting expectations: A survey experiment in China. *Public Management Review*, 24(11), 1779–1801. <https://doi.org/10.1080/14719037.2021.1930123>
- Turok, I., & Habiyaremye, A. (2020). Territorial collaboration: A novel way to spread prosperity. *Regional Studies*, 54(12), 1776–1786. <https://doi.org/10.1080/00343404.2020.1807490>
- Volden, C. (2016). Failures: Diffusion, learning, and policy abandonment. *State Politics & Policy Quarterly*, 16(1), 44–77. <https://doi.org/10.1177/1532440015589910>
- Wang, J., & Yu, J. (2026). Policy innovation by China's local governments: Beyond the “championship” model. *The China Quarterly*, 265, 127–145. <https://doi.org/10.1017/S0305741025101732>
- Wang, S., & Yang, D. Y. (2025). Policy experimentation in China: The political economy of policy learning. *Journal of Political Economy*, 133(7), 2180–2228. <https://doi.org/10.1086/734873>
- Wang, Z. (2025). Designated accountability: Pair-making mechanisms and evolving central–local relations in china's targeted poverty alleviation. *Journal of Contemporary China*, 34(152), 232–252. <https://doi.org/10.1080/10670564.2024.2333484>
- Wei, J. (2021, August 3). *Zizhiqu Yue-Gui xiezuo ban, gongzuo dui guapai [The autonomous region's Guangdong–Guangxi cooperation office and work team unveiled]*. Guangxi ribao 广西日报. <http://gx.people.com.cn/n2/2021/0803/c179430-34849604.html>
- Wei, Q., & Yang, W. (2025). Addressing uneven development through state-steered intercity cooperation? insights from Shenzhen–Ganzhou cooperation plan-making. *Regional Studies*, 59(1), 2377669. <https://doi.org/10.1080/00343404.2024.2377669>
- Xing, P., & Yang, R. (2025). The gift of mobility: Transfer networks of local decision-makers and interlocal collaboration. *Review of Policy Research*, 42(3), 682–710. <https://doi.org/10.1111/ropr.12642>
- Xiong, M., Cheng, S., & Zhao, J. Z. (2025). Policy imitation based on similarities: The diffusion of public–private partnerships. *Public Administration and Development*, 45(2), 130–145. <https://doi.org/10.1002/pad.2084>
- Yi, H., Berry, F. S., & Chen, W. (2018). Management innovation and policy diffusion through leadership transfer networks: An agent network diffusion model. *Journal of Public Administration Research and Theory*, 28(4), 457–474. <https://doi.org/10.1093/jopart/muy031>
- Yi, H., Liu, W., & Ma, L. (2024). Designed networks and the emergence of self-organizing interlocal learning network: Evidence from Chinese cities. *Public Administration*, 102(1), 21–39. <https://doi.org/10.1111/padm.12906>
- Yuan, W. (2021, July 29). *Guangxi Yue-Gui xiezuo bangongshi, Yue-Gui xiezuo gongzuo dui zhengshi guapai [Guangxi's Guangdong–Guangxi cooperation office and work team formally unveiled]*. Shenzhen shangbao 深圳商报. <https://baijiahao.baidu.com/s?id=1706619140189281535&wfr=spider&for=pc>
- Zhang, J. (2012). From Hong Kong's capitalist fundamentals to Singapore's authoritarian governance: The policy mobility of neo-liberalising Shenzhen, China. *Urban Studies*, 49(13), 2853–2871. <https://doi.org/10.1177/0042098012452455>

- Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., & Artzi, Y. (2020). BERTScore: Evaluating text generation with BERT. *International Conference on Learning Representations (ICLR)*.
- Zhu, X., & Zhang, Y. (2016). Political mobility and dynamic diffusion of innovation: The spread of municipal pro-business administrative reform in China. *Journal of Public Administration Research and Theory*, 26(3), 535–551. <https://doi.org/10.1093/jopart/muv025>