Laboratories of Autocracy: Landscape of Central–Local Dynamics in China's Policy Universe

Kaicheng Luo
Shaoda Wang
David Y. Yang*

August 31, 2025

Abstract

Using a comprehensive collection of 3.7 million Chinese policy documents and government work reports spanning the past two decades, we identify 115,679 distinct policies and systematically trace their initiation and diffusion. Our analysis reveals three key findings. First, China's policymaking has historically been highly decentralized, with local bureaucrats playing crucial roles in both creating new policies and spreading them. Second, since 2013, policymaking has become substantially more centralized, driven primarily by changing bureaucratic incentives: bottom-up innovation is no longer rewarded, while strict compliance with central directives is. Third, our analysis of industrial policies shows that centralization affects both policy suitability and effectiveness. Top-down industrial policies tend to align poorly with local conditions and are less effective at fostering industrial growth, underscoring the costs of centralization. At the same time, centralization provides offsetting benefits by mitigating distortions in decentralized policy diffusion that stem from strategic competition among local officials. Overall, our quantitative assessment indicates that the economic costs of centralization in China have substantially outweighed its benefits.

Keywords: Centralization, policy innovation and diffusion, China, industrial policy
JEL Classification: P21; H77; O25

*. Luo: Massachusetts Institute of Technology; kluo0630@mit.edu. Wang: University of Chicago, NBER, and BREAD; shaoda@uchicago.edu. Yang: Harvard University, NBER, BREAD, and J-PAL; davidyang@fas.harvard.edu. We thank Loren Brandt, Ben Faber, Hanming Fang, Fred Finan, Chang-tai Hsieh, Ruixue Jia, Hongbin Li, Ernest Liu, Mushfiq Mobarak, Jacopo Ponticelli, Scott Rozelle, Michael Song, Kjetil Storesletten, Wei Xiong, Noam Yuchtman, Xiaodong Zhu, and many seminar and conference participants for stimulating comments and suggestions. Yangsai Chen, Ruiyang He, Ruoyuan Li, Anran Tan, Meilan Zhang, and Xinyi Zhang provided outstanding research assistance.

1 Introduction

A fundamental question in political economy concerns the appropriate level for making policy decisions — a matter that has sparked extensive debates over the merits of centralization versus decentralization (Hayek 1945; Rueschemeyer, Skocpol, and Evans 1985; Bardhan 2002; Besley and Coate 2003; Mookherjee 2015). While top-down policy promotion may streamline adoption, internalize spillovers, and enhance efficiency, it often sacrifices the local suitability that bottom-up policy initiatives provide (Tiebout 1956; Oates 1972; Alesina and Spolaore 1997; Bolton and Roland 1997; Alesina, Baqir, and Hoxby 2004). Such tension is especially relevant in governing large polities with high levels of regional heterogeneity.

Despite its theoretical importance, studying the centralization of policymaking remains empirically challenging. Measuring centralization in policymaking is difficult because it requires systematically tracing the origin and diffusion patterns of *all* policies across layers of government hierarchy; as a result, most studies of (de)centralization focus their attention on the enforcement of a particular policy or the provision of a specific public good (Olken 2007; Burgess et al. 2012; Lipscomb and Mobarak 2017; Jia and Nie 2017; Dal Bó et al. 2021; Balán et al. 2022). Assessing centralization's impact on policy outcomes is even more demanding because it involves linking policies to both local conditions and intended outcomes.

In this paper, we study the centralization of policymaking and how it affects the local suitability of policies, focusing on China's policymaking across all domains over the past two decades. China is an especially compelling context, where the centralization–decentralization tradeoff is critical and the empirical challenges we highlight can be directly addressed. We investigate two questions. First, what share of local governments' policy portfolios is shaped by the central government's direct involvement — such as initiating or explicitly promoting policies in a top-down manner? Second, does the central government's direct involvement undermine policy suitability and effectiveness at the local level, a first-order concern with centralization emphasized in the theoretical literature?

We compile a comprehensive dataset of 422 thousand central government policy documents and 3.3 million local government policy documents and work reports. From this corpus, we identify 115,679 distinct policies implemented from 2004 to 2020 and trace their origins and vertical and horizontal diffusion patterns. For the subset of industrial policies aimed at promoting industrial growth, we also measure policy suitability and effectiveness.

We document four main findings that address our research questions. First, policymaking in China is highly decentralized. Over the past two decades, 82% of the policies appearing in local governments' portfolios originated as local initiatives; of these, 74% diffused solely horizontally among local governments and never involved explicit central government action. Further analysis exploiting political turnovers reveals that local bureaucrats are the primary drivers of decentralized policy initiation and diffusion, and bureaucrats are rewarded with political promotion for local policy innovation.

Second, since 2013, decentralized policymaking has declined dramatically: the share of top-down policies in local governments' portfolios has increased by 40%, the adoption rate of top-down initiatives has nearly tripled, and local replication of central policy details has more than doubled. This shift likely resulted from general changes in local bureaucrats' career incentives: after 2013, political promotion was granted to bureaucrats who most actively implemented top-down policies rather than those who pioneered new policies as observed before. Furthermore, the timing of this centralization aligns with the central government's phased roll-out of tighter top-down control across different ministries and policy domains.

Third, locally initiated and horizontally diffused policies tend to be associated with higher *ex ante* local suitability and better *ex post* economic outcomes. Focusing on industrial policies aimed at promoting sector-specific industrial growth and innovation, we find that those initiated or adopted by local governments without central involvement are better aligned with local conditions, as measured by *pre-existing* regional supply chains and private firms' *ex ante* investment preferences. In contrast, top-down industrial policies initiated or explicitly adopted by the central government show systematically weaker alignment with local conditions. Moreover, industrial policies that are better matched to local conditions prove significantly more effective, on average, in achieving policy objectives — including increased industrial output, patenting, and exports — underscoring the costs of centralization.

Fourth, centralizing policymaking curbs strategic competition among local bureaucrats that otherwise impedes learning from peer jurisdictions. When policies are initiated and diffused locally without central government's explicit involvement, local bureaucrats competing for the same political promotion opportunities may be reluctant to adopt policies from one another for fear of boosting their competitors' credentials. Since local bureaucrats with similar promotion prospects are often posted to economically comparable localities, we find that such rivalry stifles the diffusion of local policy innovations, undermines local suitability among diffused policies, and dampens economic performance. Centralization alleviates these distortions, since local bureaucrats no longer worry about

political competitors receiving credit when they implement top-down policies. On net, our back-of-the-envelope calculation suggests that the gains from centralization in promoting policy diffusion are outweighed by the losses from reduced policy suitability due to central intervention.

Since China's 1979 departure from a centrally planned economy, its extraordinary growth has been widely credited to locally driven initiatives under a decentralized framework (Oi 1995; Montinola, Qian, and Weingast 1995; Xu 2011; Chen, Li, and Zhu 2024). This paper builds on that view by demonstrating that, under well-structured political incentives, an autocracy can also function as a vibrant laboratory for policy innovation — generating new, locally suited ideas — much like Justice Brandeis's famous characterization of federalist systems as "laboratories of democracy."[1] Over the past decade, however, China's decentralization trend reversed as political incentives for policy innovation were removed. This shift demands rigorous scrutiny, given China's history of heavy-handed central planning failures (Lin 1990; Meng, Qian, and Yared 2015; Frank et al. 2024), its vast regional heterogeneity, and the mounting complexity of governance.

In this paper, we focus on policies' local suitability — an important dimension in theoretical debates on centralization versus decentralization, and one that we show to be strongly linked to policy effectiveness — as a criterion for evaluating the implications of China's shift from a decentralized to a more centralized policymaking regime. We view this as a proof-of-concept exercise to show how policy outcomes depend on both *where* policies are initiated and *how* their initiation and diffusion are incentivized. While a comprehensive evaluation of every dimension of centralization lies beyond the scope of this paper, we provide suggestive evidence on several others that may be relevant, such as economies of scale, regional spillover internalization, national security imperatives, and time horizons. We hope future studies tackle these additional rationales for centralization.

This paper relates to three strands of literature. First, it contributes to the emerging literature on policy diffusion and innovation. A large body of work demonstrates how policy innovation and diffusion in federalist societies can serve as "laboratories of democracy" (Besley and Case 2003; Bernecker, Boyer, and Gathmann 2021; Caughey, Xu, and Warshaw 2017; Grumbach 2023; DellaVigna and Kim 2022). Yet, comparatively little is known about how authoritarian regimes acquire the decentralized information necessary to design and implement effective policies, given the typically limited scope for bottom-up participation in policymaking in such settings.[2] A particularly relevant study

1. See New State Ice Co. v. Liebmann (1932).

2. A notable exception is the literature on policy experimentation in China, which highlights the role of top-down pilots in policy learning (Montinola, Qian, and Weingast 1995; Cao, Qian, and Weingast 1999; Heilmann 2008; Wang and Yang 2025).

is DellaVigna and Kim (2022), which documents patterns and dynamics of policy innovation and diffusion in the United States. Our findings in the Chinese context reveal several notable contrasts with the U.S. model: *(i)* although both countries rely heavily on decentralized policymaking, in China local bureaucrats — rather than localities *per se* — are the primary drivers; *(ii)* in both settings, political rivalries hinder policy diffusion, but through different mechanisms — partisan politics in the U.S. versus bureaucratic competition in China; and *(iii)* while political polarization undermines the efficiency of policy diffusion in the U.S., in China political centralization emerges as the most significant impediment.

Second, this paper contributes to the long-standing debate regarding centralization versus decentralization (Hayek 1945; Rueschemeyer, Skocpol, and Evans 1985; Bardhan 2002; Mookherjee 2015). On the one hand, our finding that bottom-up policies exhibit higher suitability with local conditions lends empirical support to the theoretical literature emphasizing the importance of decentralized information (Tiebout 1956; Oates 1972; Cremer, Estache, and Seabright 1994; Seabright 1996; Alesina and Spolaore 1997; Bolton and Roland 1997; Besley and Coate 2003; Alesina, Baqir, and Hoxby 2004). On the other hand, our results on strategic biases in policy diffusion illustrate the distortions that can arise from decentralized regional competition (Blanchard and Shleifer 2001; Sonin 2003; Cai and Treisman 2004; Young 2000). Unlike much of the existing empirical literature — which typically examines the impact of (de)centralization on a specific policy domain, such as safety (Jia and Nie 2017), pollution (Lipscomb and Mobarak 2017; Wang and Wang 2020), deforestation (Burgess et al. 2012), corruption (Olken 2007), agriculture (Dal Bó et al. 2021), or taxation (Balán et al. 2022) — our paper is the first to offer a holistic account of the complete policy portfolios across various levels of government, thereby shedding light on not only policy enforcement but also on *policymaking* itself.

Third, our analyses on the perils of centralized policymaking echo the literature that highlight the importance of decentralized policy learning in driving China's economic boom since reform and opening-up (Heilmann 2008; Rawski 1995; Roland 2000; Qian 2002; Wang and Yang 2025). More recently, Fang, Li, and Lu (2025), focusing on China's industrial policies, document a growing centralization trend that is consistent with the broad findings of ours. This paper also relates to the cautionary tales of over-centralization, such as Kornai (1960) and Nove (1971). Although information asymmetry between central and local governments has diminished substantially in recent years (Martinez-Bravo et al. 2022), we find that bottom-up policies remain markedly more suitable for local conditions and yield superior outcomes, supporting the findings of Chen, Li, and Zhu (2024) in their examination of policymaking during the 1980s and 1990s.

The remainder of the paper is organized as follows. Section 2 outlines the institutional background, data sources, and construction process. Section 3 provides an overview of China's decentralized policymaking environment. Section 4 documents the post-2013 shift toward centralization and examines its underlying causes. Section 5 introduces our measures of policy–locality suitability and analyzes their relationship with policy effectiveness. Section 6 evaluates the benefits and costs of centralization through the lens of policy–locality suitability. Section 7 concludes.

2 Institutional background and data

In this section, we describe the institutional background of governance and policymaking in China. We then introduce the policy documents and data construction process.

2.1 Governance and policymaking in China

2.1.1 Hierarchical structure of the government

China's administrative system comprises five tiers: the central government, 31 provincial-level units, 333 prefectures, 2,853 counties, and 40,497 townships.[3] At each level, authority is divided between two parallel branches: the government branch, led by an administrative head (for example, a prefectural mayor), and the Communist Party branch, headed by the corresponding party secretary.

The Organization Department of the Chinese Communist Party Central Committee centrally manages the appointment, evaluation, promotion, and demotion of officials at every tier. It oversees the entire bureaucracy — applying stringent performance metrics and powerful incentives — to ensure that officeholders, from provincial governors and party secretaries to township chiefs, remain firmly aligned with national priorities and central objectives.

At the apex, the State Council and the Chinese Communist Party's Politburo and its Standing Committee in Beijing formulate national policies, develop five-year plans, and set broad strategic objectives. At the sub-national level, provincial, prefectural, and county governments — and their corresponding party committees — administer policy portfolios that blend directives from the central government with initiatives conceived locally.

3. Below the township level, urban communities and rural villages establish self-governing councils that support policy implementation, although they are not formal government entities. In practice, township governments primarily execute policies delegated to them, while substantive policymaking occurs at the county level or above.

This paper examines the division of labor between central and local governments. We focus on the policy portfolios of prefectural governments — the most detailed level at which a comprehensive policy record is available — to distinguish between centrally mandated policies and those initiated locally (the latter potentially originating from any sub-national level).[4]

2.1.2 Policymaking by the central and local governments

Both central and local governments can initiate new policies. Some policy domains, such as foreign affairs and national defense, remain the exclusive domain of the central government, whereas others, like community services and local governance, fall entirely under local authority. Nevertheless, most policy domains — from economic development to education to environmental protection — are governed through a collaborative process between central and local governments, with the exact degree of centralization varying substantially by policy domain, by region, and over time.

In the case of a central policy initiative, the proposal first appears in national documents — five-year plans, work reports or central policy directives — and only later shows up in local policy documents and work reports after local adoption (if at all). Although some central directives specify firm timelines for nationwide roll-out, the majority leave both which jurisdictions participate and when they implement the policy to local discretion. This latitude produces substantial variation in the coverage and speed of central policy diffusion.

Local policy initiatives, on the other hand, emerge in local government documents and work reports. Once a policy is launched locally, neighboring jurisdictions may observe it and adopt it themselves, generating horizontal diffusion. At some point, the central government might take notice and respond in one of three ways: *(a)* veto the policy and stop it from further implementation in any locality, *(b)* endorse it for broader experimentation across other localities and then evaluate whether it is suitable for national roll-out, or *(c)* explicitly elevate it to the national level by turning it into a central policy directive.

Accordingly, any policy implemented by a given locality in a given year can be classified by the degree of central involvement as (a) centrally initiated, (b) locally initiated and later adopted by the central government, or (c) locally initiated without any central

4. As detailed in Section 2.2, prefectural governments' policy portfolios are extracted from their own summaries of implemented policies. These portfolios include policies initiated or implemented by the prefecture itself, as well as (a) policies initiated by the corresponding provincial government and implemented by the prefecture, and (b) policies initiated and implemented by county governments within that prefecture. Therefore, prefectural governments' policy portfolios comprehensively reflect the landscape of bottom-up policy initiatives.

adoption. The combined share of categories (a) and (b), relative to (c), provides a useful proxy for the level of the central government's involvement in China's policymaking landscape.

2.1.3 Policy documents as a pillar for policymaking

A critical pillar of policymaking across all levels of the Chinese government is the issuance of policy documents. These documents can take various forms, ranging from national and local Five-Year Plans, to government work reports, to specific policy directives. We focus on the policy documents issued by the executive and administrative branches of the Chinese government, hence distinguishing policies from law.[5]

These policy documents — serving as a key policy instrument — are far more than aspirational statements. They constitute the authoritative medium through which mandates are issued, monitored, and enforced across every administrative tier. While policy documents technically do not enjoy the permanent, universal status of law, their legitimacy derives from a stringent chain of command within the government system, which in turn underpins the enforceability of policies. Therefore, the issuance and enactment of these policy documents are fundamental to the state's entire policy-execution and compliance apparatus.[6]

Policy documents can be issued at every level of government, but their scope and authority vary with the issuer's rank. At the national level, the State Council promulgates overarching regulations and guidelines that set nationwide priorities and coordinate cross-sector initiatives; beneath it, individual ministries and commissions issue more narrowly tailored directives within their own policy domains. Local governments then adopt these central directives and supplement them with jurisdiction-specific implementation rules and action plans. Moreover, provincial, prefectural, and even county-level authorities — including local branches of each ministry — can issue independent policy documents that carry binding force only within their own jurisdictions; they can neither supersede higher-level directives nor apply outside the area that issued them.

5. Legal documents in China are statutes enacted by the National People's Congress or its Standing Committee, promulgated by presidential order, and bound by a rigorous legislative process that grants them stable, universal legal force. Policy documents, by contrast, are issued by government organs through more flexible, non-statutory procedures, may target specific sectors rather than the entire populace, and can be updated relatively flexibly to reflect evolving priorities.

6. The salience of policy documents in China is often described as "governing by policy documents" (*Wen Jian Zhi Guo*), reflecting the belief that these texts play a more central role than formal laws in China's governance system.

2.2 Construction of policy data

We construct our core measures of policies, and their initiation and diffusion based on policy documents. Our baseline policy documents dataset combines two primary sources.

First, we assemble 3,454,306 policy documents issued by central, provincial, and prefectural governments between 1980 and 2023 from *PKULaw*, a leading legal search engine hosted by Peking University Law School and widely used by lawyers, judges, and academics.[7] Among these documents, 421,951 were issued by the central government. Appendix Figure A.1, Panels (a) and (b), provide examples of central and local policy documents. The *PKULaw* dataset comprehensively captures all policy documents issued at prefecture level and above, regardless of the eventual status of the policy documents themselves. For example, 229,323 policy documents (6.64%) were issued and later officially voided, and they remain in the policy document database.

Second, we compile the complete set of annual prefectural government work reports for 2004–2020 from the *Renmin* database.[8] These reports systematically enumerate the policies implemented in each prefecture in a given year in standard format, which we then match to the corresponding policy document(s) and supplement with additional local socioeconomic indicators.

We describe these data sources in detail below and explain our choices for measurement and variable construction for the empirical analysis.

2.2.1 Identifying distinct policies

A critical step in our empirical strategy is to accurately identify *policies* based on the corpus of millions of government policy documents. For our baseline sample, we focus on the initiatives that governments themselves recognize as policies, rather than imposing an external definition of what is or isn't a policy. Specifically, we extract a comprehensive set of policy-related keywords from annual prefectural government reports spanning 2004–2020, and apply this lexicon to systematically flag relevant initiatives across the entire document set.

As shown in Appendix Figure A.1, Panel A, prefectural government reports follow a standardized format. Over the past 20 years, Section 1 always begins with a recap of policies implemented during the previous year, while Section 2 outlines plans for the upcoming year. For our purposes, we focus exclusively on the recap section to exclude policy ideas that are mentioned but never implemented.

7. For more details about `www.pkulaw.com`, see Wang and Yang (2025).
8. Available at: https://data.people.com.cn/.

To extract keywords from sentences, we first compile a stop-word list — including terms like "enhance," "further," and "implement" — so that we capture only the core elements that distinguish one policy from another. Each candidate keyword is validated in two stages: first through manual review by research assistants, and then by ChatGPT-4o, which assesses whether the term can stand alone as a meaningful policy keyword. For instance, we exclude "Promoting Environmental Protection" from our dataset but keep "river chief scheme" and "ecological red line policy." We regard the former as a vague and overly broad slogan, while the latter two as solid agendas that refer to specific campaigns and policy actions. Any keyword rejected in either step is removed from the master dataset.

This data construction process enables us to reconstruct the full policy portfolio of any locality for any given year. For each policy, we record its title, full text, issuing authority, effective date, area of law, and legal status as of December 2023. Overall, we identify 115,679 distinct policies *implemented* between 2004 and 2020, and on average, each prefectural government implements 1,479 policies per year during this period.

Beyond the aforementioned approach's advantages of requiring few assumptions and allowing for straightforward implementation, two additional considerations regarding our baseline policy sample are potentially relevant for the empirical analysis. First, does the semantic naming of policies by different governments reflect the most appropriate level of aggregation? If governments have incentives to oversell their policy innovations, some policy keywords may be overly broad, masking important variation. Conversely, if governments strategically differentiate otherwise similar policies by using different names, keywords may be too narrow, artificially splitting a coherent agenda.

Second, while government work reports provide a well-structured source for extracting policy keywords, one might wonder whether key information is lost by relying on these summaries rather than on the policy documents themselves. Specifically, are there locally enacted policies that do not appear in annual reports?

To address these questions and evaluate the robustness of our main findings, we explore three alternative ways to construct the policy sample. First, we disaggregate bundled policies by domain so that each policy–domain pair constitutes a distinct initiative comparable to the rest of the sample.[9] This approach yields 651,488 policies. Second, we group policies that are sufficiently similar based on the likelihood of co-occurrence in pol-

9. For example, the "Rural Revitalization Campaign," a large policy bundle initiated by multiple central government units, is broken down into "Rural Revitalization + Agriculture," "Rural Revitalization + Transportation," "Rural Revitalization + Education," "Rural Revitalization + Commerce," etc. For each bundled policy, its corresponding domains are defined as the set of local department types that issue documents related to its topic.

icy documents for keyword pairs. This approach reduces our baseline sample to 101,966 policies. Third, we extract keywords directly from the universe of policy titles — bypassing government work reports. This approach yields 167,705 distinct policies during our sample period.

Our main findings remain highly robust under these three alternative policy definitions, suggesting that the baseline approach sufficiently captures key variations in the data without requiring strong assumptions. In Appendix B, we describe these alternative constructions in greater detail.

2.2.2 Tracing policy's origins and diffusion

Once the policy keywords are extracted, we search for these keywords across the full corpus of policy documents to identify their initiation and track their diffusion. Specifically, we use the Aho–Corasick algorithm to trace each policy across different policy documents.[10]

Linking various policy documents based on policy keywords allows us to define the entire life cycle of each policy ever implemented in China between 2004 and 2020 — from local policy initiation, to horizontal diffusion, to vertical adoption, and eventually to national roll-out, if at all. Note that since we match policies to policy documents starting from 1980, we are able to capture the origin of policies much earlier than the 2004-2020 implementation window. To illustrate the output of this process, Appendix Table A.1 presents a random sample of policy keywords along with their year and location of initiation, while Appendix Table A.2 further summarizes the main contents of these policies.

On average, each policy idea appears in 22.92 documents issued by distinct government branches. Once initiated, policy ideas take on average 5.38 years to reach the first half of localities that adopt the policy, and a total of 9.47 years to reach all adopted localities. Among local policy ideas that are eventually adopted by the central government, on average, it takes 5.73 years from a policy's initiation for the central government to become involved.

2.2.3 Supplementary data

Data on policy outcomes: economic performance Policy objectives are inherently multidimensional. To evaluate policy effectiveness, Section 6 focuses on industrial policies

10. The Aho and Corasick (1975) algorithm is a linear-time string-search algorithm for finding all occurrences of multiple patterns $P_1, P_2, \ldots, P_k$ in a text T. It constructs a finite automaton (trie with failure links) in $\mathcal{O}(\sum |P_i|)$ time and processes T in $\mathcal{O}(|T|)$ time, reporting all pattern matches. The automaton augments a trie of the input patterns with failure links—pointers that mimic the fallback behavior of the KMP algorithm, allowing efficient backtracking when partial matches fail.

in our sample — specifically, those explicitly aimed at fostering the growth of particular industries. We classify industries using the four-digit codes of China's national standard GB/T 4754 (2017), a hierarchical system maintained by the National Bureau of Statistics for categorizing economic activities.[11] This classification enables us to match industrial policies with the economic performance measures of the corresponding industries.

We draw equity investment data from the Business Registration Database, maintained by the State Administration for Industry and Commerce. This comprehensive registry covers over 250 million firm records from 1980 to 2023 and includes detailed information on firm location, ownership structure, legal representatives, shareholders, executives, registered capital, industry classification, founding year, and subsequent updates. Each firm is assigned a four-digit industry code corresponding to its primary business activities.

We measure firm-level performance using the Annual Survey of Industrial Enterprises, maintained by the National Bureau of Statistics. It contains detailed input and output data for all manufacturing firms with annual revenue exceeding 5 million RMB before 2011 and 20 million RMB thereafter. Each firm is classified into a four-digit industry code by the National Bureau of Statistics.

We use export revenue data from the Detailed Records of Imports and Exports, compiled by the General Administration of Customs of China, which report city of origin and HS codes for all export transactions from 2000 to 2016. HS codes are mapped to UN ISIC codes and then converted into four-digit industrial codes using official Chinese documentation. To improve accuracy, we apply NLP methods, computing BERT-based embeddings of code descriptions and matching codes by semantic proximity.

Finally, we collect patent information from the China National Intellectual Property Administration. The patent dataset covers approximately 11 million patents filed by Chinese companies between 1990 and 2020. Each patent is linked to an industry classified by a four-digit code. Additionally, we link the patent data to the Business Registration Database based on the filing party's identifier.

Data on politicians Following Wang and Yang (2025) and Wang, Zhang, and Zhou (2020), we compile detailed biographical information on the universe of Chinese central ministers and local (provincial and prefectural) leaders over our two-decade sample period. For each politician, we record hometown, date of birth, education level, current title, past work history, and other observable characteristics.

11. This system is analogous to NAICS in the U.S. and NACE in the EU, with the four-digit level providing the most granular and widely used classification.

3 Decentralized policymaking

In this section, we describe the landscape of decentralized policymaking in China, and investigate the driving forces behind decentralized policy innovation and diffusion.

3.1 Patterns of decentralized policymaking

Origins of policies We begin by examining the origins of the 115,679 distinct policies identified between 2004 and 2020. Of these, 20,994 (18.15%) were first introduced by the central government. For example, in 2005 the Ministry of Education announced a policy to provide "full tuition waivers for primary and secondary education in rural areas" — a centrally initiated policy with no local precedent.

The vast majority of the policies, 94,685 (81.85%), originated with local governments. Among these locally initiated policies, 24,322 (25.68%) were eventually promoted as national policies, or introduced as centrally-led policy experimentation for further evaluation (Heilmann 2008; Wang and Yang 2025).[12]

Table 1, Panel A, reports the decomposition of policies by origins. Panel B disaggregates the statistics by policy domain (e.g., industrial policy, finance, education, public health), and Panel C uses alternative definitions of policy as described in Section 2.2.

Diffusion of policies Among the 94,685 locally initiated policies, 29,957 (31.64%) were "one-off" measures implemented only in their place of origin and never got adopted elsewhere. For instance, Guangdong's 2018 policy on "Industrial Clusters for Petrochemicals, Energy, and Advanced Materials" was not taken up by other jurisdictions. The other 64,728 policies (68.36%) diffused to at least one additional locality. Zhejiang's 2005 "Village Shareholder System" is one such case, later adopted by 25 other provinces. On average, a locally initiated policy is adopted by 3.76 other localities within its first three years. In comparison, policies initiated by the central government spread more widely, but still fall far short of national diffusion: an average top-down policy reaches 15.74 prefectures within three years of its introduction.

This incomplete reach of top-down initiatives, combined with the predominance of bottom-up policy initiation, produces highly decentralized local policy portfolios: in a given year, 62.9% of the policies implemented in a prefecture originated from bottom-up sources and never involved any central-government endorsement.

12. Many policies spread widely before central involvement. For example, Zhejiang's 2006 "Domestic Waste Disposal System" reached multiple provinces and cities between 2007 and 2009 before being adopted nationally in 2010. Since then, more than 130 prefectures have implemented it.

The discretion for local bureaucrats to selectively enforce central directives on the extensive margin — echoing O'Brien and Li (2017) — is further amplified by intensive margin adjustments. For any given policy, we compute the textual similarity between the central document and each local version using the cosine similarity measure, TF-IDF weighting scheme, and standard Chinese stop-word removal procedure. The resulting score ranges from 0 (completely different) to 1 (identical). The sample average of 0.141 indicates substantial local adaptation and tailoring by grassroots officials in the policy diffusion process.

All three measures — limited diffusion of central policies, a low share of centrally originated policies in local portfolios, and a high degree of local tailoring during diffusion — illustrate the decentralized nature of China's policymaking landscape.

3.2 Local bureaucrats' roles in decentralized policymaking

Next, we examine local bureaucrats' roles in policymaking. Section 3.2.1 focuses on their roles in initiating new policies, and Section 3.2.2 examines their roles in facilitating the diffusion of policies across localities.

3.2.1 Local bureaucrats and local policy innovation

Measuring locality's policy innovation and compliance We construct two indices to measure each locality's policy innovativeness and compliance in a given year. The *bottom-up policy innovation index* is defined as:

$$\text{Innovation}_{i,t} = \frac{1}{|U|} \sum_{p \in U} \frac{\text{totalAdopt}_p}{\text{ranking}_{i,p}}, \tag{1}$$

where U is the set of bottom-up policies that prefecture i carried out in year t, totalAdopt_p is the total number of prefectures adopting policy p, and $\text{ranking}_{i,p}$ is the order in which i adopted p. This measure simultaneously captures how quickly a locality acts on a policy idea, and how important that idea turns out to be: for example, if prefecture i is the initiator of policy p ($\text{ranking}_{i,p} = 1$), and policy p eventually becomes a national policy (totalAdopt_p = the total number of prefectures in China), then the innovation index will be driven up accordingly.

The *top-down compliance index* is defined analogously for centrally initiated policies:

$$\text{Compliance}_{i,t} = \frac{1}{|\tilde{U}|} \sum_{p \in \tilde{U}} \frac{\text{totalAdopt}_p}{\text{ranking}_{i,p}}, \tag{2}$$

where $\tilde{U}$ is the set of top-down policies that prefecture i carried out in year t.

These two indices allow us to quantify the extent to which politicians or localities initiate new policy agendas or adhere to central directives in a given year. For example, during Xi Jinping's tenure as Zhejiang's Party Secretary, the province's innovation index reached 3.31, ranking first among all provinces.[13] Xi also exhibited a high level of compliance, ranking sixth nationwide.

Appendix Figure A.2 plots average innovation and compliance indices across localities. Innovation displays greater regional inequality, with activity concentrated in coastal areas and a clear east–west divide. Compliance is more evenly distributed, with less geographic clustering. Several top-performing prefectures are located in Hebei, indicating that strong adherence to central policy is not confined to the most developed regions.

Localities vs. bureaucrats in driving policy innovation A natural question is whether policy innovation is primarily driven by innovative bureaucrats or if some localities inherently provide a more nurturing environment for innovation. To address this, we exploit the fact that local bureaucrats in China are frequently rotated across localities, and follow the approach described in Abowd, Kramarz, and Margolis (1999) to separately identify bureaucrat fixed effects and locality fixed effects in driving policy innovations:

$$Y_{ijt} = \alpha_i + \Psi_{j(i,t)} + \gamma_t + \varepsilon_{ijt}, \tag{3}$$

where Y_{ijt} is the policy innovation (or compliance) index of prefecture i, led by bureaucrat j, in year t. α_i is prefecture fixed effect, $\Psi_{j(i,t)}$ is bureaucrat fixed effect, and γ_t is year fixed effect. Standard errors are estimated via nonparametric bootstrap based on 1,000 replications. Those parameters are identified from bureaucrats who moved across localities.

As shown in Table 2, Panel A, bureaucrat fixed effects explain five times more variation in the innovation index than locality fixed effects.[14] This suggests that bureaucrats, rather than localities, play the central role in shaping bottom-up policy innovation. According to Table 2, Panel B, qualitatively similar patterns emerge for the top-down compliance index, where bureaucrat fixed effects also explain more variation than locality fixed effects. These patterns are robust to alternative definitions of policies (see Appendix Table A.3).

Intriguingly, in addition to bureaucratic fixed effects, year fixed effects also appear

13. This was driven by policies he initiated that later diffused widely, such as the fiscal expenditure performance evaluation" program (launched in 2006 and adopted nationally in 2011) and the commercialization of technological innovation" program (launched in 2005 and eventually adopted by 24 other provinces).

14. As Andrews et al. (2008) point out, if prefectures are weakly connected to one another because of limited mobility of politicians across localities, the AKM estimates of the contribution of locality effects to variance of innovation are biased *upwards*, while estimates of the contribution of politician effects to variance are biased *downwards*. Therefore, our estimates of the relative importance of politician fixed effects with respect to locality fixed effects can be interpreted as a lower bound.

important in shaping policy innovation and compliance, highlighting the importance of evolving policymaking dynamics, which we discuss in greater detail in Section 4.

3.2.2 Local bureaucrats and decentralized policy diffusion

In addition to driving bottom-up policy innovation, local bureaucrats may also play a significant role in shaping decentralized policy diffusion across jurisdictions. To investigate this, we examine *(i)* how local political turnover affects the diffusion of policies initiated in a locality, and *(ii)* how political rivalry between bureaucrats influences the diffusion of policies between their jurisdictions.

Political turnovers and policy diffusion First, we examine how policy diffusion evolves after a prefectural leader departs from his position. Specifically, we estimate the following event study model:

$$Y_{pit} = \sum_{T} \beta_T T_{it} + \phi_p + \lambda_t + \varepsilon_{pit}, \tag{4}$$

where Y_{pit} is the number of adoptions in year t for policy p, which was initiated by prefecture i. T_{it} represents the event study dummy variables: T_{it} equals one if, in year t, T years have passed since prefecture i experienced a turnover of its political leader, and zero otherwise. We further control for the full set of policy fixed effect and year fixed effect, and cluster the standard errors at the prefecture level. To demonstrate robustness, Appendix Figure A.3 shows that the results hold when we, in Panel A, follow Sun and Abraham (2021) to account for heterogeneous treatment effects in staggered event-study designs, and in Panel B, apply more restrictive definitions of "promotion."

As shown in Figure 1, the future departure of a local bureaucrat — resulting from rotation, promotion, demotion, or retirement — is orthogonal to prior trends in the diffusion of locally initiated policies. However, once the local bureaucrat leaves their position, there is a stark 41.6% reduction in the speed of diffusion, even when it is the same policy, and it never recovers to baseline levels in the subsequent years. This pattern exists for all types of bureaucratic turnover, and is particularly pronounced when the departing bureaucrat is *not* promoted to a higher position.[15] These patterns are again robust to alternative definitions of policies (Appendix Figure A.5).

This finding suggests that local policy diffusion is largely driven by incentivized bureaucrats actively promoting their own innovations. Once the originator of a new policy

15. Similar patterns appear we examine the departure of central government officials. Appendix Figure A.4 shows that the number of adopters decreases by 22.9% following the departure of the minister who initiated the policy.

can no longer claim full political credit for its subsequent success, the incentive to promote it diminishes, and diffusion effectively stalls.

Economic and political determinants of decentralized policy diffusion We next examine the decentralized policy diffusion among local governments more generally. We define policy similarity for all prefecture pairs in China, and investigate its determinants. Specifically, for each prefecture i in each year t, we construct a vector representing its entire policy portfolio covering all policy dimensions: $\vec{V}_{it} = (v_{i1t}, v_{i2t}, v_{i3t}, ..., v_{iNt})$. We then calculate, for each prefecture pair in a given year, their similarity in policy portfolios, as measured by the (opposite) distance between their policy portfolio vectors: $S_{ijt} = -||\vec{V}_{it} - \vec{V}_{jt}||$.

For each pair of prefectures, we further compute their economic proximity in a given year ($\text{Proximity}_{ijt,econ}$), as measured by similarity in per capita GDP;[16] and political proximity in a given year ($\text{Proximity}_{ijt,pol}$), as measured by the (opposite) Mahalanobis distance between two politicians' key characteristics.[17]

We estimate the following equation:

$$S_{ijt} = \alpha \cdot \text{Proximity}_{ijt,m} + \lambda_t + \gamma_i + \sigma_j + \varepsilon_{ijt}, \tag{5}$$

where $m \in \{\text{Econ}, \text{Pol}\}$ denotes different measures of proximity. λ_t stands for year fixed effects, γ_i and σ_j represent prefecture fixed effects. The standard errors are two-way clustered at the origin and destination levels.

As shown in Panel A of Figure 2, economic similarity between prefectures is a strong, positive predictor of policy portfolio's similarity. This is consistent with policy diffusion as a result of predicted or observed policy outcomes: if a policy has been demonstrated to work in another locality with similar socioeconomic conditions, it is more likely to succeed here.

In stark contrast, as shown in Panel B, when two prefectural leaders are more politically similar — indicating increased competition for promotions — policy portfolio similarity declines significantly. This suggests that local bureaucrats, in order to avoid enhancing their competitors' credentials, tend not to learn from their peers who are close competitors, thereby distorting the policy diffusion process. These findings echo ample anecdotal accounts of how political rivalry prevents the diffusion of suitable policies — for example, why Beijing rejected Shanghai's widely praised program of automobile li-

16. Our findings are robust to using alternative measures of economic similarity, based on fiscal income, unweighted GDP, fiscal expenditure etc.

17. Specifically, their start-age, hierarchical status, gender, ethnicity, education, central government experience as well as previous experience.

cense plate auctions (see Appendix C for details).

Appendix Table A.4 quantifies these patterns: a one-standard-deviation increase in economic similarity is associated with a 1.04-point (2.5%) rise in policy similarity, whereas a one-standard-deviation increase in political similarity corresponds to a 0.08-point (0.2%) decline in policy similarity. For the political similarity analysis, we can control for a stringent set of prefecture-pair fixed effects, holding constant the baseline rate of policy diffusion between any two localities and exploiting only the variation in political similarity generated by bureaucratic turnovers over time. Our findings remain robust under this more demanding specification, suggesting a causal role for political competition in policy diffusion.

Our results are also robust to alternative measures of policy diffusion. Rather than comparing policy portfolios across localities each year, we can explicitly account for the direction of diffusion. Specifically, we define an instance of policy diffusion (Diffusion_{ijt}) as prefecture i adopting, in year t, a policy that was previously initiated by prefecture j. We then estimate the following model:

$$\text{Diffusion}_{ijt} = \beta \cdot \text{Proximity}_{ijt,m} + \gamma_{ij} + \lambda_t + \varepsilon_{ijt}. \tag{6}$$

Appendix Table A.5 presents the results. We find that while economic proximity again positively predicts policy diffusion, greater political competition reduces its likelihood. In more stringent specifications, we include politician-by-prefecture fixed effects, isolating variation arising solely from political turnover in adopting prefectures. This approach allows us to assess whether changes in local political leadership affect a prefecture's propensity to adopt policies from an origin prefecture whose leadership remains unchanged. The results remain robust: a one-standard-deviation increase in political competition corresponds to a 1.2% decrease in the probability of policy diffusion.

Importantly, economic similarity between two prefectures is strongly and positively associated with their political similarity, as shown in Appendix Figure A.6.[18] This indicates that strategic political competition could impede the most beneficial form of policy diffusion — namely, diffusion among economic neighbors. Moreover, Appendix Table A.6 reveals a negative interaction between political and economic similarity in facilitating policy diffusion. This implies that strategic competition is most intense when potential returns are highest, further exacerbating the distortion.

Taken together, the findings in this section underscore the pivotal role local bureau-

18. Politicians with similar backgrounds tend to be assigned by the central government to localities with comparable economic conditions — an expected pattern if the central government aims to maximize an objective that depends on both local conditions and bureaucratic capabilities.

crats play in China's decentralized policymaking process. They are the driving force behind local policy innovation and shape the diffusion of policies across jurisdictions. However, their competitive incentives can also introduce distortions into decentralized policy diffusion.

4 The turn toward centralization

As foreshadowed by Table 2 in the previous section, local policymakers' innovation and compliance tendencies appear to have undergone significant temporal shifts. In this section, we document and explain the notable changes in China's policymaking process over the past decade. We begin by describing the salient trend toward centralization in policymaking (Section 4.1), and then examine how it may have resulted from broader shifts in local bureaucrats' political incentives (Section 4.2)

4.1 Increased policymaking centralization after 2013

We first examine the diffusion of new top-down and bottom-up policies, measured by the number of prefectures that adopt each within three years of its introduction. We plot this year by year in Figure 3, Panel A. Before 2013, a typical central policy reached about ten prefectures; since then, that number has nearly tripled, reflecting a substantial rise in local compliance with the central policy agenda. By contrast, bottom-up policies have continued to diffuse to roughly five prefectures — a rate essentially unchanged before and after 2013. This pattern underscores that the post-2013 rise in centralization stems primarily from greater local adherence to centrally initiated policies, effectively crowding out local innovation.

Next, we examine how local bureaucrats allocated effort between innovating bottom-up policies and complying with top-down policies, and how this balance has changed over time. In Panel B, we plot the relative share of bottom-up versus top-down policies in an average prefectural government's portfolio in each year. The share of top-down policies remained relatively stable at around 30% prior to 2012, then rose sharply to 42% between 2013 and 2020 — a 40% increase from the baseline, suggesting a notable shift toward policy centralization over the past decade.

Panel C turns to textual similarity between central policy documents and their local counterparts, capturing the degree to which local enforcement mimics the central directive. Similarity roughly doubles after 2013. This is not merely a case of "window-dressing": the pattern holds when calculated using only substantive policy content, ex-

cluding generic slogans (Appendix Figure A.7). While local governments are expected to adapt policies to regional conditions, the post-2013 increase in similarity points to a narrowing of local discretion.

Taken together, the trends in portfolio composition, diffusion, and textual similarity reveal a marked decline in bottom-up innovation relative to compliance with top-down directives, signifying an accelerated shift toward policy centralization.

Importantly, the sharp turn toward centralization we observe in policy diffusion and localization is largely obscured when we examine only coarse time-series indicators, such as the annual share of centrally initiated policies or the volume of centrally published documents (Appendix Figure A.8). These aggregate measures overlook the pivotal role of individual bureaucrats: when we trace complete policy lifecycles, link them to bureaucrats' tenure, and shift the unit of analysis from the year to the bureaucrat-year level, the trend toward centralization becomes markedly sharper. This suggests that to understand the driver of the observed centralization, one has to focus on the (evolving) roles that local bureaucrats play in the policymaking process, which we turn to next.

4.2 Understanding the turn toward centralization

Qualitative accounts suggest that, after Xi Jinping assumed power in late 2012, authority was rapidly consolidated within the central government, suppressing local policy initiatives and experimentation (Heilmann 2018; Naughton 2021). In particular, the heightened rewards for local compliance and the broader diffusion of central directives mirror Beijing's own critique at the time that "government orders never leave Zhongnanhai" (the central leadership compound in Beijing) — a problem Xi prioritized resolving upon taking office.[19] We view this as a broad shift that redirected local bureaucrats' incentives away from policy innovation and toward compliance with central authority. We explore a number of key manifestations of such change.[20]

Promotion incentives: reward for innovation vs. compliance In order to stand out among their peers in the policymaking process, local bureaucrats can potentially allocate their effort across two visible dimensions: (a) initiate innovative new policies from bottom-up that can diffuse widely or even get picked up by the central government for

19. See: Voice of America; source: https://www.voachinese.com/a/regulatory-system-china-20130123/1589175.html.

20. While these qualitative and quantitative lines of evidence consistently highlight the top leadership's role in shifting bureaucratic incentives toward compliance and centralized policymaking, it is important to acknowledge that other factors — such as evolving internal and external conditions of the Chinese economy — may also have influenced policy outcomes during this period.

national roll-out; and (b) be an early adopter of top-down policies assigned by the central government to demonstrate compliance and loyalty.

To understand local policy decisions, it is therefore important to examine the incentive schemes faced by local bureaucrats: are they rewarded for policy innovation, or policy compliance?

To answer this question, for each given year, we identify all the prefectural leaders that finished their terms within a 5-year moving window around it, and compare the prefectural leaders that received promotion versus those that did not by estimating the following equation:

$$Promotion_i = \alpha + \beta_1 \cdot innovation_i + \beta_2 \cdot compliance_i + X_i\Gamma + \varepsilon_i, \tag{7}$$

where $Promotion_i$ is a binary variable indicating whether prefectural leader i was promoted to a higher position by the end of their term. For the right-hand side variable, $innovation_i$ and $compliance_i$, we compute the average innovation/compliance index (defined in Section 3.2) for each prefectural leader i over their entire term in office. Control variables X_i include their year of departure and their official rank within the hierarchy.

Figure 4 plots the estimated coefficients by year. As shown in Panel A, from 2005 to 2012, the innovation index positively predicted subsequent promotion, with this effect intensifying over time. After 2013, however, the relationship between promotion and the innovation index declined rapidly and converged to zero within three years. These patterns support the interpretation that, prior to 2013, local bureaucrats were rewarded by the central government for policy innovation, a practice that ceased after 2013.

Panel B presents a mirror image for the rewards to the compliance index. Between 2005 and 2012, there was no significant correlation between the compliance index and subsequent promotion; however, after 2013, this correlation became positive and significant. This finding suggests that the political incentive scheme shifted from rewarding policy innovation to rewarding compliance after 2013.

Appendix Table A.7, Panel A, quantifies these graphical patterns. Before 2013, a one standard deviation increase in the innovation index was, on average, associated with an 8.2% increase in the likelihood of promotion, while the correlation became statistically indistinguishable from zero in the post-2013 era. In contrast, prior to 2013, the compliance index was uncorrelated with promotion; thereafter, a one standard deviation increase in the compliance index was associated with a 7.8% higher chance of promotion. These estimates remain robust after controlling for the bureaucrats' cohort, the hierarchical level of the prefectural cities, as well as the bureaucrats' age and education levels.

Importantly, the shift from innovation to compliance in predicting promotion is ro-

bustly observed even after controlling for local GDP growth, a strong, positive predictor of promotion due to the underlying political tournament. In other words, variation in policy innovation and compliance across localities does not merely influence local bureaucrats' political promotion through changes to corresponding GDP performance. Instead, the central government additionally rewards innovation for its positive information externalities and compliance as a signal of loyalty, beyond their roles in driving economic performance.[21]

Panel B shows that our results are robust to alternative proxies for innovation and compliance. Specifically, instead of the innovation index, we count the number of locally initiated policies that became national policies in the last three years; and instead of the compliance index, we count the number of times a locality adopted a national policy within the first three years of inception. Our main empirical patterns persist under these alternative measures. Panel C further demonstrates that the main findings remain unchanged when we construct new innovation and compliance indices accounting for the total number of policies implemented in each locality–year.

In Appendix Table A.8, we repeat the same exercise, replacing the outcome variable with whether an official is investigated for corruption. This is motivated by the popular perception that China's anti-corruption campaign since 2013 was used by top leadership for factional struggle and power consolidation (e.g., Lorentzen and Lu (2021)) — suggesting that "sticks," in addition to "carrots," might be deployed to advance the central policy agenda. However, we find *no* significant association between policy innovation or compliance and anti-corruption investigations, either before or after 2013, which is inconsistent with this hypothesis.

Consolidation of central policymaking bodies A key intervention toward centralization in policymaking was Xi's establishment of various Central Leading Groups. For example, in 2013 Xi created the Central Leading Group for Comprehensively Deepening Reform, which he chaired himself, enabling it to bypass the State Council and directly advance his policy agenda. The creation of such a body could send a salient signal to local officials that particular policies were personally supported by Xi, making compliance in corresponding domains a high-stake signal of loyalty.

Exploiting the staggered establishment of six such groups across different policy domains (excluding those irrelevant to local policy, such as the Central Leading Group on

21. Appendix Figure A.9 shows the relationship between local GDP growth performance and corresponding bureaucrats' promotion prospects. One observes a positive relationship between GDP growth performance and promotion throughout the period from 2004 to 2020, though the relationship becomes notably weaker in recent years.

National Defense), Appendix Figure A.10 presents event-study estimates showing the degree of policy centralization relative to the timing of the leading groups' establishment. We find that these groups were targeted at policy domains experiencing increasing decentralization, and once the leading groups were established, they were in general effective in reversing those decentralization trends toward significant centralization.

5 Policy-locality suitability

So far, we have documented that China's policymaking process relies substantially on bottom-up innovation (Section 3), but has been quickly shifting toward greater centralization since 2013 (Section 4). What are the impacts of this turn toward centralization in policymaking?

In this section, we introduce our measures of policy–locality suitability — a key concept for understanding centralization — which allow us to evaluate how shifts toward centralization affect policy effectiveness. The performance of a policy in a given locality depends critically on how well its design and implementation align with local conditions. For example, promoting the mass installation of solar panels in a rainy area is likely to prove futile. Likewise, data centers become more expensive to operate without a cool ambient climate and reliable electricity.

We begin by describing the sample of industrial policies, among which we could measure policy-locality suitability and policy effectiveness (Section 5.1). We then describe the measures of policy-locality suitability (Section 5.2), and we document the key relationship between policy-locality suitability and policy effectiveness (Section 5.3).

5.1 Zooming in on industrial policies

To systematically evaluate the consequences of policy centralization, in the remainder of this paper, we restrict our analysis to industrial policies.

Our baseline sample covers around 115,679 distinct policies across all policy domains, which makes it particularly challenging to holistically quantify policy effectiveness. For instance, consider an education policy that provides free lunches to middle school students; its goal could be to improve students' health outcomes, raise their grades, or reduce dropout rates, etc. Given the multiplicity of policy goals — and without knowing the government's exact objective function — it is difficult to evaluate a policy's effectiveness. Furthermore, comparing policies across domains is even trickier. For example, even if we have a clear measure of an education policy's effectiveness, we lack a system-

atic benchmark to compare it with the effectiveness of a health or environmental policy, which focus on entirely different outcomes.

To circumvent these challenges, we focus on industrial policies. The government explicitly states that its primary goal for enacting industrial policies is to promote the growth of a specific industry, as measured by industrial output, exports, and patent filings — common indicators used by the central government to evaluate industrial policy success (Fu 2015, Glaser 2022, Fang, Li, and Lu 2025).

Specifically, for each 4-digit industry — the most granular classification in the National Bureau of Statistics coding system — we search both its name and that of the corresponding 3-digit industry across all policy documents, coding industrial policy enforcement at the prefecture–industry–year level. We view this baseline definition of industrial policy as transparent and decision-free, and it yields 122,104 instances of industrial policy adoption nationwide during our sample period, spanning 1,215 distinct 4-digit industries. For robustness, we further restrict the sample to economic policy documents that specify policy actions promoting these industries; this refined definition — similar in spirit to Fang, Li, and Lu (2025) and Juhász et al. (2022) — further amplifies our empirical results. We then classify each policy as top-down or bottom-up based on whether the central government had explicitly promoted the industry by that year. As shown in Appendix Figure A.11, industrial policies have exhibited a centralization trend since 2013, similar to that observed for all policies.

We retain this definition of industrial policy at the prefecture–industry–year level throughout the subsequent analysis, as it allows us to quantify policy–locality suitability and examine its effect on policy outcomes. This definition differs from the baseline definition of policies in two notable ways. First, in the baseline sample, each distinct policy keyword is counted separately — for example, "subsidizing EV production" and "constructing EV charging stations" are treated as two distinct policies — whereas for the industrial policy analysis they are grouped together as "industrial policy promoting EV." As a result, 18,733 policies in the baseline sample correspond to the promotion of 1,215 industries. Second, in the baseline sample, a policy implemented in both Shanghai and Beijing would still be treated as the same policy. In contrast, in the industrial policy analysis we distinguish between the two adoptions in order to compare how suitable the same industry is across localities. Using this definition, we identify 122,104 instances of industrial policy adoption at the prefecture–industry–year level.

5.2 Measures of policy-locality suitability

We measure the local suitability of an industrial policy in two different (but related) ways.

Alignment with pre-existing regional supply chains The first approach is to directly measure the pre-determined observable characteristics that make a locality particularly conducive to a specific industry. In particular, for a given industry to grow, one important condition is having easy access to its key upstream suppliers, as demonstrated by the extensive literature on industrial agglomeration effects (Greenstone, Hornbeck, and Moretti 2010; Kline and Moretti 2014).

Motivated by this, for each industrial policy aimed at promoting industry p in prefectural city c, we first directly measure the extent to which c has *pre-existing* strength in the supply chain of industry p:[22]

$$S_{cp} = \sum_{j} \alpha_{jp} \cdot I_{cj}, \tag{8}$$

where α_{jp} represents the share of key upstream industry j in the input composition of industry p, as extracted from China's national input-output table. I_{cj} denotes the accumulated investment in industry j in prefectural city c over the decade preceding the industrial policy. The weighted average S_{cp} thus reflects the *absolute* strength of prefectural city c in industry p, from the perspective of supply chain access.[23]

Furthermore, we benchmark S_{cp} against both prefectural city c's strengths in other industries and other prefectural cities' strengths in industry p:

$$Supply\text{-}chain\ suitability_{cp} = \frac{S_{cp} / \sum_{p' \in P} S_{cp'}}{\sum_{c' \in C} S_{c'p} / \sum_{c' \in C, p' \in P} S_{c'p'}}, \tag{9}$$

where the numerator $\frac{S_{cp}}{\sum_{p' \in P} S_{cp'}}$ measures prefectural city c's strength in industry p's supply chain relative to all other industries, and the denominator measures the rest of the country's strength in industry p's supply chain relative to all other industries. Therefore, the ratio captures the *relative* strength of prefectural city c in industry p, from the perspective of *pre-existing* supply chain access.

22. We retain only the key upstream industries—those supplying more than 10 percent of a downstream industry's inputs. We then refine our focus to spatially sensitive industry pairs, defined as follows. For each downstream–upstream industry pair, we identify the five largest prefectures in each industry. For each downstream prefecture, we compute its distance to the five largest prefectures in the upstream industry. If every downstream prefecture lies within 500 km of at least one upstream prefecture, the pair is classified as spatially sensitive.

23. Our findings remain quantitatively similar when we measure supply chain strength within a 500 km radius rather than within the same prefecture. These results are available upon request.

Alignment with entrepreneurs' revealed preferences From China's business registration records, we extract every equity investment made in a given prefecture-industry-year, capturing both the entry of new firms and additional investments in existing firms. Specifically, we compute I_{cp}, which represents the accumulated equity investment in prefectural city c in industry p during the decade preceding the initiation of an industrial policy. This *ex ante* investment accumulation reflects the absolute level of entrepreneurial enthusiasm for industry p in prefectural city c, as indicated by their business investment decisions.

Based on I_{cp}, we further calculate the entrepreneurs' relative enthusiasm for a given industry in a specific prefecture:

$$Investment\ suitability_{cp} = \frac{I_{cp} / \sum_{p' \in P} I_{cp'}}{\sum_{c' \in C} I_{c'p} / \sum_{c' \in C, p' \in P} I_{c'p'}}, \tag{10}$$

where the numerator $\frac{I_{cp}}{\sum_{p' \in P} I_{cp'}}$ captures, for a given prefecture c, the concentration of investments in industry p relative to other industries. The denominator measures this relative concentration for the rest of the country, so that the resulting ratio reflects the "extra entrepreneur preference" for a particular industry in a specific locality compared to all other industry-locality clusters.

Given this definition, for each industrial policy — whether top-down or bottom-up — that is aimed at promoting industry p in prefectural city c, we can measure the extent to which the policy aligns with or deviates from the preferences of entrepreneurs, as revealed by their business investment decisions prior to the policy's initiation.[24]

Correlations between the two measures of suitability The two suitability measures — one based on supply-chain strength and the other on investment flows — are strongly positively correlated, as shown in Appendix Figure A.12. Importantly, both measures are constructed from data collected *before* the industrial policy was introduced in each prefecture–industry, so neither was directly influenced by the policy itself. In other words, the observed high correlation cannot be driven by the policy's effects; instead, it is consistent with supply chain strength being an important consideration in entrepreneurs' investment decisions.

24. Our construction of "entrepreneurs' revealed relative preference" based on business registration records is similar in spirit to that constructed in Fang, Li, and Lu (2025), but with two important distinctions: (a) we restrict our calculation to investments made within the decade *before* the issuance of an industrial policy, thereby capturing the revealed preference of business investors in the *absence* of policy interventions—preferences that likely reflect local fundamentals such as natural endowments and regional supply chains; and (b) we disaggregate investments by private versus state-owned enterprises to examine whether they align differently with top-down versus bottom-up industrial policies.

Appendix Figure A.13 plots the spatial distribution of these measures of local suitability using the automobile industry as an example. This figure further verifies that the suitability measures — whether inferred from revealed preferences or defined based on supply chains — are highly correlated. It is worth noting, however, that when we measure relative strength based on *ex ante* investment, the versions based on private enterprise investment and state-owned enterprise investment, while correlated, show notable differences. This pattern likely indicates that these two types of firms may have different preferences when making investment decisions, and this distinction allows us to examine whether an industrial policy is more closely aligned with the revealed preferences of private firms or those of SOEs.

5.3 Policy-locality suitability and policy effectiveness

Do industrial policies that are more suitable to local economic conditions actually deliver more desirable policy outcomes to the locality?

As *prima facie* evidence, Appendix Figure A.14 plots the average suitability of adopting localities against years since policy inception. We observe a clear downward trend, indicating that high-suitability localities tend to be early adopters under decentralized policymaking. This pattern suggests that local policymakers indeed value alignment between industrial policies and their localities' economic conditions.

To answer this question more directly, we compare the dynamic impacts of industrial policies that promote industries in line with local strengths versus those that defy local strengths. Specifically, we estimate the following empirical model:

$$Y_{pct} = \sum_{T} \alpha_T T_{pct} + \sum_{T} \beta_T T_{pct} Suitability_{pc} + \phi_{cp} + \gamma_{pt} + \lambda_{ct} + \varepsilon_{pct}, \tag{11}$$

where Y_{pct} are the outcomes of interest for industry p in prefectural city c in year t. T_{pct} represents the event-study dummy variables, which equal one if, in year t, T years have passed since prefectural city c implemented a policy promoting industry p, and zero otherwise. $Suitability_{pc}$ is the continuous suitability score for industry p in prefectural city c, measured using either equity investment or local supply chain, prior to the initiation of an industrial policy. β_T represents the coefficients of interest, capturing the role of industry-locality suitability in determining the effectiveness of an industrial policy. We control for full sets of two-dimensional fixed effects: prefecture-by-year, industry-by-year, and prefecture-by-industry. The standard errors are two-way clustered at the prefecture and industry levels.

We examine three main outcome variables: industrial output, export values, and patent

filings, which are the most frequently mentioned target outcomes of industrial policies in China, and also the common metrics used to evaluate local officials' effectiveness in promoting industrial development.[25]

Figure 5 plots the β_T coefficients over time, across the two suitability measures and three outcomes of interest. One observes a consistent pattern throughout: industrial policies aligned with local strengths, according to pre-existing supply chains or business investment stock, are significantly more effective in delivering industrial growth compared to those that are unsuitable for local conditions.[26] This suitability premium appears to be growing over time, suggesting that the gap in industrial policy effectiveness is likely persistent.[27] This heterogeneity is unlikely to result from endogenous policy selection, since policy suitability exhibits no correlation with pre-intervention trends in policy outcomes.

Appendix Table A.10 quantifies these patterns. Specifically, when an industrial policy is introduced in a locality with a one-standard-deviation higher *ex ante* investment-based suitability measure, it leads on average to 29.2% greater industrial output, 32.3% higher exports, and 19.3% more patent filings. Likewise, a one-standard-deviation increase in suitability based on pre-existing supply chains is associated with 8.0% greater industrial output, 0.9% higher exports, and 8.7% more patent filings following policy initiation.

6 Trade-offs of centralized policymaking

In Sections 6.1 and 6.2, we assess, respectively, the costs and benefits of centralized policymaking through the lens of policy-locality suitability. Leveraging the findings reported in these two sections, in Section 6.3, we quantitatively compare the costs and benefits of China's policy centralization since 2013.

Our analysis of the trade-offs associated with policy centralization focuses on policy suitability for local conditions, a key concern highlighted in the theoretical literature (Tiebout 1956; Oates 1972; Alesina and Spolaore 1997; Bolton and Roland 1997; Besley and

25. It is worth noting that these measures of industrial policy outcomes reflect policy effectiveness through the lens of the (local) government's own goals, which are not necessarily equivalent to maximizing the broader welfare of society.

26. In Appendix Figure A.15, we separately plot event-study estimates for industrial policies in the top decile of suitability versus those below this threshold. Consistent with our triple difference-in-differences results, we observe marked industrial growth following the initiation of high-suitability policies, but see no analogous trend for low-suitability counterparts.

27. Interestingly, Appendix Table A.9 shows that the heterogeneity disappears when we measure policy–locality suitability based on SOEs' revealed preferences (i.e., pre-policy investments by SOEs rather than by private firms). This finding is consistent with the interpretation that, on average, the private sector is better than the state at understanding local economic conditions — such as whether a particular sector may be profitable given regional supply chains — and that following private sector investment patterns can lead to more desirable outcomes in terms of promoting industrial growth and innovation.

Coate 2003; Alesina, Baqir, and Hoxby 2004), and likely of first-order importance in this context. Although examining all potential considerations related to centralization is beyond the scope of this paper, we also briefly discuss several alternative policy objectives that might relate to centralization and assess their empirical relevance, in Section 6.4.

6.1 Centralization's negative impact on policy suitability

Centralized policymaking, by sacrificing valuable local information and initiative and introducing greater incentives for policy compliance with the central government, can reduce the suitability of policies to local conditions and thereby undermine their effectiveness.

Ample qualitative evidence illustrates this risk. For example, China's wind energy industry initially thrived under bottom-up promotion, concentrated in the northwest where high wind density and low land costs supported large-scale farms. After 2013, however, central directives drove expansion into less suitable provinces such as Hunan and Hubei, where weak wind resources and limited transmission capacity produced "ghost wind farms." See Appendix C for a detailed description of this industry.

To assess whether such declines in suitability are systematic, we next turn to our industrial policy sample and quantitatively compare the average policy–locality suitability of top-down versus bottom-up initiatives. Specifically, for each industrial policy, we compute the extent to which it complies with local conditions among the adopters, according to the two suitability measures defined in Section 5. As shown in Table 3, Panel A, among all the industrial policies implemented by a given prefecture in a given year, those assigned from the central government are significantly less suitable for local economic conditions. The results are robust across alternative measures of policy–locality suitability: whether based on *pre-existing* regional supply chain strength or *ex ante* business investment flows, and whether measured as continuous or binary. On average, compared to bottom-up industrial policies, top-down ones promote industries that are 18-22% less suitable for the given prefecture.

To address concerns that some localities may adopt policies in a merely "performative" manner — such as issuing documents that echo central directives without making substantive efforts to promote the targeted industry — we restrict, in Panel B, the sample to industrial policies that explicitly stipulate monetary subsidies for eligible local firms. Within this subsample, the suitability gap between top-down and bottom-up industrial policies remains unchanged, suggesting that our results are not driven by the cheap talk

of local bureaucrats.[28]

Consistent with the cross-sectional comparison between top-down and bottom-up industrial policies, we find that when an industrial policy is explicitly promoted by the central government, the policy–locality suitability of subsequent adopters declines sharply relative to earlier adopters who had taken up the *same* policy through horizontal diffusion among local governments (see Appendix Table A.12).[29] This supports the interpretation that explicit central involvement prompts a number of less-suitable localities to adopt the policy in order to signal compliance and loyalty, even though they would not otherwise do so given their lack of local suitability.

To the extent that policy-locality suitability is strongly associated with policy effectiveness, as documented in Section 5.3, top-down industrial policies would on average end up being less effective in generating industrial growth than their bottom-up counterparts. If the central government aims to effectively promote the growth of the specific industries, it should not, when choosing localities *within* China, allocate industrial policies to those lacking relevant supply chains or deemed less suitable by private investors.

6.2 Centralization's positive impact on policy suitability

In Section 3.2, we document that political competition among local bureaucrats distorts and impedes decentralized policy diffusion.[30] When local politicians compete for the same promotion opportunities, they may avoid learning from one another — so as not to enhance their competitors' credentials — and thereby forgo opportunities to adopt otherwise suitable and effective policies. Holding the policy capacity of a locality constant, strategic political competition leads to substitutions in policy adoption. Instead of adopting policies from one's political competitors — which have been proven effective in jurisdictions with comparable socioeconomic conditions — a prefectural leader may opt for alternative policies from jurisdictions that are less socioeconomically similar, or adopt central government policies that are less tailored to local needs.

To evaluate how centralization in policymaking may mitigate the distortions embedded in decentralized policy diffusion, we proceed in two steps. First, we quantify the extent to which political competition affects the suitability of local policies to local condi-

28. Furthermore, as shown in Appendix Table A.11, the results remain robust when we restrict the sample to policies labeled as "economic policies" by *PKULaw*, or when we focus exclusively on the agricultural and manufacturing sectors.

29. Similar patterns appear in Appendix Figure A.16. In an event study at the prefecture–industry–year level, we observe that once the central government adopts a locally initiated industrial policy and promotes it as national, the suitability of subsequently adopting localities drops sharply.

30. More broadly, it has been well documented that decentralized competition among Chinese local bureaucrats leads to various distortions (Jia and Nie 2017; Wang and Wang 2020; Wang and Yang 2025).

tions. Second, we examine whether and to what extent these effects are mitigated by the post-2013 shift toward centralization.

Political competition and policy suitability We rank each prefecture's 30 closest economic peers — those with the most similar GDP per capita — and count how many of them are led by the prefecture leader's 30 closest political competitors — those with the most similar backgrounds. We then test whether greater political competition among economically comparable prefectures — which should limit the menu of viable policies — actually leads to less well-suited policy portfolios.

Table 4, Panel A shows the results. We find that having one additional political competitor among economic neighbors is associated with a 0.74% loss in average policy-locality suitability in that year, confirming the hypothesis that strategic competition prevents the diffusion of beneficial policies. In Panels B and C, we expand the definitions of economic and political neighbors to the top 40 and top 50, respectively, the qualitative findings remain consistent, and as expected, the estimated effects are attenuated. Reassuringly, in Appendix Table A.13, Panel A, we verify that the level of political competition with economic neighbors is uncorrelated with the overall number of policies adopted in a given year, confirming that the extensive margin remains unaffected.

This finding, together with the observed link between lower policy suitability and reduced policy effectiveness, underscores the economic costs of decentralized policymaking and policy diffusion. Next, we examine the extent to which this bias has been mitigated by the centralization in policymaking since 2013.

Centralization mitigating decentralized distortions Figure 2, Panels C and D suggest that centralized policymaking indeed changed policy diffusion patterns. After 2013, the positive correlation between economic similarity and policy diffusion weakens, consistent with centralization leading to the adoption of policies less tailored to local conditions. Interestingly, Panel D shows that the negative association between political proximity and policy diffusion also significantly weakens after 2013, confirming that centralization helps mitigate biases in decentralized diffusion.

The interaction terms in Table 4 quantify how the changes in policy diffusion patterns after 2013 affect policy suitability: the negative correlation between average policy-locality suitability and political competition with economic neighbors is completely muted post-2013. This suggests that centralization — shifting incentives from bottom-up policy initiatives toward top-down compliance — mitigates the distortions inherent in decentralized policy diffusion.

To the extent that strategic distortion in decentralized policy diffusion is driven by political competition among local bureaucrats for career advancement, the post-2013 removal of political rewards for bottom-up policy innovation (documented in Section 4.2) has inadvertently tamed these distortions. Under more centralized policymaking after 2013, local bureaucrats have become less concerned about boosting their peers' credentials and, consequently, more willing to adopt policies proven effective in peer localities with similar socioeconomic conditions.[31]

6.3 Net effect of centralization: back-of-the-envelope calculation

As shown in Sections 6.1 and 6.2, the centralization of policymaking has mixed impacts on policy-locality suitability. By linking these countervailing forces to the relationship between suitability and policy effectiveness documented in Section 5.3, we perform a back-of-the-envelope calculation to quantitatively compare trade-offs associated with centralization. Below, we outline main components of the calculation; we provide more details in Appendix D.

Compared to the 2012 level of decentralization, the post-2013 centralization trend converted 2,562 prefecture-level industrial policies from bottom-up to top-down. Based on our estimates in Section 6.1, each top-down policy is 22% less suitable for local conditions than a bottom-up one. Linking this suitability gap to the differential effectiveness of suitable versus unsuitable policies in driving industrial growth and innovation, we calculate that the *yearly* cost of lowered policy suitability that can be attributed to the post-2013 centralization in policymaking is 580 billion RMB in industrial output, 437 billion in exports and 10,486 patent filings.

Meanwhile, Section 6.2 shows that for each of the top 30 political rivals within a prefecture's top 30 economic neighbors, post-2013 centralization mitigates competition-induced suitability loss by 0.007. Given that the average prefecture has 2.78 such rivals, this amounts to a 0.02 increase in policy suitability. Applying this figure to the 9,376 bottom-up industrial policies implemented annually — and using our differential effectiveness estimates — we calculate yearly benefits of 121 billion RMB in industrial output, 91 billion RMB in exports, and 2,194 additional patent filings.

These calculations, which use the more conservative estimates based on the investment-suitability measure, suggest that — through the lens of policy suitability — the costs of centralization consistently exceed its benefits by more than 400%, cautioning against over-

31. Though we note an overall decrease in horizontal diffusion of locally initiated policies, suggesting a reduction in local bureaucrats' policymaking motivation.

centralization in policymaking, at least in this context.[32]

6.4 Other trade-offs associated with centralization

So far, we have evaluated policy centralization primarily through the lens of policy–locality suitability, a dimension emphasized in the theoretical literature. While our analysis shows that centralization can have important consequences for policy outcomes, it should not be taken as a comprehensive assessment of all the trade-offs it entails. In this section, we therefore provide suggestive evidence on the empirical relevance of several additional trade-offs associated with centralization, with the aim of informing future research on the topic.

First, we examine whether the central government's industrial policies tend to target more "ambitious" industries. Using ComTrade data, we compute the 2024 global market size (total import value summed across all countries) for each industry and map this information to China's industrial policies. As shown in Appendix Table A.14, Panel A, top-down industrial policies do not appear to select significantly larger industries in the long run, compared to their bottom-up counterparts.[33]

Second, we investigate whether top-down policies are more focused on industries pertinent to national security. Specifically, we identify industries included in the export ban list to China by the U.S. government.[34] We observe that the Chinese central government is indeed more likely to promote such strategic industries than local governments (Panel B of Appendix Table A.14).[35]

Third, the central government may have a longer policy horizon than local governments, causing it to favor more forward-looking industries — those in which China currently lacks a comparative advantage, but have high potentials in the long run. To test this, we use UN Comtrade data to calculate China's revealed comparative advantage

32. By contrast, calculations based on the supply-chain-suitability measure imply even larger discrepancies between the costs and benefits of centralization.

33. Given that local governments initiate significantly more industrial policies than the central government, in Panel A, we further restrict the bottom-up sample to local policies targeting industries with the highest future output values, ensuring that the number of central and local policies is balanced. Under such comparison, local governments actually target *more* ambitious industries.

34. The sanction list is sourced from the "Critical Supply Chains Products" released by the U.S. International Trade Administration (available at: https://www.trade.gov/data-visualization/draft-list-critical-supply-chains). This list comprises 2,409 HS6 products, predominantly spanning four critical domains: Public Health, Energy, Information Technology/Communications Technology (IT/CT), and Critical Minerals and Materials.

35. That said, as shown in Appendix Table A.15, even after accounting for these national security-related industries, our main findings regarding the reduced suitability for top-down policies persist (with even larger magnitudes), suggesting that the costs of centralization cannot be attributed to national security considerations.

(RCA) in exporting, in both 2000 and 2024.[36] We then calculate, for each industry, the difference in these two RCA measures, which serves as a proxy for China's long-term growth in it. As shown in Panel C of Appendix Table A.14, the central government does not appear to target industries with higher long-run growth, as compared to their local counterparts.[37]

Fourth, as pointed out by Liu (2019), within a production network, market distortions accumulate via backward demand linkages, causing upstream sectors to suffer the greatest size distortions. In such settings, there may be a rationale for the central government to promote these upstream sectors to maximize national output. As shown in Panel D of Appendix Table A.14, we find no evidence that the central government targets such industries more than local governments do.

Fifth, the central government might additionally promote sectors with high economies of scale to maximize national aggregate productivity — an objective that may lie beyond the scope of individual local governments. Using the sector-level economies of scale estimated by Atkin, Costinot, and Fukui (2021), we find no evidence that the central government is more likely to target sectors with higher economies of scale (Panel E of Appendix Table A.14).

Finally, certain industries generate pollution externalities that affect neighboring jurisdictions. To better coordinate spatial environmental spillovers, it might be more efficient for the central government, rather than decentralized local governments, to take the lead in promoting such industries. Following He, Wang, and Zhang (2020), we identify key manufacturing industries classified as polluting by the Ministry of Ecology and Environment. As reported in Panel F of Appendix Table A.14, the central government is indeed more likely to promote these sectors than local governments.[38]

The above patterns suggest that centralization in policymaking may entail a series of high-stakes trade-offs beyond policy–locality suitability, underscoring the need for more holistic assessments of centralized policymaking in future work.

36. Export-based RCA is defined as: $RCA_{cp} = \frac{E_{cp} / \sum_{p' \in P} E_{cp'}}{\sum_{c' \in C} E_{c'p} / \sum_{c' \in C, p' \in P} E_{c'p'}}$. The numerator captures, for a given country c, the export in industry p relative to other industries. The denominator measures this share for all other countries in the world.

37. Furthermore, Panel B reveals that when the number of central and local policies is held equal (i.e., stratifying local policies by highest long-run growth), industries with the highest long-run potential appear to be promoted more by the local governments instead of the central government. These findings suggest that central authorities are not inherently more forward-looking than local governments when selecting industrial policies.

38. Appendix Table A.15 shows that accounting for such polluting sectors does not alter our baseline finding that top-down policies have lower suitability.

7 Conclusion

In this paper, we map the landscape of China's policymaking process. By tracing the full arc of over 115,000 distinct policies — from inception to diffusion and adoption — we uncover salient features of an institutional setup once characterized by decentralized experimentation. We then document the political incentive changes that fueled the shift to an increasingly centralized policymaking regime. We offer an empirical investigation into the trade-offs associated with centralized policymaking, highlighting a tension between the suitability of locally-tailored policies and the distortions arising from strategic competition in decentralized policy diffusion.

While scholars have long speculated about a turn toward centralization in China in recent years, this shift has advanced in subtler ways than conventional aggregate indicators reveal. The central government is not necessarily issuing more policies relative to local governments, perhaps due to bureaucratic capacity constraints (a fascinating topic for future inquiry). Rather, centralization becomes saliently evident only when one traces the full policymaking cycle and links them to bureaucrats' tenures, which highlights the evolving role of local officials in driving or constraining policy initiatives. This underscores the value of large-scale textual analysis for uncovering economic and political dynamics that might otherwise remain hidden.

The trade-offs studied in this paper are not unique to China. As governments around the world grapple with challenges that demand both coordination and customization — from climate mitigation, to education policy, to industrial strategy — understanding the optimal hierarchical level for decision-making and the associated trade-offs becomes increasingly imperative. By illuminating the mechanisms and consequences of centralization in China's policymaking, this paper provides new evidence and calls for a reconsideration of how polities of varying sizes can design institutions that balance local initiative with system-wide integration.

While we find that political incentives explain both the quantity of policy innovations and the quality of policy diffusion, we offer little empirical insight into the quality of policies that could have been invented but were "strangled in the cradle," since we observe only those that ultimately appeared on paper. Moreover, our policy–locality suitability measure is inherently relative, allowing us to study mismatches between policies and localities within a given industry but ill-suited for assessing the absolute quality of the overall pool of industrial policies. Investigating how institutional changes shape both the direction and quality of policy innovation remains an exciting avenue for future research.

References

Abowd, John M, Francis Kramarz, and David N Margolis. 1999. "High wage workers and high wage firms." *Econometrica* 67 (2): 251–333.

Aho, Alfred V, and Margaret J Corasick. 1975. "Efficient string matching: an aid to bibliographic search." *Communications of the ACM* 18 (6): 333–340.

Alesina, Alberto, Reza Baqir, and Caroline Hoxby. 2004. "Political Jurisdictions in Heterogeneous Communities." *Journal of Political Economy* 112 (2): 348–396.

Alesina, Alberto, and Enrico Spolaore. 1997. "On the Number and Size of Nations." *Quarterly Journal of Economics* 112 (4): 1027–1056.

Andrews, Martyn J, Len Gill, Thorsten Schank, and Richard Upward. 2008. "High wage workers and low wage firms: negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society Series A: Statistics in Society* 171 (3): 673–697.

Atkin, David, Arnaud Costinot, and Masao Fukui. 2021. *Globalization and the Ladder of Development: Pushed to the Top or Held at the Bottom?* Technical report. National Bureau of Economic Research.

Balán, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L Weigel. 2022. "Local elites as state capacity: How city chiefs use local information to increase tax compliance in the democratic republic of the Congo." *American Economic Review* 112 (3): 762–797.

Bardhan, Pranab. 2002. "Decentralization of governance and development." *Journal of Economic Perspectives* 16 (4): 185–205.

Bernecker, Andreas, Pierre C Boyer, and Christina Gathmann. 2021. "The role of electoral incentives for policy innovation: Evidence from the us welfare reform." *American Economic Journal: Economic Policy* 13 (2): 26–57.

Besley, Timothy, and Anne Case. 2003. "Political institutions and policy choices: evidence from the United States." *Journal of Economic Literature* 41 (1): 7–73.

Besley, Timothy, and Stephen Coate. 2003. "Centralized versus Decentralized Provision of Local Public Goods: A Political Economy Approach." *Journal of Public Economics* 87 (12): 2611–2637.

Blanchard, Olivier, and Andrei Shleifer. 2001. "Federalism with and without political centralization: China versus Russia." *IMF staff papers* 48 (Suppl 1): 171–179.

Bolton, Patrick, and Gerard Roland. 1997. "The Breakup of Nations: A Political Economy Analysis." *Quarterly Journal of Economics* 112 (4): 1057–1090.

Burgess, Robin, Matthew Hansen, Benjamin A. Olken, Peter Potapov, and Stefanie Sieber. 2012. "The Political Economy of Deforestation in the Tropics." *Quarterly Journal of Economics* 127 (4): 1707–1754. https://doi.org/10.1093/qje/qjs034.

Cai, Hongbin, and Daniel Treisman. 2004. "State corroding federalism." *Journal of Public Economics* 88 (3-4): 819–843.

Cao, Yuanzheng, Yingyi Qian, and Barry R Weingast. 1999. "From federalism, Chinese style to privatization, Chinese style." *Economics of Transition* 7 (1): 103–131.

Caughey, Devin, Yiqing Xu, and Christopher Warshaw. 2017. "Incremental democracy: The policy effects of partisan control of state government." *The Journal of Politics* 79 (4): 1342–1358.

Chen, Heng, Bingjing Li, and Xiaodong Zhu. 2024. "Bottom-Up Institutional Change and Growth in China."

Cremer, Jacques, Antonio Estache, and Paul Seabright. 1994. *The decentralization of public services: lessons from the theory of the firm.* Vol. 1345. World Bank Publications.

Dal Bó, Ernesto, Frederico Finan, Nicholas Y Li, and Laura Schechter. 2021. "Information technology and government decentralization: Experimental evidence from paraguay." *Econometrica* 89 (2): 677–701.

DellaVigna, Stefano, and Woojin Kim. 2022. *Policy diffusion and polarization across US states.* Technical report. National Bureau of Economic Research.

Fang, Hanming, Ming Li, and Guangli Lu. 2025. *Decoding China's Industrial Policies.* Technical report. National Bureau of Economic Research.

Frank, Eyal G, Shaoda Wang, Xuebin Wang, Qinyun Wang, and Yang You. 2024. *Campaigning for Extinction: Eradication of Sparrows and the Great Famine in China.* Technical report. Working Paper.

Fu, Xiaolan. 2015. *China's path to innovation.* Cambridge University Press.

Glaser, Bonnie S. 2022. *Made in China 2025 and the Future of American Industry.* Center for Strategic and International Studies (CSIS).

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. "Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings." *Journal of political economy* 118 (3): 536–598.

Grumbach, Jacob M. 2023. "Laboratories of democratic backsliding." *American Political Science Review* 117 (3): 967–984.

Hayek, Friedrich A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35 (4): 519–530.

He, Guojun, Shaoda Wang, and Bing Zhang. 2020. "Watering down environmental regulation in China." *The quarterly journal of economics* 135 (4): 2135–2185.

Heilmann, Sebastian. 2008. "From local experiments to national policy: the origins of China's distinctive policy process." *The China Journal*, no. 59, 1–30.

———. 2018. *Red swan: how unorthodox policy-making facilitated China's rise.* The Chinese University of Hong Kong Press.

Jia, Ruixue, and Huihua Nie. 2017. "Decentralization, collusion, and coal mine deaths." *Review of Economics and Statistics* 99 (1): 105–118.

Juhász, Réka, Nathan Lane, Emily Oehlsen, and Verónica C Pérez. 2022. "The who, what, when, and how of industrial policy: A text-based approach." *What, When, and How of Industrial Policy: A Text-Based Approach (November 20, 2022).*

Kline, Patrick, and Enrico Moretti. 2014. "Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority." *The Quarterly journal of economics* 129 (1): 275–331.

Kornai, Janos. 1960. "Overcentralization in Economic Administration: A Critical Analysis Based on Experience in Hungarian Light Industry." *Soviet Studies* 11, no. 4 (April): 421–424.

Li, Hongbin, and Li-An Zhou. 2005. "Political turnover and economic performance: the incentive role of personnel control in China." *Journal of public economics* 89 (9-10): 1743–1762.

Li, Shanjun. 2018. "Better lucky than rich? Welfare analysis of automobile licence allocations in Beijing and Shanghai." *The Review of Economic Studies* 85 (4): 2389–2428.

Lin, Justin Yifu. 1990. "Collectivization and China's agricultural crisis in 1959-1961." *Journal of Political Economy* 98 (6): 1228–1252.

Lipscomb, Molly, and Ahmed Mushfiq Mobarak. 2017. "Decentralization and Pollution Spillovers: Evidence from the Re-drawing of County Borders in Brazil." *Review of Economic Studies* 84 (1): 464–502. https://doi.org/10.1093/restud/rdw023.

Liu, Ernest. 2019. "Industrial policies in production networks." *The Quarterly Journal of Economics* 134 (4): 1883–1948.

Lorentzen, Peter L, and Xi Lu. 2021. "Personal ties, meritocracy, and China's anti-corruption campaign."

Martinez-Bravo, Monica, Gerard Padró i Miquel, Nancy Qian, and Yang Yao. 2022. "The rise and fall of local elections in China." *American Economic Review* 112 (9): 2921–2958.

Meng, Xin, Nancy Qian, and Pierre Yared. 2015. "The institutional causes of China's great famine, 1959–1961." *The Review of Economic Studies* 82 (4): 1568–1611.

Montinola, Gabriella, Yingyi Qian, and Barry R Weingast. 1995. "Federalism, Chinese style: the political basis for economic success in China." *World politics* 48 (1): 50–81.

Mookherjee, Dilip. 2015. "Political decentralization." *Annual review of economics* 7 (1): 231–249.

Naughton, Barry. 2021. *The rise of China's industrial policy, 1978 to 2020.* Universidad Nacional Autónomica de México, Facultad de Economía México.

New State Ice Co. v. Liebmann. 1932. *New State Ice Co. v. Liebmann.* 285 U.S. 262. Opinion by Justice Brandeis, decided May 23, 1932.

Nove, Alec. 1971. "An Economic History of the U.S.S.R." *The Journal of Economic History* 31 (2): 514–515. https://doi.org/10.1017/S002205070009135X.

O'Brien, Kevin J, and Lianjiang Li. 2017. "Selective policy implementation in rural China." In *Critical readings on the communist party of China (4 Vols. Set)*, 437–460. Brill.

Oates, Wallace E. 1972. *Fiscal Federalism.* New York: Harcourt Brace Jovanovich.

Oi, Jean C. 1995. "The role of the local state in China's transitional economy." *The China Quarterly* 144:1132–1149.

Olken, Benjamin A. 2007. "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 115 (2): 200–249. https://doi.org/10.1086/517935.

Qian, Yingyi. 2002. "How reform worked in China."

Rawski, Thomas G. 1995. "Implications of China's reform experience." *China Q.*, 1150.

Roland, Gerard. 2000. *Transition and economics: Politics, markets, and firms.* MIT press.

Rueschemeyer, Dietrich, Theda Skocpol, and Peter B Evans. 1985. *Bringing the state back in.* Cambridge University Press Cambridge.

Seabright, Paul. 1996. "Accountability and decentralisation in government: An incomplete contracts model." *European economic review* 40 (1): 61–89.

Sonin, Konstantin. 2003. "Why the rich may favor poor protection of property rights." *Journal of comparative economics* 31 (4): 715–731.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of econometrics* 225 (2): 175–199.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditure." *Journal of Political Economy* 64 (5): 416–424.

Wang, Shaoda, and Zenan Wang. 2020. "The environmental and economic consequences of internalizing border spillovers." *University of Chicago: Chicago, IL, USA.*

Wang, Shaoda, and David Yang. 2025. "Policy experimentation in China: The political economy of policy learning." *Journal of Political Economy.*

Wang, Shenhao, and Jinhua Zhao. 2017. "The distributional effects of lotteries and auctions—License plate regulations in Guangzhou." *Transportation Research Part A: Policy and Practice* 106:473–483.

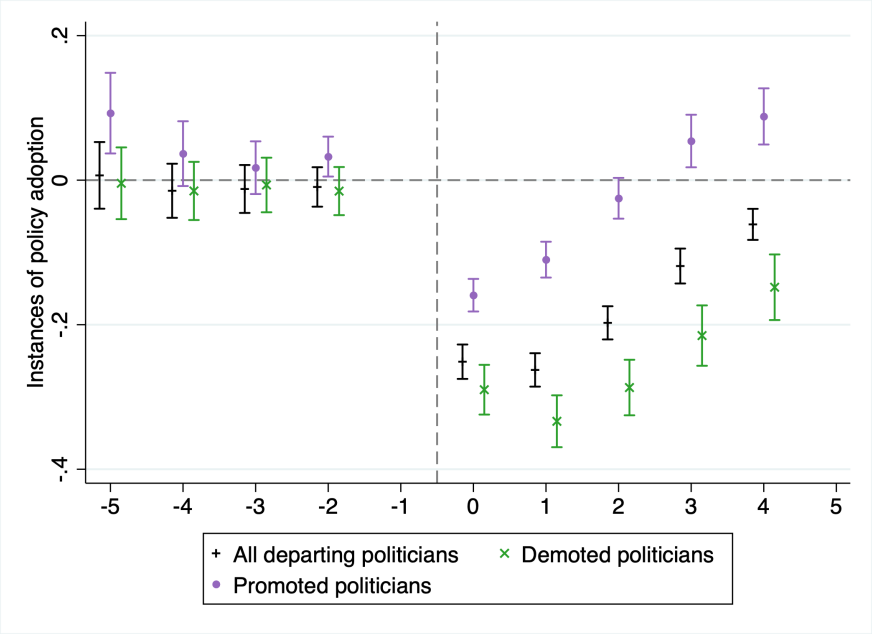
Wang, Zhi, Qinghua Zhang, and Li-An Zhou. 2020. "Career incentives of city leaders and urban spatial expansion in China." *Review of Economics and Statistics* 102 (5): 897–911.

Xu, Chenggang. 2011. "The fundamental institutions of China's reforms and development." *Journal of economic literature* 49 (4): 1076–1151.

Young, Alwyn. 2000. "The razor's edge: Distortions and incremental reform in the People's Republic of China." *The Quarterly Journal of Economics* 115 (4): 1091–1135.

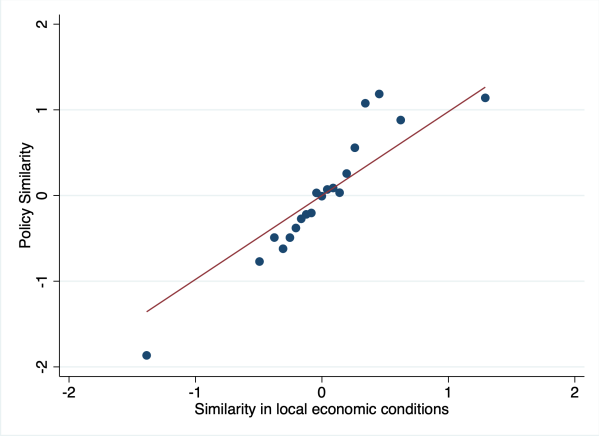
Figures

Figure 1: Bureaucratic turnover and policy diffusion

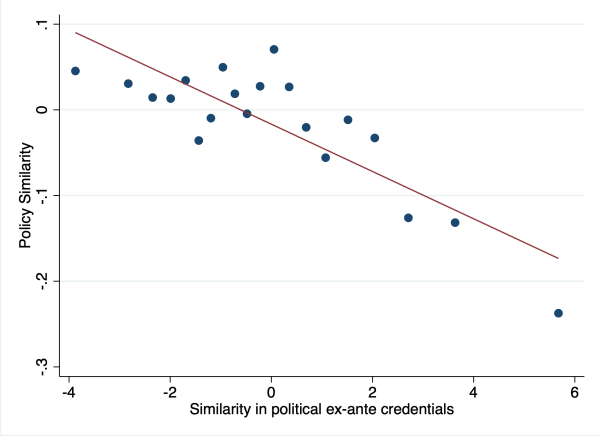


Notes: This figure presents event-study estimates illustrating the decrease in policy adoptions after prefectural party-secretary departure. Specifically, we run a policy-year level regression with two-way fixed effects, focusing on instances of adoption ± 5 years around the departure of the politician who initiated the policy. We cluster standard errors at the policy level and compare baseline estimates with cases where the departing politicians got promoted or demoted.

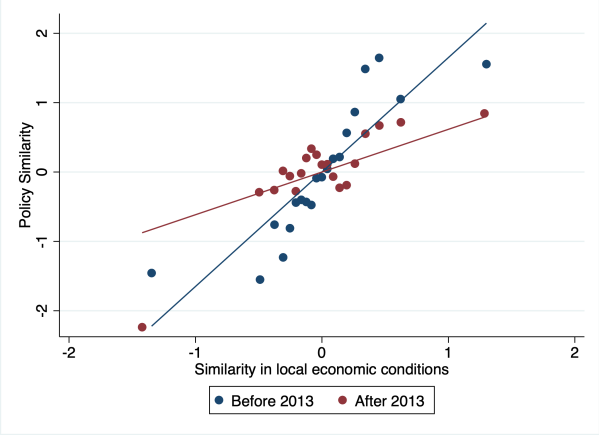
Figure 2: Policy diffusion based on economic vs. political similarity



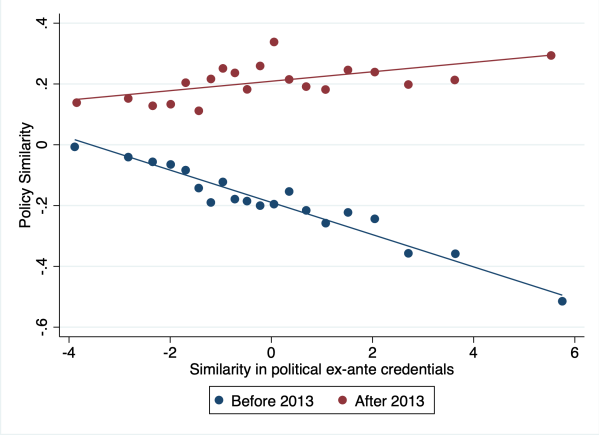
Panel A: Policy portfolio and economic proximity



Panel B: Policy portfolio and political proximity



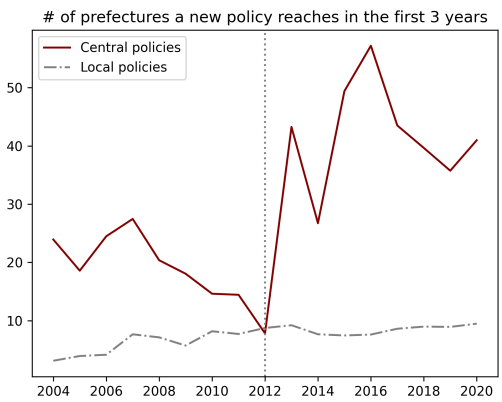
Panel C: Policy portfolio and economic proximity: Before and After 2013



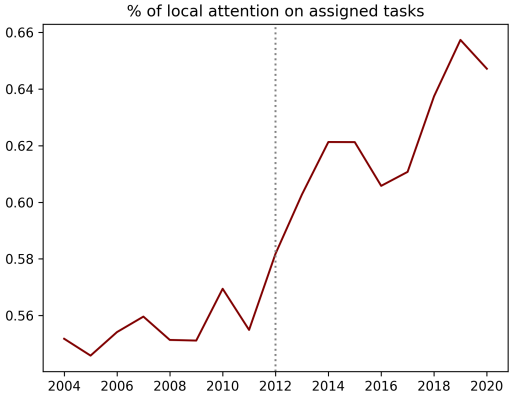
Panel D: Policy portfolio and political proximity: Before and After 2013.

Notes: This figure illustrates the similarity of policy space between city pairs as a function of their economic and political proximity. The dependent variable, policy similarity, is measured as the Euclidean distance between two vectors of dummy variables, v_{pct}, indicating whether policy p was implemented in city c in year t. The key independent variable, economic similarity (Panel A), is defined as the standardized difference in GDP per capita. Political credential similarity (Panel B), on the other hand, is defined as the Mahalanobis distance between the full set of observable characteristics of local politicians prior to their appointment, including age at entry, hierarchical rank, gender, ethnicity, education, central government experience, and prior local experience. Panel A plots policy similarity against the economic proximity, controlling for origin, destination, and year fixed effects. Panel B plots policy similarity against political proximity, controlling for prefecture pair fixed effects and year fixed effects. Panels C and D contrast the relationship before and after 2013.

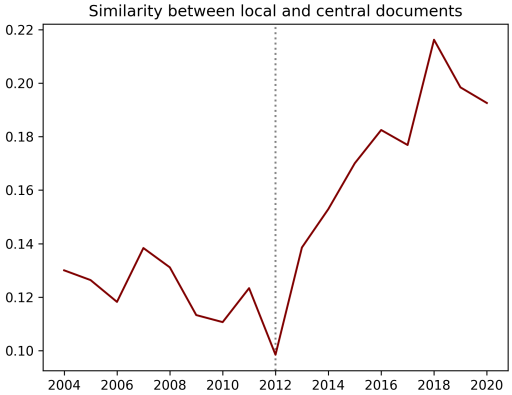
Figure 3: Centralization trend in policymaking



Panel A



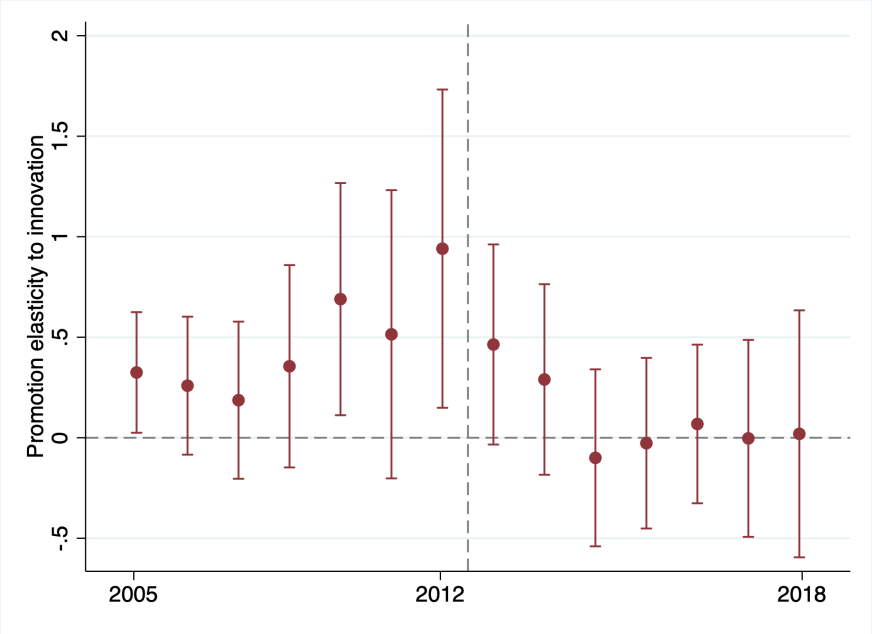
Panel B



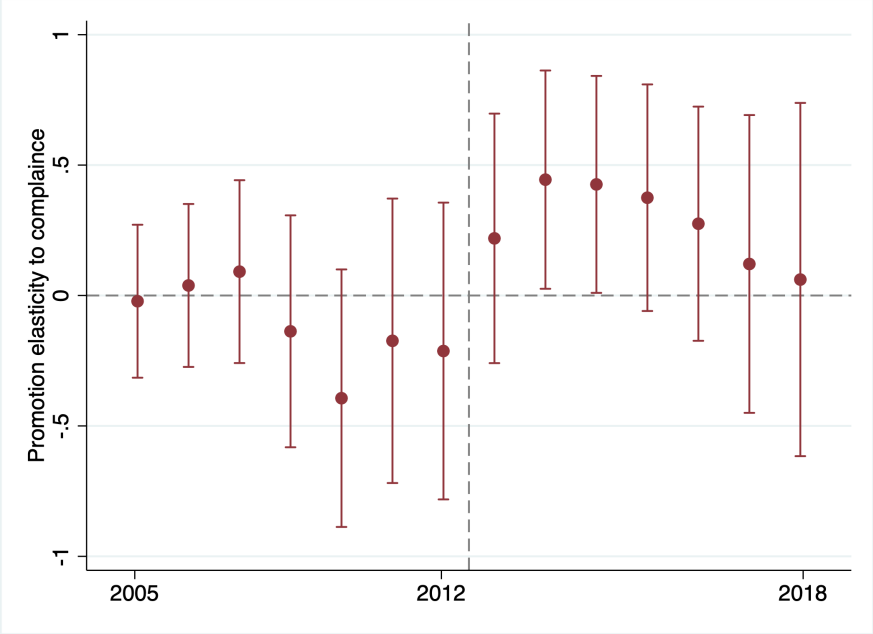
Panel C

Notes: This figure illustrates the increasing centralization of policymaking after 2013. In Panel A, we focus on central-government initiated policies, and plot the number of prefectures adopting the policy in the first 3 years by year of policy-initiation. Panel B shows the annual trend in local governments' attention to centrally assigned tasks. For each locality-year, we calculate the share of implemented policies that had been promoted by the central government beforehand, and then average this measure across all localities within the same year. In panel C, we compute the average textual similarity between the first central document on a certain topic and local follow-ups with TF-IDF algorithm, cosine similarity, and standard stop-word removal.

Figure 4: Career incentives and policy innovation vs. compliance



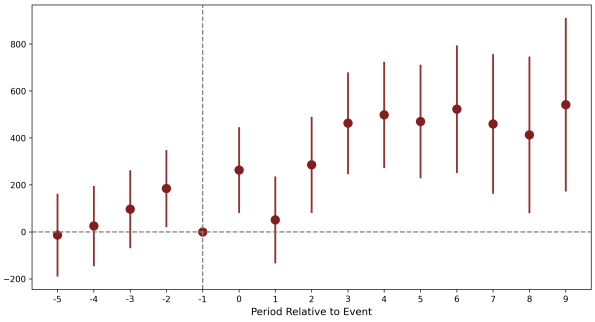
Panel A: Promotion elasticity wrt. innovation



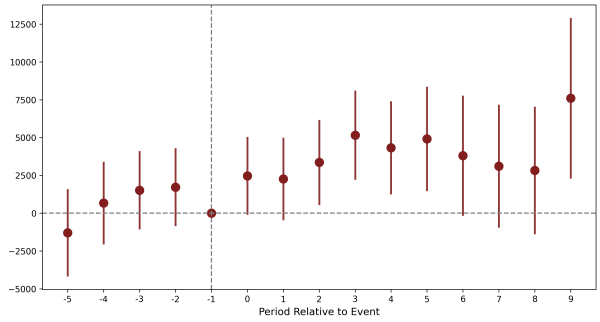
Panel B: Promotion elasticity wrt. compliance

Notes: The two figures above plot the point estimates and 90% confidence intervals of a series of cross-sectional regressions. In each regression, we focus on a sliding-window of politicians who depart from office during $[t-2, t+2]$. We regress their job outcome (promotion$\in \{0,1\}$) on both an index for innovation and another for compliance. We parse out the effect of GDP growth on promotion, and control for both tenure-start-year $\times$ hierarchy and tenure-end-year $\times$ hierarchy fixed effects. The dependent variables in both figures are scaled by its mean respectively so that the estimated coefficients can be interpreted as the elasticity of promotion probability with respect to changes in policy innovation (compliance). Standard errors are clustered at prefecture level.

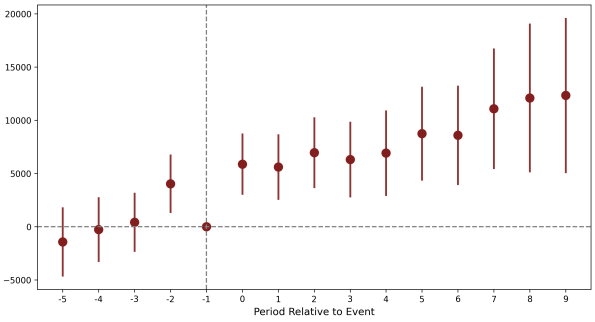
Figure 5: Local suitability and industrial policy effectiveness



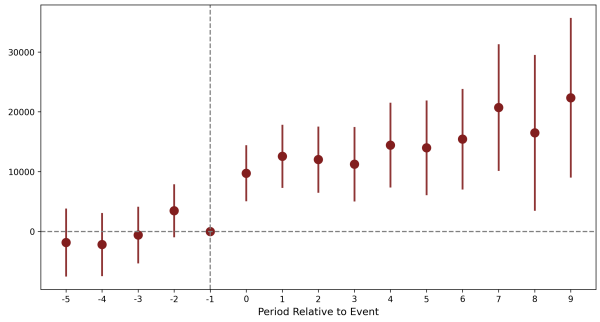
Panel A.1: Export value (1 thousand yuan), Supply-chain suitability



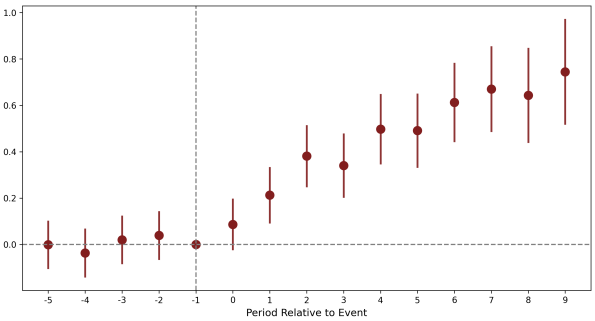
Panel B.1: Export value (1 thousand yuan), Investment suitability



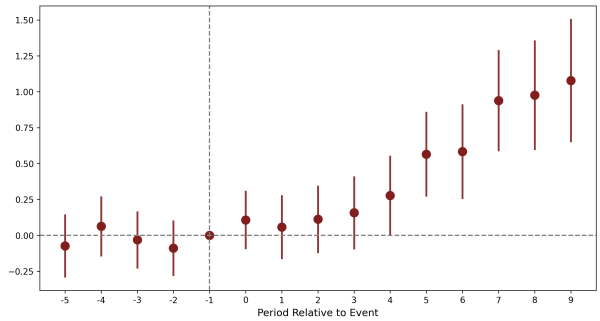
Panel A.2: Total revenue (1 thousand yuan), Supply-chain suitability



Panel B.2: Total revenue (1 thousand yuan), Investment suitability



Panel A.3: Patent (counts), Supply-chain suitability



Panel B.3: Patent (counts), Investment suitability

Notes: This figure illustrates how policy effectiveness varies with *ex ante* local suitability, based on a triple-difference strategy. We define the treatment as the implementation of any industrial policy in prefecture p during year t targeting industry i. We include dummy variables for 5 lead periods and 10 lag periods relative to that treatment, as well as their interactions with policy-locality suitability (supply-chain or investment). We report the estimated coefficients and standard errors before the interaction term. Panels A.1–A.3 present results based on *ex ante* supply-chain suitability, Panels B.1–B.3 focus on *ex ante* investment suitability. Prefecture-by-year, prefecture-by-industry, and industry-by-year fixed effects are controlled in each regression. Standard errors are clustered at the prefecture level.

Table 1: Summary statistics of policymaking

	# of policies	% top-down	% nationalized	Avg. # adoption
	(1)	(2)	(3)	(4)
Full sample	115679	0.181	0.392	35.104
by issuing ministry:				
Agriculture	9515	0.244	0.532	42.144
Business, finance, and economics	9279	0.252	0.518	46.124
Development and reform	17880	0.196	0.466	44.747
Domestic affairs and public security	3407	0.222	0.485	40.925
Environment	6094	0.229	0.518	51.958
Judiciary supervision	1453	0.275	0.560	43.109
Science, education and culture	10691	0.279	0.566	48.995
Transportation	1429	0.191	0.427	36.703
Others	55931	0.126	0.265	23.902
Alternative policy definitions				
Aggregating similar keywords	101966	0.202	0.438	38.248
Exploding by policy×domain	651488	0.206	0.355	9.375
Extracted from policy titles	276479	0.356	0.688	61.313

Notes: In this table we report some summary statistics – that include the number of policies within each category, the percentage of them initiated by the central government (with a central document preceding all local follow-ups), the percentage of policies eventually rolled out to the entire country as of 2024, and the average number of provinces implementing the policies. In Panel A, we report summary statistics for the full sample, that in each policy domain, and that in each alternative policy definitions. In Panel B, we focus on the industrial policy sample, defining an industrial policy keyword as one for which, among all documents containing the keyword, more than 50% also mention a specific four-digit industry name. We then report those statistics in each industry categories, and finally counting the instances of policy adoption.

Table 2: Decomposition of innovation between bureaucrats and locality

Decomposing innovation and compliance			
	$\tau_{\text{politician}}$	$\tau_{\text{prefecture}}$	τ_{year}
	(1)	(2)	(3)
Panel A: Bottom-up innovation index			
Variation of Y explained	0.304***	0.059*	0.132***
	(0.037)	(0.033)	(0.014)
Panel B: Top-down compliance index			
Variation of Y explained	0.196***	0.088**	0.308***
	(0.046)	(0.043)	(0.016)

Notes: In Panel A, we follow Abowd, Kramarz, and Margolis (1999) to decompose the bottom-up innovation index into bureaucrat fixed effects, locality fixed effects, and calendar year fixed effects. In Panel B, we repeat the same exercise for the top-down compliance index. Akin to employer-employee matched design, identification exploits variation of innovation index within (rotating) bureaucrats across the places they hold office. After the decomposition, we report the share of variance of Y associated with each set of fixed effects ($\text{cov}(\tau, Y)/\text{var}(Y)$). In the parentheses below the estimates, we report standard errors estimated with 1,000 bootstraps.

Table 3: Centralization and policy suitability

	Investment suitability		Supply-chain suitability	
	Continuous	Top 10%	Continuous	Top 10%
	(1)	(2)	(3)	(4)
Panel A: All policies				
Top-down	-0.274***	-0.031***	-0.201***	-0.027***
	(0.069)	(0.008)	(0.041)	(0.008)
# of obs.	118,104	118,104	114,484	114,484
Panel B: Policies with subsidy				
Top-down	-0.275***	-0.036***	-0.193***	-0.028***
	(0.081)	(0.009)	(0.044)	(0.010)
# of obs.	41,415	41,415	40,457	40,457
Mean of DV	1.23	0.23	1.09	0.19
SD of DV	3.34	0.17	1.51	0.39
Prefecture $\times$ Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the difference in policy suitability between top-down and bottom-up local industrial policies. Specifically, we regress policy suitability, based either on investment (columns 1–2) or supply-chain linkages (columns 3–4), on an indicator for whether an industrial policy has been promoted by the central government in the past 3 years before its local implementation. Columns 1 and 3 report estimates using the continuous suitability measure, whereas columns 2 and 4 use an indicator variable that equals one if suitability falls within the top 10 percent of its distribution. In Panel A, we identify the adoption of industrial policies using the full policy-document sample. In Panel B, we focus on a subset of policy documents where keywords or amount of subsidies are explicitly mentioned, concerning that some policy adoptions might be performative without substantiate fiscal input. All regressions include prefecture-by-year fixed effects. Standard errors are clustered at the prefecture level.

Table 4: Political competition among economic neighbors and policy suitability

	Investment suitability		Supply-chain suitability	
	Continuous	Top 10%	Continuous	Top 10%
	(1)	(2)	(3)	(4)
Panel A: among top 30 economic neighbors				
# of 30 top political competitors	-0.0074***	-0.0024**	-0.0076	-0.0053***
	(0.0022)	(0.0011)	(0.0070)	(0.0020)
# $\times \mathbb{1}$\{post 2013\}	0.0077***	0.0014	0.0039	0.0059**
	(0.0027)	(0.0013)	(0.0082)	(0.0024)
Panel B: among top 40 economic neighbors				
# of 40 top political competitors	-0.0057***	-0.0018**	-0.0095*	-0.0029**
	(0.0016)	(0.0007)	(0.0049)	(0.0014)
# $\times \mathbb{1}$\{post 2013\}	0.0083***	0.0022**	0.0076	0.0036**
	(0.0019)	(0.0010)	(0.0056)	(0.0017)
Panel C: among top 50 economic neighbors				
# of 50 top political competitors	-0.0041***	-0.0012**	-0.0046	-0.0024**
	(0.0012)	(0.0006)	(0.0038)	(0.0011)
# $\times \mathbb{1}$\{post 2013\}	0.0062***	0.0016**	0.0046	0.0027**
	(0.0014)	(0.0007)	(0.0041)	(0.0064)
# of obs.	3,824	3,824	3,731	3,731
Mean of DV	1.23	0.23	1.09	0.19
Prefecture FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: This table presents the effects of political competition among a local government's economic neighbors on the suitability of the policies it adopts. Political similarity is measured as the standardized Mahalanobis distance between the portfolios of two politicians. We rank each prefecture's 30, 40, 50 closest economic peers — those with the most similar GDP per capita, and count how many of them are led by the prefecture leader's 30, 40, 50 closest political competitors — those with the most similar backgrounds. Column (1) reports results for continuous investment suitability. Column (2) reports results for a binary variable indicating whether investment suitability is among the top 10% nationwide. Columns (3) and (4) repeat the same exercise using supply-chain suitability. All specifications include prefecture and year fixed effects. Standard errors are clustered at the prefectural level.

Online Appendix

Appendix A Additional figures and tables

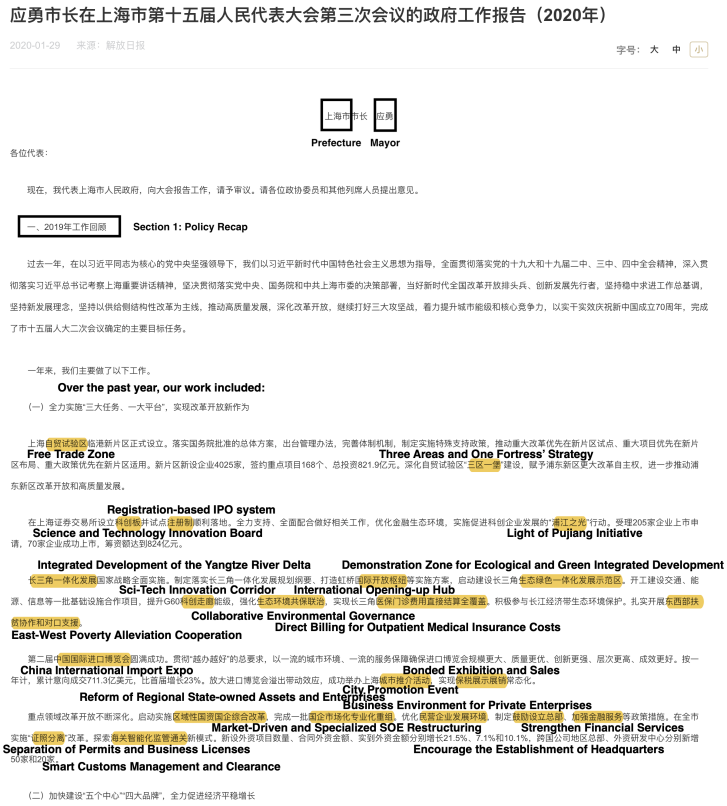
Figure A.1: Sample of policy documents



Panel A: Prefectural policy document



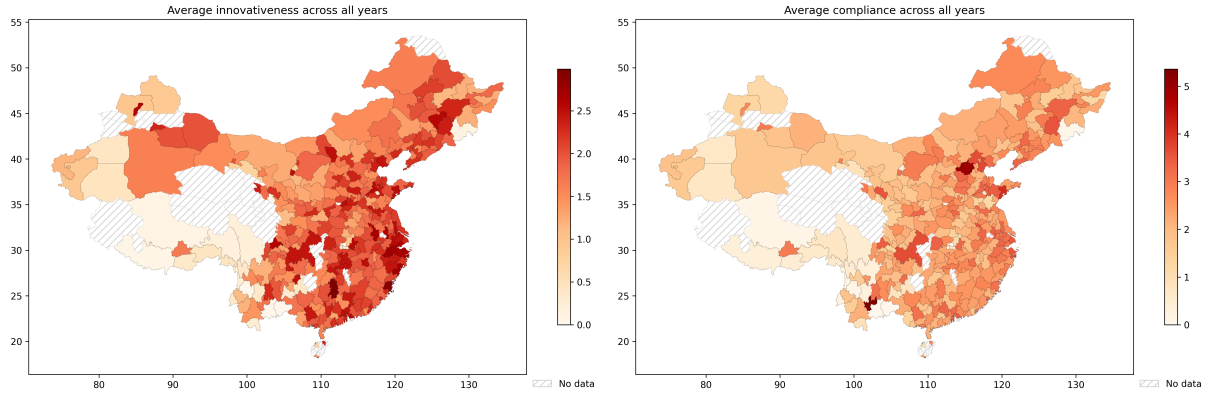
Panel B: Central policy document



Panel C: Example of government report (Shanghai, 2020)

Notes: Panel A shows a prefectural policy document from PKULaw, highlighting the policy keyword "developing a water-saving society." Panel B presents an example of a central government document from PKULaw containing the same policy keyword. Panel C shows a screenshot of the Shanghai government work report (2020), illustrating the policy-keyword extraction procedure.

Figure A.2: Spatial variation in policy innovation and compliance

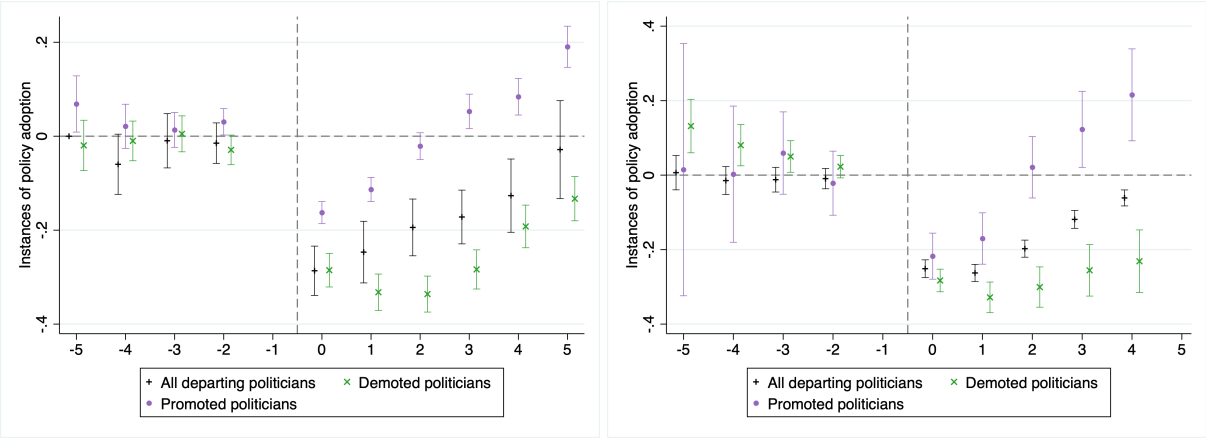


Panel A: Bottom-up innovation

Panel B: Top-down compliance

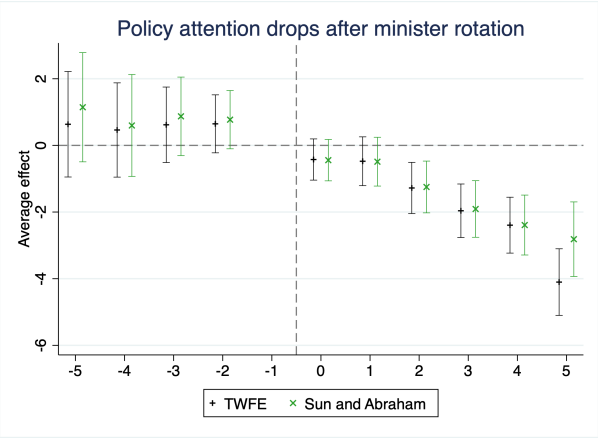
Notes: Panel A shows the geographical variation of average innovativeness between 2004 and 2020. Panel B shows the geographical variation of average compliance index across all years. White-shaded areas represent autonomous prefectures or sub-prefectural/county-level administrative units for which no data are available.

Figure A.3: Bureaucratic turnover and policy diffusion: robustness



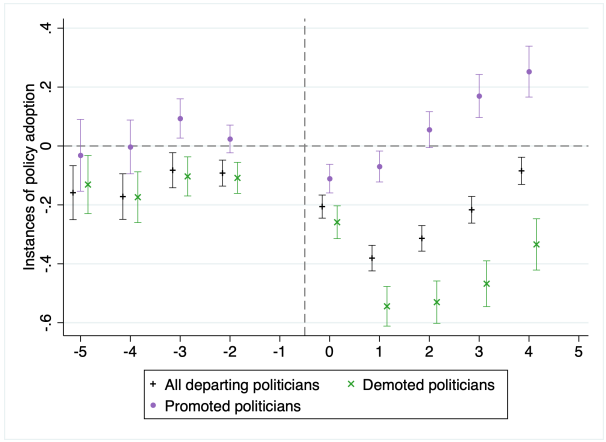
Notes: These figures plot event-study estimates illustrating the slowdown in policy diffusion following the departure of the prefectural party secretary who introduced the policy. Standard errors are clustered at the prefectural level, and we compare baseline estimates with cases in which departing officials were promoted or demoted. In Panel A, we follow Sun and Abraham (2021) to adjust for negative weights in two-way fixed-effects (TWFE) regressions, treating the final departure cohort in our data (2022) as the never-treated control group. In Panel B, we adopt a more stringent definition of promotion, counting only upward moves to (vice) provincial governor or party secretary (3% of cases).

Figure A.4: (Central) bureaucratic turnover and policy diffusion

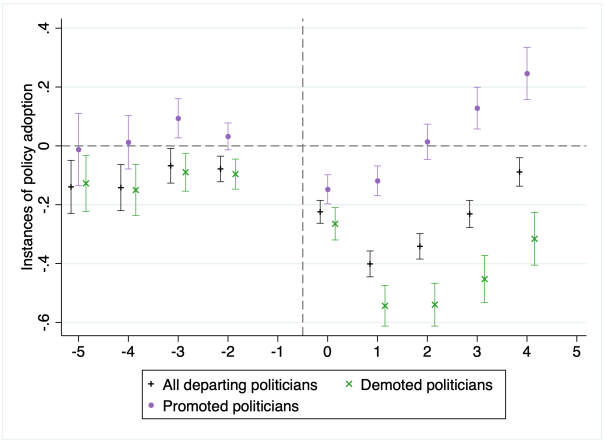


Notes: These figures plot event-study estimates illustrating the slowdown in policy diffusion following the departure of the central minister who introduced the policy. Standard errors are clustered at the ministry level, and we follow Sun and Abraham (2021) to adjust for negative weights in two-way fixed-effects (TWFE) regressions.

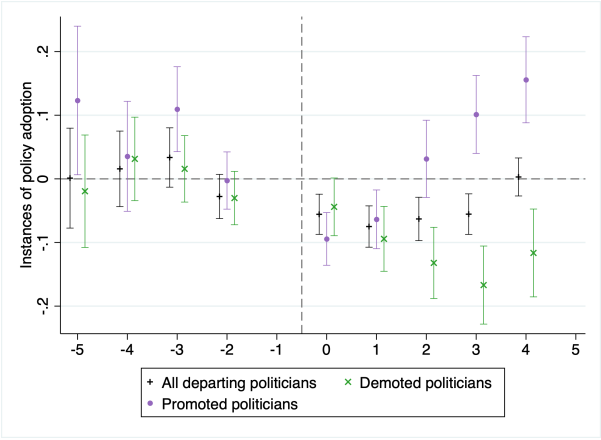
Figure A.5: Bureaucratic turnover and policy diffusion: alternative policy definitions



Panel A. Aggregated version



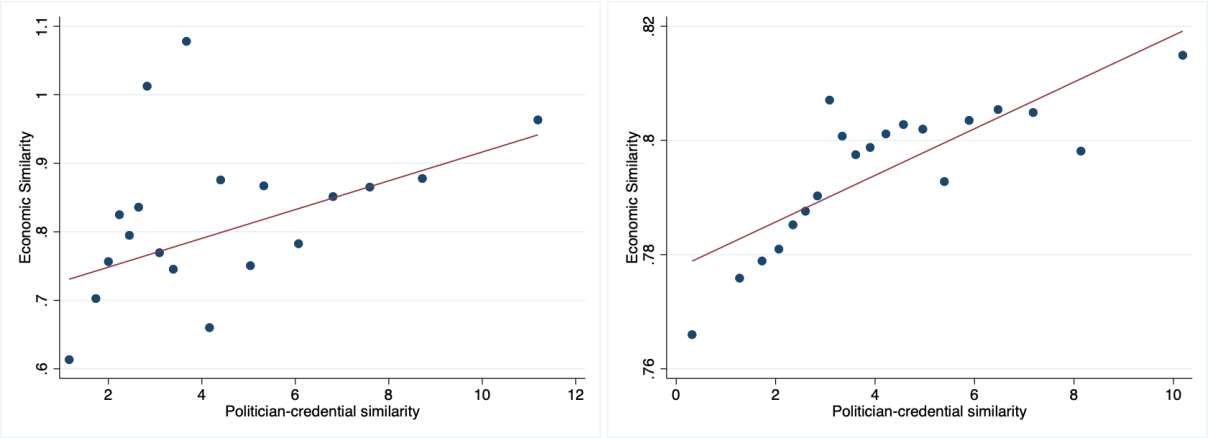
Panel B. Disaggregated version



Panel C. Title keywords

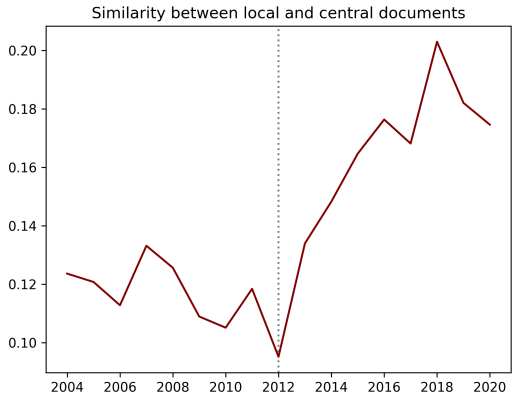
Notes: These figures plot event-study estimates illustrating the slowdown in policy diffusion following the departure of the prefectural party secretary who introduced the policy. Standard errors are clustered at the prefectural level, and we compare baseline estimates with cases in which departing officials were promoted or demoted. In Panel A, we aggregate similar policies into clusters. In Panel B, we disaggregate large policy bundles into units comparable with the rest of the sample by treating each policy $\times$ domain combination, rather than each policy, as the unit of observation. In Panel C, we replicate the analysis using a new sample extracted from policy titles.

Figure A.6: Political similarity and economic similarity



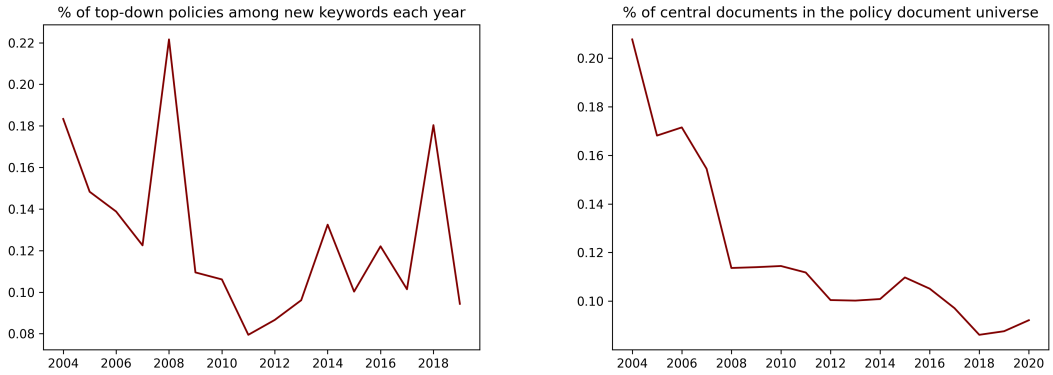
Notes: This figure illustrates the correlation between economic conditions and politician characteristics across prefecture-pairs. Panel A presents a binned scatter plot of economic distance (measured by GDP per capita dispersion) against political distance (measured by the Mahalanobis distance between politicians' feature vector). Panel B parses out origin prefecture, destination prefecture, and year fixed effects.

Figure A.7: Textual similarity between central document and local follow-up: robustness



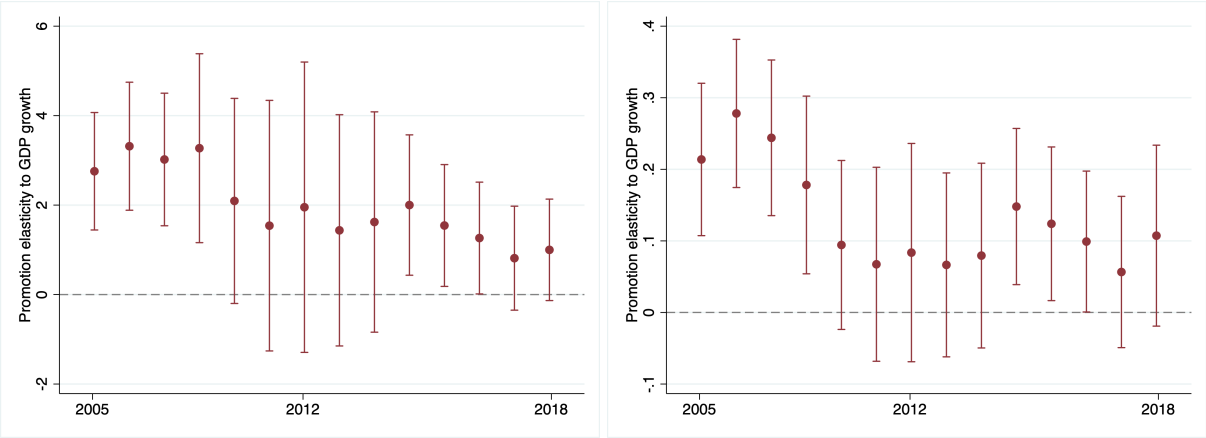
Notes: Similar to 3, Panel C, in this figure we compute the average textual similarity between the first central document on a certain topic and local follow-ups with TF-IDF algorithm, cosine similarity, and standard stop-word removal. We exclude the first opening paragraph from each policy document so that we can focus on concrete policy agendas instead of slogans, quotes and guidelines.

Figure A.8: Absence of the trend toward centralization looking at new documents



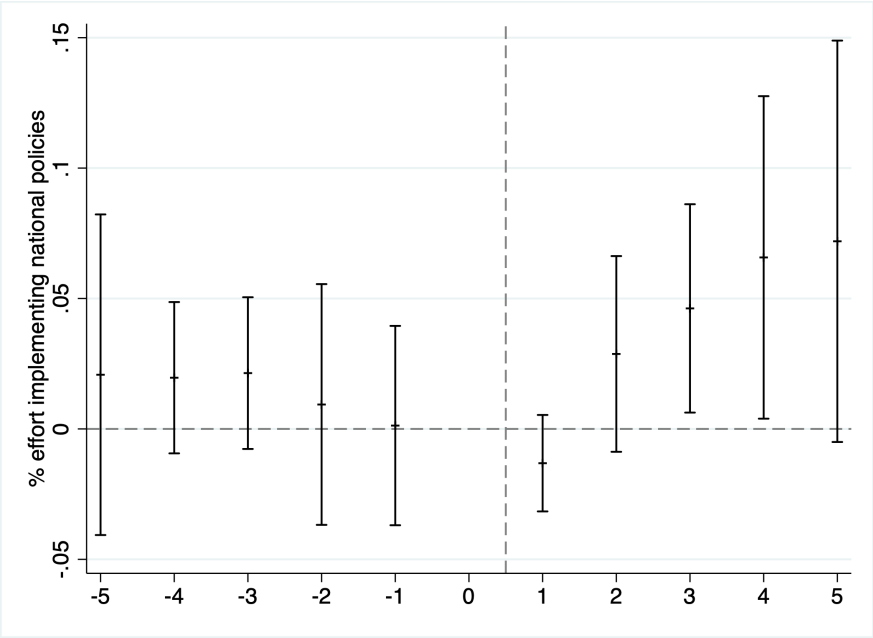
In the left panel, we show the share of new policy keywords that are initiated by the central government each year. In the right panel, we show, in the universe of policy document corpus, the % of documents issued by central government each year. In both cases, one fails to recognize a trend toward centralized policymaking.

Figure A.9: Career incentives and GDP



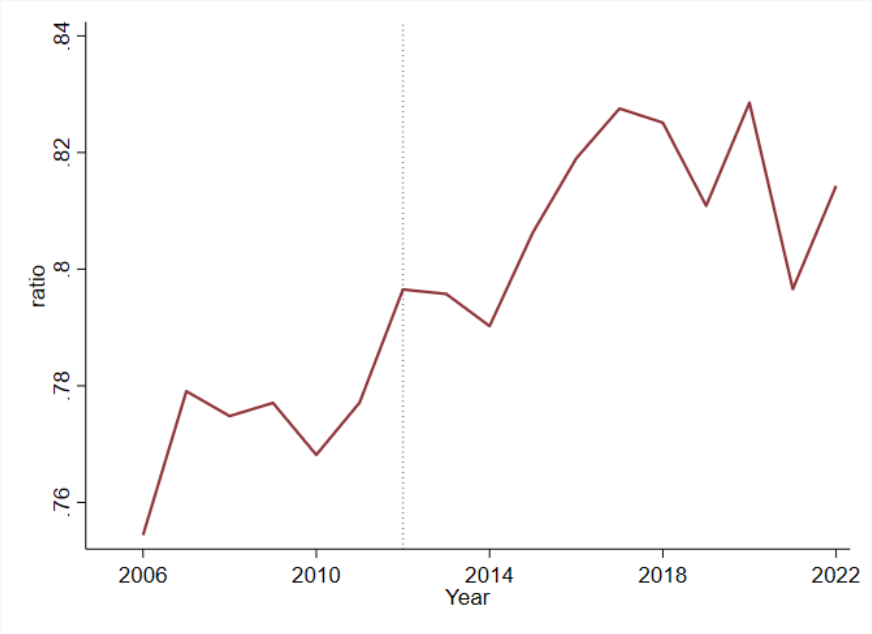
Notes: This set of figures plots the point estimates and 95% confidence intervals from a series of cross-sectional regressions. In each regression, we focus on politicians departing office within a sliding window $[t-2, t+2]$. We regress their career outcome (promotion$\in \{0,1\}$) on an innovation index, a compliance index, and a GDP growth index—the canonical predictor à la Li and Zhou (2005). In the left panel, GDP growth is measured in raw percentage terms; in the right panel, it is standardized within each year. All indices are scaled by the mean of the dependent variable so that the coefficients can be interpreted as elasticities. Standard errors are clustered at the prefecture level.

Figure A.10: Working groups and centralization



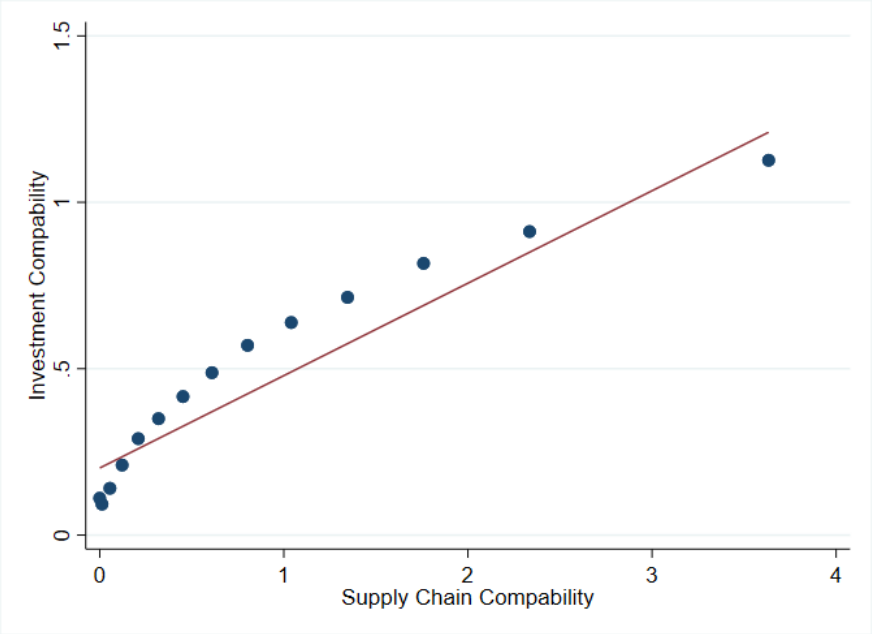
Notes: This figure showcases the increase in local politicians' effort of compliance within the policy domains where a central government working group, chaired by Xi, has been organized. Specifically, we compiled data at policy-domain-by-year level and plotted the TWFE estimates as well as adjustments à la Sun and Abraham (2021) of a regression of compliance effort on year dummies relative to the formation of working groups. The regression is weighted by the number of policy documents issued in each domain within each year. Standard errors are clustered at policy domain level.

Figure A.11: Centralization of industrial policies



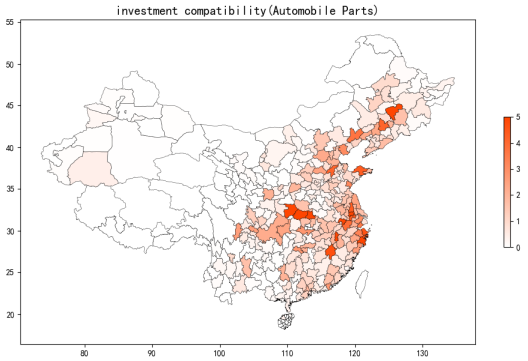
Notes: This figure plots the share of local industrial policies that are assigned from top-down. For each locality–year, we compute the share of top-down industrial policies among all implemented industrial policies and then average that share across localities, and observe a sharp increase in centralization over time.

Figure A.12: Correlation between supply chain and investment suitabilities

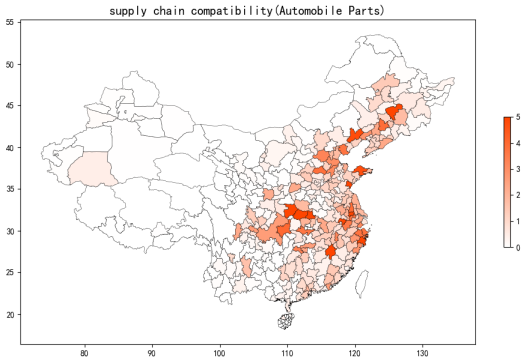


Notes: This figure plots the correlation between Investment suitability with supply chain suitability. For a given locality, the two measures are positively associated with each other.

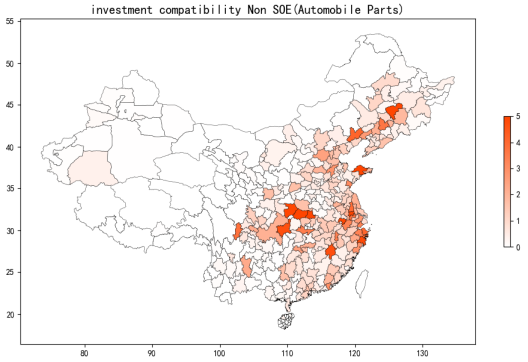
Figure A.13: Visualization of policy suitability



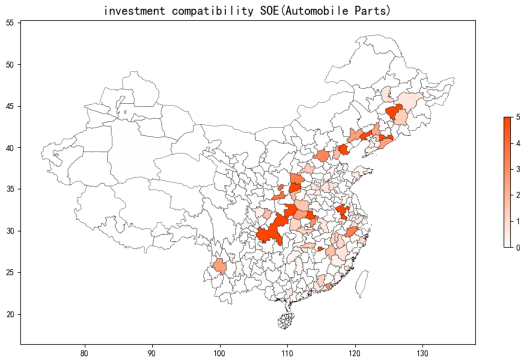
Panel A: Investment suitability of automobile sector



Panel B: Supply chain suitability of automobile sector



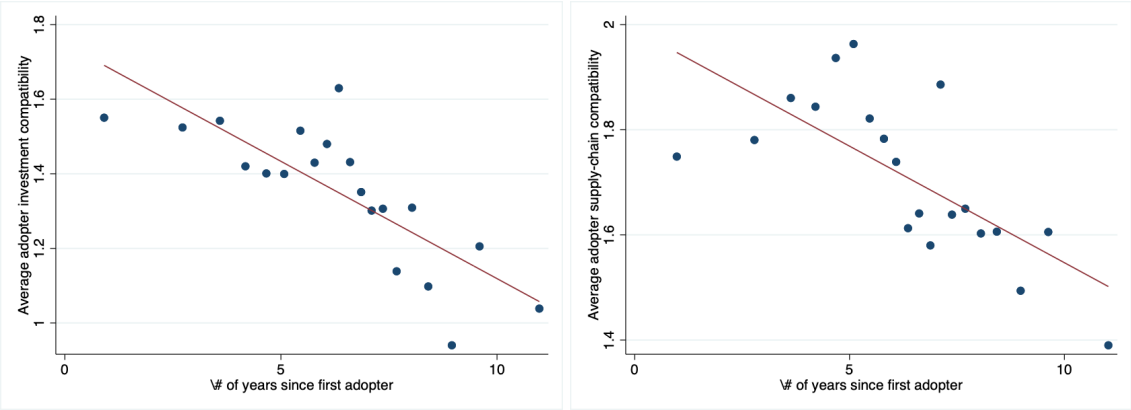
Panel C: Investment suitability of automobile sector, non-SOE



Panel D: Investment suitability of automobile, SOE

Notes: This figure maps automobile-industry suitability across China. Panel A measures suitability using pre-existing investment flows; Panel B measures suitability based on pre-existing supply-chain strength. Panels C and D measure investment-based suitability separately for private firms and SOEs, respectively.

Figure A.14: Policy suitability and sequence of adoption

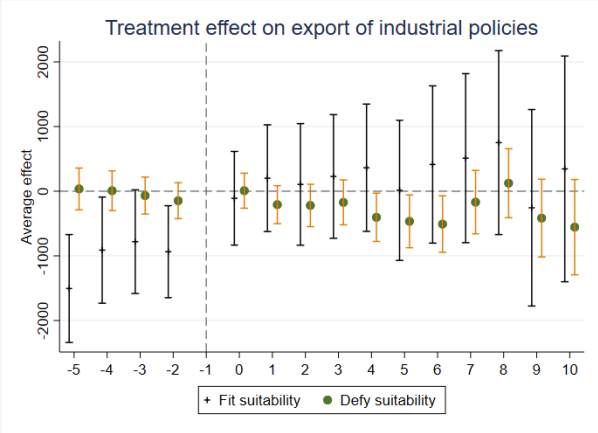


Panel A: Investment suitability

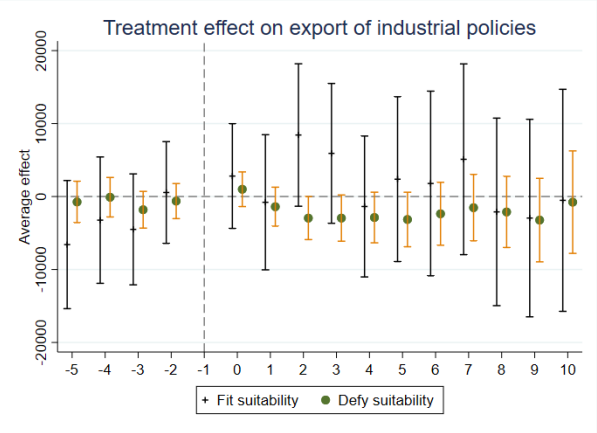
Panel B: Supply chain suitability

Notes: This figure plots the relationship between policy suitability and the order of policy adoption, showing that early adopters of an industrial policy exhibit, on average, higher suitability than later adopters. Prefecture, industry and year fixed effects are parsed out, before we plot the binned scatter plot. Panel A uses the investment-based suitability measure, while Panel B uses the supply-chain-based suitability measure.

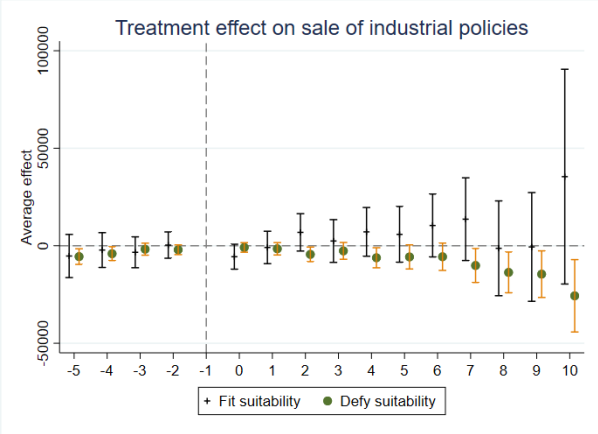
Figure A.15: Local suitability and industrial policy effectiveness



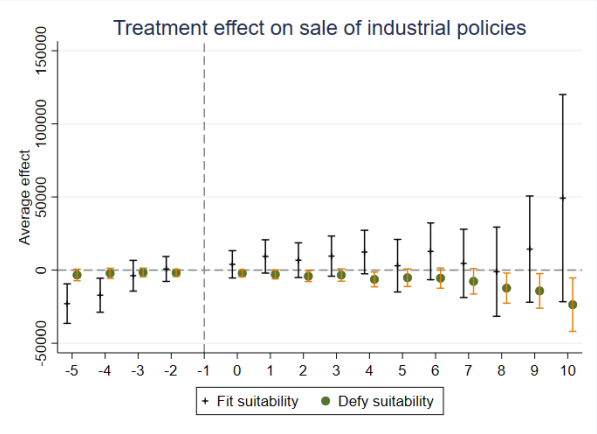
Panel A.1: Export value (1 thousand yuan), Supply-chain suitability



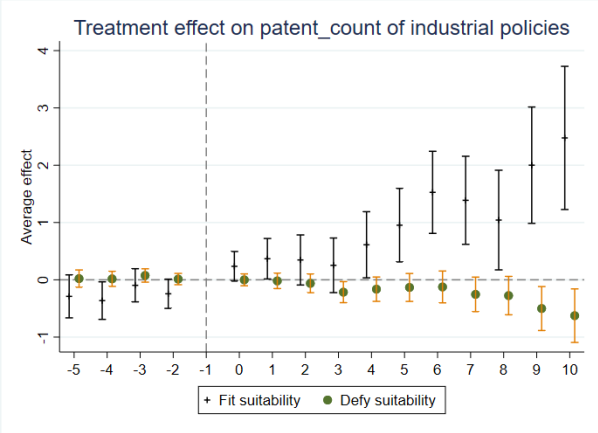
Panel B.1: Export value (1 thousand yuan), Investment suitability



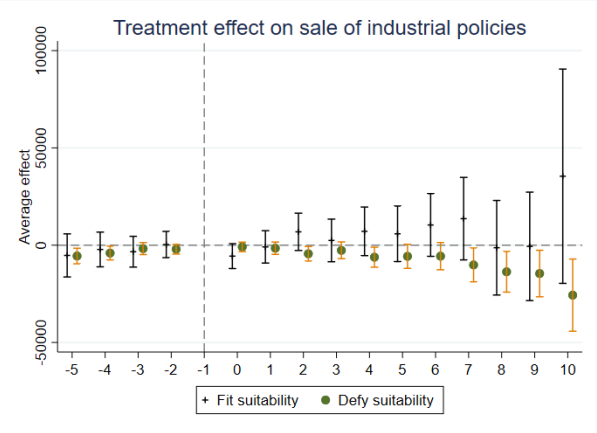
Panel A.2: Total revenue (1 thousand yuan), Supply-chain suitability



Panel B.2: Total revenue (1 thousand yuan), Investment suitability



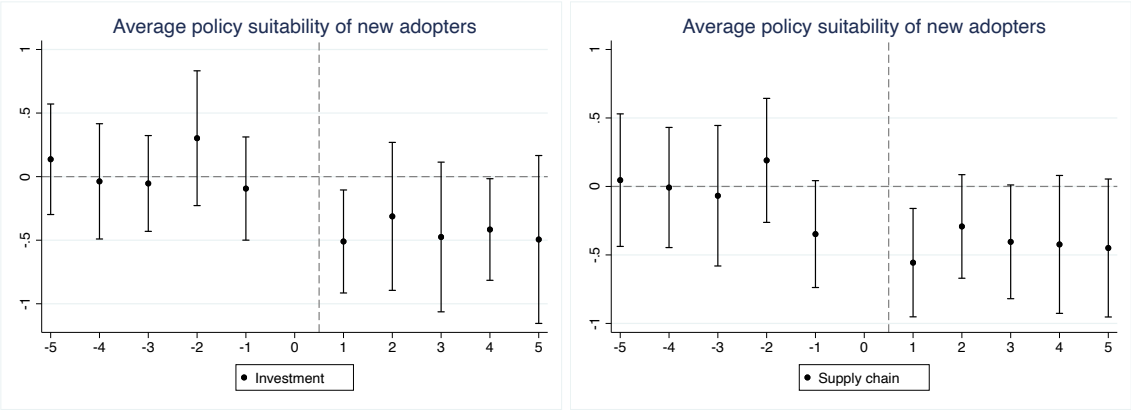
Panel A.3: Patent (counts), Supply-chain suitability



Panel B.3: Patent (counts), Investment suitability

Notes: This figure plots event-study coefficients estimating how industrial policies influence key outcomes under conditions of either fit or defy local suitability. Specifically, fit suitability refers to cases where local suitability falls within the top 10% of all observations, while defy suitability refers to those within the remaining 90%.Panels A.1–A.3 present results based on *ex ante* supply-chain suitability, while Panels B.1–B.3 show results based on *ex ante* investment suitability. Outcomes include export value, firm revenue, and patent registrations. All regressions include prefecture-by-year, prefecture-by-industry, and industry-by-year fixed effects, with standard errors clustered at the prefecture level.

Figure A.16: Central adoption and policy suitability



Panel A: Investment suitability

Panel B: Supply chain suitability

Notes: This figure plots the event-study coefficients, showing that central-government adoption of a given industrial policy is associated with lower policy suitability among subsequent adopters. Panel A uses the investment-based suitability measure, while Panel B uses the supply-chain-based suitability measure.

Table A.1: Policy examples

Policy	Inception year	First adopter
Primary and Secondary School Teacher Title System Reform	2009	Shaanxi, Shanghai
Cross-Region Housing Provident Fund Loans	2010	Liaoning, Hubei
Long-Term Management of Village Environment	2012	Jiangsu, Fujian
Control and Demolition of Illegal Constructions	2004	Hunan
Oil and Gas Recovery at Gas Stations	2007	Tianjin
Pilot of New Rural Cooperative Finance	2014	Shandong, Beijing, Hebei
Detailed Survey of Soil Pollution	2013	Jiangsu
River Chief System	2003	Zhejiang
Overseas Chinese Investment and Talent Introduction	2002	Chongqing
End-to-End Online Processing	2008	Shanxi

Notes: This table lists 10 policies randomly selected from our sample, showing each policy's name, inception year, and first adopter(s).

Table A.2: Policy Content and Domain

Policy	Content	Domain
Primary and Secondary School Teacher Title System Reform	The policy unified the previously separate ranking systems for middle and elementary school teachers into a single five-level hierarchy.	Education
Cross-Region Housing Provident Fund Loans	The policy allows the use of housing provident funds across different regions.	Social Welfare
Long-Term Management of Village Environment	The policy is aimed at maintaining rural environmental improvements, covering sanitation, river maintenance, road upkeep, greening, and public facilities maintenance.	Environment
Control and Demolition of Illegal Constructions	This policy is aimed at identifying, stopping, and removing unauthorized buildings.	Urban Management
Oil and Gas Recovery at Gas Stations	This policy involves installing vapor recovery systems to reduce air pollution from gasoline evaporation.	Environment
Pilot of New Rural Cooperative Finance	This policy aims to expand formal financial services in rural areas through cooperative models, improving credit access for farmers and small businesses.	Rural Development
Detailed Survey of Soil Pollution	This policy is an investigation to systematically assesses soil contamination levels, particularly focusing on agricultural and industrial lands.	Environment
River Chief System	This policy introduces a water governance mechanism that assigns government leaders at all levels as "river chiefs" responsible for comprehensive waterway management, integrating pollution control, and cross-regional coordination.	Environment
Overseas Chinese Investment and Talent Introduction	The policy leverages diaspora resources by attracting overseas Chinese capital, expertise, and innovation networks to support regional development priorities.	Economic Development
End-to-End Online Processing	This policy encourages a digital government service model that allows applicants to complete entire approval processes online without visiting physical service centers.	Governance

Notes: The table lists policies in Table A.1 , summarizing their content and corresponding policy domain.

Table A.3: Decomposition of innovation between bureaucrats and locality

Decomposing innovation and compliance: alternative policy definitions			
	$\tau_{\text{politician}}$	$\tau_{\text{prefecture}}$	τ_{year}
	(1)	(2)	(3)
Panel A: Aggregating similar keywords			
Variation of innovativeness explained	0.304***	0.056*	0.140***
	(0.038)	(0.033)	(0.015)
Variation of compliance explained	0.175***	0.112**	0.301***
	(0.046)	(0.044)	(0.016)
Panel B: Disaggregating by domain $\times$ keyword			
Variation of innovativeness explained	0.250***	0.155***	0.085***
	(0.038)	(0.035)	(0.009)
Variation of compliance explained	0.223***	0.063**	0.293***
	(0.035)	(0.031)	(0.016)
Panel C: Keywords extracted directly from titles			
Variation of innovativeness explained	0.262***	0.042	0.449***
	(0.048)	(0.036)	(0.032)
Variation of compliance explained	0.338***	0.085	0.131***
	(0.056)	(0.054)	(0.012)

Notes: In all three panels, we follow Abowd, Kramarz, and Margolis (1999) to decompose the bottom-up innovation index and compliance index into bureaucrat fixed effects, locality fixed effects, and calendar year fixed effects. In Panel A, we aggregate similar policies into clusters. In Panel B, we disaggregate large policy bundles into units comparable with the rest of the sample by treating each policy $\times$ domain combination, rather than each policy, as the unit of observation. In Panel C, we replicate the analysis using a new sample extracted from policy titles. Akin to employer-employee matched design, identification exploits variation of innovation index within (rotating) bureaucrats across the places they hold office. After the decomposition, we report the normalized correlation between each set of fixed effects and the outcome variable. In the parentheses below the estimates, we report standard errors with 1,000 bootstraps.

Table A.4: Policy portfolio and proximity

	Similarity in policy portfolio		
	(1)	(2)	(3)
Panel A: Economic Proximity			
-Δ GDP per capita	2.226***	0.979***	0.348***
	(0.037)	(0.025)	(0.022)
Panel B: Political Proximity			
- Δ politician characteristics	-0.561***	-0.007**	-0.027***
	(0.007)	(0.003)	(0.003)
Year FE	No	Yes	Yes
Prefecture FE	Yes	Yes	No
Prefecture-pair FE	No	No	Yes

Notes: The dependent variable is the Euclidean distance between vectors denoting policy portfolios across all prefecture pairs from 2003-2020. In Panel A, we measure economic proximity by the difference between city pairs of GDP per capita in units of standard deviation each year. Normalizing by year eliminates the mechanical variation attributable to increasing scale. In Panel B, we measure political proximity by the Mahalanobis distance between all observable features of politician, which we describe in detail in Section 2, drawing intuition from the fact that politicians of similar age with similar backgrounds might be compared against one another in the political tournament. Across columns both panels, we sequentially control for prefecture fixed effects, calendar year fixed effects, as well as prefecture-pair fixed effects. Standard errors are always clustered at the prefecture-pair level.

Table A.5: Politician competition and directed policy diffusion

	Policy diffusion		
	(1)	(2)	(3)
Panel A: Economic proximity			
Economic proximity	0.007***	0.013***	0.004***
	(0.001)	(0.001)	(0.001)
Economic indicator	GDP per capita	GDP	Population
Origin prefecture FE	Yes	Yes	Yes
Destination prefecture FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Panel B: Political proximity			
Political proximity	-0.001***	-0.001***	-0.002***
	(0.0001)	(0.0001)	(0.0001)
Origin prefecture FE	Yes	No	No
Destination prefecture FE	Yes	No	No
Prefecture pair FE	No	Yes	No
Prefecture (d) $\times$ prefecture (o) FE	No	Yes	No
Year FE	Yes	Yes	Yes

Notes: We define policy diffusion as an indicator equal to 1 when prefecture i adopts, in year t, a policy first introduced by prefecture j. In Panel A, economic proximity is the absolute difference between city pairs in GDP per capita, total GDP, and fiscal expenditure, each expressed in standard-deviation units; normalizing by year removes mechanical variation from scale growth. In Panel B, political proximity is the Mahalanobis distance between politician characteristics. Across columns, we sequentially add controls: origin- and destination-prefecture fixed effects; prefecture-pair fixed effects; politician $\times$ target-location fixed effects; and year fixed effects. Standard errors clustered at the prefecture-pair level are reported below the estimates.

Table A.6: Policy portfolio and proximity: interaction term

	Similarity in policy portfolio		
	(1)	(2)	(3)
Economic proximity	0.655***	3.574***	0.187***
	(0.026)	(0.063)	(0.021)
Political proximity	-0.060***	-0.026***	-0.027***
	(0.004)	(0.004)	(0.005)
Economic proximity × political proximity	-0.053***	-0.036***	-0.022***
	(0.002)	(0.005)	(0.003)
Economic indicators	GDP per capita	Fiscal revenue	% tertiary GDP
Year FE	Yes	Yes	Yes
prefecture FE	Yes	Yes	Yes

Notes: The dependent variable is the Euclidean distance between vectors denoting policy portfolios across all prefecture pairs from 2003-2020. We measure economic proximity by the difference between city pairs of GDP per capita, fiscal revenue, or % of GDP contribution from tertiary industry in units of standard deviation each year. Normalizing by year eliminates the mechanical variation attributable to increasing scale. We measure political proximity by the Mahalanobis distance between all observable features of politician, which we describe in detail in Section 2, drawing intuition from the fact that politicians of similar age with similar backgrounds might be compared against one another in the political tournament. Across columns both panels, we control for origin and destination fixed effects and calendar year fixed effects. Standard errors are clustered at the prefecture-pair level.

Table A.7: Innovation, compliance and promotion likelihood

	Promotion					
	Before 2012			After 2012		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Baseline: innovation / compliance index						
Average innovation index	0.202***	0.196***	0.173**	0.011	0.020	0.020
	(0.072)	(0.072)	(0.072)	(0.076)	(0.080)	(0.080)
Average compliance index	-0.069	-0.053	-0.065	0.177**	0.166**	0.166**
	(0.067)	(0.068)	(0.068)	(0.079)	(0.082)	(0.082)
Δ GDP	1.157***	1.145***	1.241***	0.357**	0.340**	0.341**
	(0.331)	(0.339)	(0.331)	(0.159)	(0.161)	(0.162)
Panel B: Sum of successful innovations / early follow-ups						
Successful innovation	0.005***	0.005***	0.004***	0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Early adopters	-0.003	-0.000	-0.001	0.011**	0.011**	0.011**
	(0.009)	(0.009)	(0.009)	(0.005)	(0.005)	(0.005)
Panel C: Robustness: alternative measures of innovation / compliance						
$\sum$ innovation index	0.131***	0.126***	0.106***	-0.022	-0.028	-0.030
	(0.041)	(0.041)	(0.041)	(0.040)	(0.041)	(0.042)
$\sum$ implementation index	-0.068*	-0.060	-0.049	0.068	0.076	0.077
	(0.039)	(0.038)	(0.038)	(0.046)	(0.047)	(0.048)
Δ GDP	1.091***	1.112***	1.182***	0.294*	0.287*	0.290*
	(0.355)	(0.362)	(0.354)	(0.158)	(0.160)	(0.162)
Panel D: Placebo: muted effects for politicians $>$ 55 years old						
Average innovation index	0.038	0.148	0.118	0.138	0.165	0.148
	(0.366)	(0.361)	(0.376)	(0.211)	(0.203)	(0.203)
Average compliance index	0.285	0.248	0.237	0.004	-0.003	0.003
	(0.303)	(0.297)	(0.298)	(0.226)	(0.228)	(0.232)
Δ GDP	-0.915	-0.493	-0.498	0.608*	0.639*	0.648*
	(1.987)	(2.165)	(2.220)	(0.332)	(0.372)	(0.353)
# of obs.	571	571	571	872	872	872
Mean of DV	0.401	0.401	0.401	0.317	0.317	0.317
Start cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Finish cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Start Age FE	No	Yes	Yes	No	Yes	Yes
Level of education FE	No	No	Yes	No	No	Yes

Notes: Each column within each panel comes from a politician-level regression. In Panel A, we regress promotion indicators on the logged innovation and compliance indices, which capture both the speed of action and policy success. In Panel B, we unpack these indices: instead of the innovation index, we count the number of locally initiated policies that became national policies in the past three years; instead of the compliance index, we count how often a locality adopted a national policy within its first three years. "Start cohort" is the year a politician takes office interacted with his or her position hierarchy; "Finish cohort" is the year a politician leaves office interacted with position hierarchy; "Start age" is the age at the start of tenure; and "Level of education" is a dummy for having obtained a postgraduate degree prior to taking office. Standard errors are clustered at the prefectural level.

Table A.8: Innovation, compliance and investigation likelihood

	Investigation					
	Before 2012			After 2012		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Baseline: innovation / compliance index						
Average innovation index	-0.025	-0.024	-0.023	-0.070	-0.069	-0.070
	(0.027)	(0.028)	(0.028)	(0.053)	(0.053)	(0.053)
Average compliance index	0.005	0.001	0.002	0.032	0.026	0.025
	(0.027)	(0.027)	(0.028)	(0.060)	(0.060)	(0.060)
Δ GDP	-0.195*	-0.198*	-0.202*	-0.003	-0.023	-0.021
	(0.114)	(0.120)	(0.120)	(0.077)	(0.083)	(0.083)
Panel B: Robustness: alternative measures of innovation / compliance						
$\sum$ innovation index	-0.003	0.000	0.001	-0.025	-0.023	-0.025
	(0.015)	(0.015)	(0.015)	(0.024)	(0.024)	(0.024)
$\sum$ compliance index	-0.001	-0.004	-0.005	0.013	0.009	0.010
	(0.013)	(0.014)	(0.014)	(0.025)	(0.026)	(0.026)
Δ GDP	-0.185*	-0.187	-0.190	-0.001	-0.020	-0.017
	(0.111)	(0.116)	(0.117)	(0.074)	(0.080)	(0.081)
# of obs.	613	611	611	807	807	807
Start cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Finish cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Start Age FE	No	Yes	Yes	No	Yes	Yes
Level of education FE	No	No	Yes	No	No	Yes

Notes: Each column within each panel comes from a politician-level regression. In Panel A, we regress the probability of officials being subject to a corruption investigation on the logged innovation and compliance indices. In Panel B, we use the sum of the innovation and promotion indices. "Start cohort" is the year a politician takes office interacted with his or her position hierarchy; "Finish cohort" is the year a politician leaves office interacted with his or her position hierarchy; "Start age" is the age of the politician at the start of tenure; and "Level of education" is a dummy variable indicating whether they obtained a post-graduate degree prior to their term in office. Standard errors are clustered at the prefectural level.

Table A.9: Divergence in Policy Effectiveness: Triple Differences based on SOE

	Exports		Sales		Patents	
	Total	Within SOE	Total	Within SOE	Total	Within SOE
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Investment suitability						
DDD	696.40***	21.41*	-25.42	7.87	0.02*	-0.00
	(257.80)	(12.10)	(456.80)	(21.93)	(0.01)	(0.00)
# of obs.	2,118,464	2,118,464	2,481,648	2,481,648	3,477,600	3,477,600
Panel B: Supply chain suitability						
DDD	159.50	1.17	3,243	-11.10	0.10	0.00
	(167.30)	(1.61)	(1,974)	(24.45)	(0.10)	(0.00)
# of obs.	1,765,008	1,765,008	2,042,456	2,042,456	2,884,000	2,884,000
Mean of DV	26,745	26,745	45,199	45,199	3.41	3.41
Prefecture × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Prefecture × Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports triple-difference estimates of how policy effectiveness varies with preference revealed by State-owned-enterprise (SOE) companies across three outcomes—exports (cols. 1–2), sales (cols. 3–4), and patents (cols. 5–6). In odd-numbered columns, the outcome covers all firms; in even-numbered columns, it is restricted to SOE-driven activity. Suitability in columns 1–4 is calculated using SOE investment or supply-chain data; columns 5–6 use full-sample suitability measures. Panel A focuses on investment suitability; Panel B on supply-chain suitability. All regressions include prefecture-by-year, prefecture-by-industry, and industry-by-year fixed effects, with standard errors clustered at the prefectural level. Robust standard errors in parentheses

Table A.10: Divergence in Policy Effectiveness: Triple Difference

	Exports	Sales	Patents
	(1)	(2)	(3)
Panel A: Investment suitability			
Policy × Suitability	5,495***	14,421***	0.454***
	(1,044)	(2,284)	(0.102)
# of obs	2,118,464	2,481,648	3,477,600
Panel B: Supply chain suitability			
Policy × Suitability	339.2***	8,740***	0.479***
	(69.76)	(1,423)	(0.0537)
# of obs	1,765,008	2,199,568	2,884,000
Mean of DV	26,745	45,199	3.4
Mean of DV for treated group	56,823	165,163	8.2
Prefecture × Year FE	Yes	Yes	Yes
Prefecture × Industry FE	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes

Notes: This table reports triple-difference estimates of how policy effectiveness varies with suitability across three outcomes—exports (col. 1), sales (col. 2), and patents (col. 3). Panel A focuses on investment suitability; Panel B on supply-chain suitability. All regressions include prefecture-by-year, prefecture-by-industry, and industry-by-year fixed effects, with standard errors clustered at the prefecture level.

Table A.11: Centralization and policy suitability

	Investment suitability		Supply-chain suitability	
	Continuous	Top 10%	Continuous	Top 10%
	(1)	(2)	(3)	(4)
Panel A:Documents about economic policies				
Top-down	-0.357***	-0.044***	-0.197***	-0.028**
	(0.083)	(0.008)	(0.052)	(0.011)
# of obs.	67,423	67,423	63,838	63,838
Panel B: Agricultural & manufacturing policies				
Top-down	-0.284***	-0.028***	-0.208***	-0.012
	(0.096)	(0.008)	(0.058)	(0.012)
# of obs.	33,316	33,316	31,082	31,082
Mean of DV	1.23	0.23	1.09	0.19
SD of DV	3.34	0.17	1.51	0.39
Prefecture $\times$ Year FE	Yes	Yes	Yes	Yes

Notes: Panel A restricts the sample to policy documents directly related to economic matters, such as fiscal policy and regional development, while excluding those concerning social security and public security. Panel B further narrows the scope to industrial policies in manufacturing and agriculture, excluding other sectors. All other specifications follow Table 3.

Table A.12: Centralization and policy suitability

	Investment		Supply-chain	
	% suitable		% suitable	
	(1)	(2)	(3)	(4)
Central adoption	-0.322***	-0.571***	-0.312**	-0.448**
	(0.118)	(0.161)	(0.133)	(0.214)
# relative years	-0.042**	-0.051**	-0.023	-0.027
	(0.020)	(0.021)	(0.018)	(0.022)
Central adoption × # relative years		0.029		0.016
		(0.018)		(0.027)
# of obs.	15,028	15,028	15,028	15,028
Prefecture FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of central-government policy adoption on policy suitability among subsequent local adopters. Data are organized at the prefecture–industry level, with each observation capturing the first time industry i was promoted in prefecture p. The dependent variable is suitability in year $t-1$. Standard errors are clustered at the prefectural level.

Table A.13: Politician competition and policy suitability

	Number of local policies		
	(1)	(2)	(3)
Panel A: Effects of economic policy intensity			
Competitors among top-30 neighbors	0.316		
	(0.264)		
Competitors among top-40 neighbors		0.221	
		(0.140)	
Competitors among top-50 neighbors			0.128
			(0.111)
# of obs.	3,871	3,871	3,871
Mean of DV	18.94	18.94	18.94
Prefecture FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: In this table, we regress the number of policies adopted by a given prefecture in a given year on the number of economic neighbors who are also political contenders to probe at extensive margin effects of political competition. Standard errors clustered at prefecture level are reported below the estimates.

Table A.14: Central Policy vs. Local Policy

	All policies	Stratifying local policies
Panel A: Long-run market potential		
	Market Value in 2024	Market Value in 2024
Top-down	18.7	-139.8**
	(49.6)	(65.1)
# of obs.	378	262
Panel B: strategic importance		
	% in Sanction List	% in Sanction List
Top-down	0.083*	-0.046
	(0.043)	(0.053)
# of obs.	376	258
Panel C: Comparative Advantage		
	RCA Growth	RCA Growth
Top-down	0.022	-0.705***
	(0.119)	(0.113)
# of obs.	376	258
Panel D: market distortion		
	Market Distortion	Market Distortion
Top-down	-0.053	-0.502***
	(0.065)	(0.061)
# of obs.	794	648
Panel E: economies of scale		
	Economies of Scale	Economies of Scale
Top-down	0.126	-0.630***
	(0.091)	(0.078)
# of obs.	427	250
Panel F: pollution intensity		
	% in Pollution List	% in Pollution List
Top-down	0.079**	0.002
	(0.031)	(0.036)
# of obs.	646	443

Notes: Market Value in 2024 refers to the global market value of a given industry in 2024. % in Sanction List indicates whether the industry appears on the U.S. International Trade Administration's sanction list for China. RCA Growth measures the change in China's Revealed Comparative Advantage in a given industry between 2000 and 2024. Market Distortion is measured by frictions accumulated through backward input–output linkages, following Liu (2019). Economies of Scale for a given industry are measured as the average level across France, Germany, Belgium, and the United Kingdom, as in Atkin, Costinot, and Fukui (2021). Pollution is a dummy variable equal to one if an industry is classified as pollution-intensive in the 2021 Comprehensive Directory of Environmental Protection, published by the Ministry of Ecology and Environment of China. The first column compares all bottom-up policies with top-down ones, while the second column compares equal numbers of local and central policies, stratifying local policies by the outcome of interest.

Table A.15: Centralization and policy suitability: robustness check

	Investment		Supply-chain	
	Continuous	Top 10 %	Continuous	Top 10 %
	(1)	(2)	(3)	(4)
Top-down	-0.193***	-0.025***	-0.158***	-0.014*
	(0.062)	(0.007)	(0.036)	(0.007)
Sanction	0.817***	0.052***	0.332***	0.104***
	(0.239)	(0.015)	(0.101)	(0.022)
Pollution	0.760***	0.053***	0.413***	0.126***
	(0.200)	(0.016)	(0.097)	(0.023)
# of obs.	118,104	118,104	114,484	114,484
Mean of DV	1.233	0.23	1.09	0.19
SD of DV	0.17	3.34	0.39	1.51
Prefecture $\times$ Year FE	Yes	Yes	Yes	Yes

Notes: This table reports point estimates of the difference in policy suitability between top-down and bottom-up industrial policies. Controls include industry characteristics — strategic importance (sanction-list inclusion) and pollution status. All specifications include prefecture–year fixed effects. Standard errors, clustered at the prefecture level, are reported in parentheses.

Appendix B Alternative definitions of policies

For the baseline analysis, we use keywords extracted from government reports as our basic unit of analysis. The clear advantage of this approach is that we do not need to impose an external definition of "what constitutes a policy"; instead, we follow governments' own definitions, as revealed by their summaries of policy initiatives in annual work reports.

Beyond this benefit, two additional issues are relevant for the empirical analysis and warrant further consideration. First, the labels that governments choose may not match the optimal level of granularity for our analysis: broad keywords risk obscuring meaningful heterogeneity, while overly narrow terms can fragment what is in practice a unified policy agenda.

Second, by drawing exclusively on annual work-report summaries, we may miss measures that never make it into those overviews—local regulations or directives recorded only in standalone policy documents and therefore absent from the corpus of extracted keywords.

To address the first set of considerations, we implement two complementary robustness checks: an aggregation exercise and a disaggregation exercise. The aggregation exercise groups together keywords that typically co-occur or refer to closely related policy changes. For instance, "Value Added Tax (VAT) reform" is often jointly discussed with "Abolition of Business Tax," reflecting a unified fiscal reform agenda; enabling us to group them together. Conversely, the disaggregation exercise breaks down broad agendas into more granular components. For example, although "VAT reform" appears as a single keyword, the actual rollout was sequential and domain-specific: starting with a pilot in the transportation sector in 2013, expanding to technology and business services in subsequent years, and culminating in a national policy in 2017. To capture this evolution, we disaggregate "VAT" into multiple agenda items by interacting it with policy domains, allowing each stage of reform to be counted as a distinct innovation.

To answer the second set of considerations, we approach the policy documents as if we don't know anything about government work reports, and rebuild our master dataset using only keywords extracted from policy document titles themselves. Policy documents in China tend to be short in general (most are below 5 pages), and address specific matters. The titles almost always convey valuable information on the topics and names of the campaigns. There's a stringent structure followed by all hierarchies regarding nomenclature — more than 98% of the titles says "(XX department)'s (**YY type of document**) about

(*ZZ topic*)".[1] With the new sample, we replicate our main results.

In all three alternative exercises, our main findings remain robust, suggesting that our core results are not sensitive to the level of agenda granularity or the source of keyword extraction. Below, we describe each of these strategies in greater detail.

Appendix B.1 Grouping similar policies into broader agendas

On the aggregation side, we identify similar policies via co-occurrence — for each keyword-pair, we compute the Jaccard similarity of the vector of documents they each appear in. For instance, if "value-added tax" appears in documents 1 through 100, "business tax" in documents 5 through 105. It translates into a Jaccard similarity of

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} \simeq 0.914$$

Based on this calculation, we group together all keyword pairs with a Jaccard similarity > 0.8, and consider them duplicates. In our final sample, we find 7.24% of the keywords substantially similar to at least one other policy, and therefore each of them is grouped with another policy in our aggregation.

Computing all pairwise Jaccard similarities between keyword vectors is computationally prohibitive. Let N be the number of keyword vectors, and let L be the average number of documents (i.e., non-zero entries) per vector. Calculating the exact Jaccard similarity between two vectors requires computing the size of their intersection and union — an $\mathcal{O}(L)$ operation. Therefore, a naïve all-pairs comparison has total time complexity

$$T_{\text{exact}}(N, L) = \binom{N}{2} \cdot \mathcal{O}(L) = \mathcal{O}(N^2 L)$$

In our case, we have $N > 110{,}000$ keywords and $L \approx 260$, leading to

$$\binom{110{,}000}{2} \approx 6.0 \times 10^9 \quad \text{pairs},$$

and

$$T_{\text{ops}} \approx 6.0 \times 10^9 \times 260 = 1.6 \times 10^{12} \quad \text{operations}.$$

This process is highly time-consuming, making the brute-force method infeasible at scale.

To address this computational bottleneck, we use MinHash signatures and LSH indexing (Broder, 1997; Indyk and Motwani, 1998). Each keyword vector is mapped to a compact signature of length k using k independent hash functions. Computing a Min-

1. For example, in Appendix Figure A.1, Panel C, it reads "Tianshui Prefectural Government's **Implementation Guidelines** for *Advancing a Water-Conserving Society*"

Hash signature is an $\mathcal{O}(kL)$ operation, and building the LSH index as well as querying it remains near-linear in N. The total complexity becomes:

$$T_{\text{LSH}}(N, L, k) = \mathcal{O}(NkL + N \log N),$$

with $k \ll L$. In our implementation, we use $k = 128$, which balances computational efficiency and accuracy for a Jaccard threshold of 0.8. On a single workstation, the full process finishes in under 4 minutes. This allows us to identify and remove near-duplicate policy keywords efficiently, without compromising the integrity of our analysis.

Appendix B.2 Unpacking policy bundles

On the disaggregation side, we decompose each policy into the Cartesian product of the keyword itself and the set of domains to which the policy applies. We group all ministries and departments in China into 16 main policy domains, à la Wang and Yang (2025). For example, "VAT" is decomposed into "VAT $\times$ transportation," "VAT $\times$ agriculture," "VAT $\times$ industrial technology," and so on. We then trace the inception, diffusion, and central — government action for each keyword $\times$ domain. In our final dataset, we identified 651,488 keyword $\times$ domain observations from approximately 110,000 raw policy keywords.

Appendix B.3 Alternative keywords from policy titles

To obtain an alternative dataset independent of government annual reports, we first extract policy-relevant keywords from the titles of all Chinese government documents. Specifically, we leverage GPT-4o to process batches of 100 titles at a time, using the following prompt:

> "I'm going to pass you one hundred titles of Chinese government documents; please extract one policy keyword from each title. Try to ensure they are specific policy terms rather than generic concepts."

This human-in-the-loop approach ensures consistent, semantically grounded identification of policy concepts across a diverse corpus.

Using the resulting set of keywords, we then trace the inception and diffusion of each policy idea across the full set of government documents. To do so efficiently, we implement the Aho–Corasick string-matching algorithm, which allows us to identify all keyword occurrences in the corpus in linear time. This approach yields a dynamic, fine-grained map of how policy ideas emerge and propagate through the bureaucratic system over time.

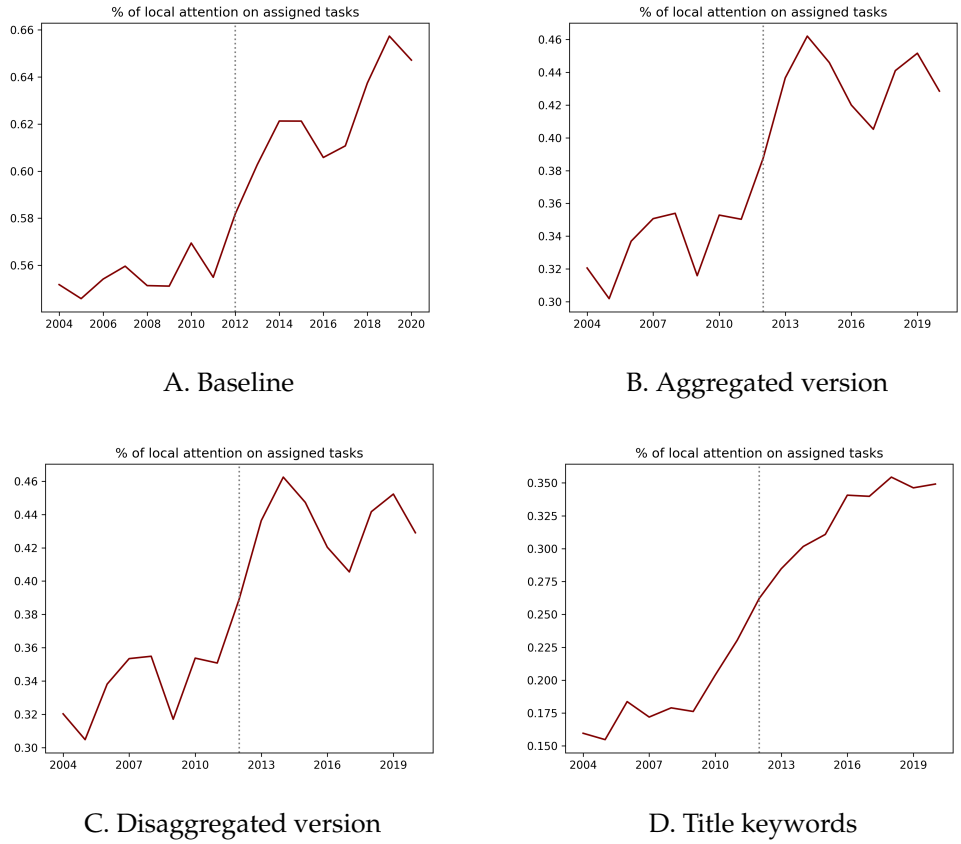
Our final dataset consists of 252,861 policy agendas spanning 1980-2023, 141,338 of which were initiated during 2004-2020. Reassuringly, this largely automated approach yields a figure of the same order of magnitude as the 115,679 agendas extracted from government reports during 2004-2020.

However, while this exercise allows us to map the full landscape of policy discourse, our baseline method—relying on keywords drawn from annual government work reports—remains preferable. These reports offer a curated summary of policy priorities as perceived and endorsed by political decision-makers themselves. Consequently, the keywords extracted from them tend to reflect the agendas that officials consider most salient, rather than incidental or administrative terms that may appear in the broader universe of documents. This process substantially reduces semantic noise and enhances the precision of our analysis, enabling a more accurate tracing of core policy ideas as they emerge, diffuse, and evolve over time.

Appendix B.4 Robustness of empirical results

All the main results in this paper are robust to these alternative sample-construction procedures. In Appendix Figure A.17, we plot the trend of centralization across the three alternative samples. In each case, we observe a discrete jump around 2013 of quantitatively similar magnitude in the effort politicians allocate to complying with the top-down policy agenda.

Figure A.17: Robustness of centralization trend



A. Baseline

B. Aggregated version

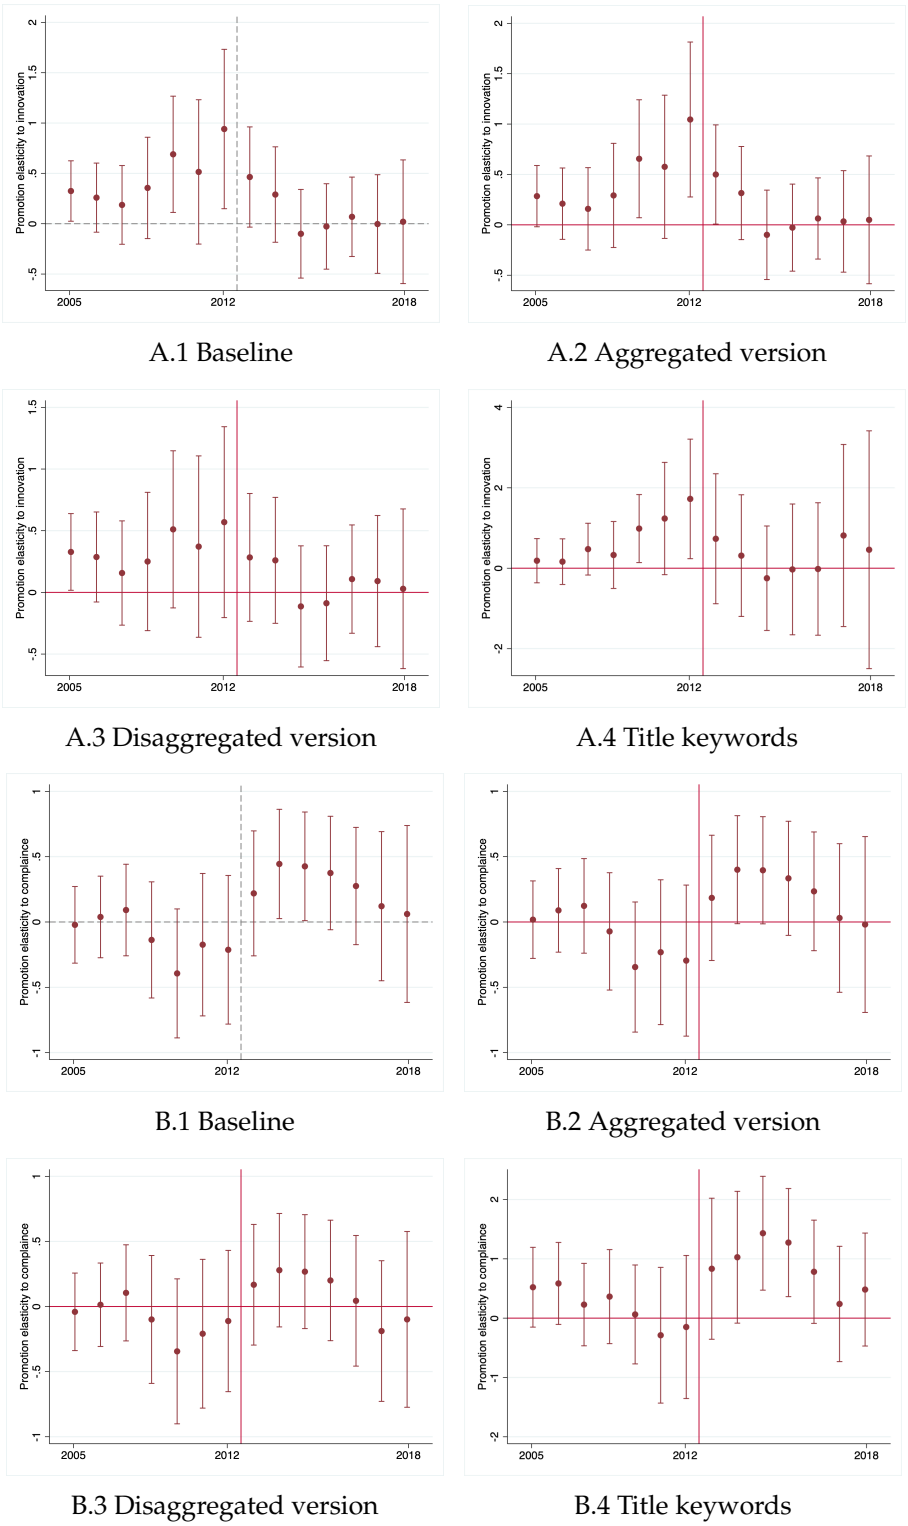
C. Disaggregated version

D. Title keywords

Notes: This figure illustrates the increasing centralization of policymaking under alternative samples. All four panels plot the percentage of local attention devoted to assigned tasks, by year. Specifically, for each locality-year, we compute the share of top-down policies among all implemented policies and then average these shares across localities. In Panel A, we reproduce the baseline estimates (as in Figure 3). In Panel B, we aggregate similar policies into clusters. In Panel C, we disaggregate large policy bundles into units comparable with the rest of the sample by treating each policy $\times$ domain combination, rather than each policy, as the unit of observation. In Panel D, we replicate the analysis using a new sample extracted from policy titles.

In Appendix Figure A.18, we replicate our career-incentive analysis using these alternative definitions. Using our baseline dataset, we observe that the promotion incentive for innovation gradually faded after 2013, while the reward for compliance became increasingly salient. Reassuringly, the same patterns emerge in the alternative samples.

Figure A.18: Robustness of evolving political incentives



A.1 Baseline

A.2 Aggregated version

A.3 Disaggregated version

A.4 Title keywords

B.1 Baseline

B.2 Aggregated version

B.3 Disaggregated version

B.4 Title keywords

Notes: Each subfigure shows point estimates and 90% confidence intervals from cross-sectional regressions. Panel A examines innovation rewards; Panel B examines compliance rewards. In each panel, subfigure 1 reproduces the baseline estimates (Figure 4); subfigure 2 aggregates similar policies; subfigure 3 disaggregates large policy bundles by treating each policy × domain as the unit of observation; and subfigure 4 replicates the analysis on a sample extracted from policy titles. Standard errors are clustered at the prefecture level.

Appendix C Anecdotal examples

Appendix C.1 Anecdotal example: vehicle restriction policy

In the 1990s and 2000s, vehicle ownership grew rapidly in China, causing severe congestion and pollution problems in major cities. In response, Shanghai initially experimented with various forms of a license plate auction system, eventually launching a multi-unit, discriminatory (pay-as-you-bid), dynamic auction in which residents bid for plates in monthly auctions, with revenues directed toward public transportation upgrades. The policy is widely praised for its allocative efficiency as well as the substantial financial support generated for public transportation (which benefits the poor).

Instead of adopting the Shanghai model — which had proven effective — Beijing opted to implement an alternative system in 2011. This system used free, random lotteries to allocate license plates to registered citizens. Due to the low lottery success rate, many citizens without an urgent need for vehicles participated preemptively in the lottery. Moreover, since license plates cannot be "banked" or "resold," citizens with low willingness to pay (WTP) for vehicles often purchased cars ahead of their high-WTP counterparts, thereby generating substantial welfare loss due to misallocation (Li 2018).

Many observers speculate that Beijing intentionally deviated from the Shanghai model of license plate auctions because of the political rivalry between the two major cities. When challenged on this point, a government official asserted that Beijing would *never* auction license plates, claiming that it "aims to protect the interests of the poor" (Wang and Zhao 2017).

Appendix C.2 Anecdotal example: wind energy development

A salient example of policy centralization leading to reduced policy suitability can be found in China's development of wind power.

In the 2000s, wind energy was an emerging industry in China that relied heavily on bottom-up industrial policy promotions. The decentralized promotion of wind energy initially concentrated in regions with favorable natural conditions — namely, the northwestern provinces such as Gansu, Xinjiang, and Inner Mongolia — where the vast Gobi deserts and steppes, characterized by high wind density (>300 W/m^2) and minimal land acquisition costs, provided an ideal setting for constructing large-scale wind farms. Under this bottom-up approach, by 2010, China had surpassed the U.S. as the world's largest wind power installer, with over 75% of its capacity concentrated in the northwest.

However, after 2013 the central government began aggressively promoting wind energy development from the top down. Some low-wind provinces, such as Hunan and Hubei, eager to signal responsiveness to central policy initiatives, rushed to replicate the northwestern model. As a result, many of these newly constructed wind installations — built in regions with wind densities below 200 W/m^2 and lacking access to ultra-high voltage transmission lines — operated at low capacity and were eventually abandoned, earning the moniker "ghost wind farms."[2]

2. Source: https://news.bjx.com.cn/html/20170221/809616-1.shtml

Appendix D Quantifying tradeoffs of centralization

Having documented the negative (Table3) and positive (Table 4) impacts of centralization on policy suitability, as well as the association between policy suitability and policy effectiveness (Table A.10), we quantify the impacts of centralization by projecting these effects onto the number of policies that shifted from bottom-up to top-down under this centralization trend. Specifically, we use the following formulas:

$$\text{Cost}_y = \Delta N_{topdown} \times \alpha_{cost} \times \beta_y \times \bar{y}$$

$$\text{Benefit}_y = N_{localpolicy} \times \alpha_{benefit} \times \beta_y \times \bar{y}$$

where $\Delta N_{topdown}$ is the excess number of top-down policies introduced by post-2013 centralization, and $N_{localpolicy}$ is the number of policies adopted via horizontal diffusion among local governments after 2013. The coefficient α_{cost} captures how much more top-down policies violate local conditions relative to bottom-up policies, while $\alpha_{benefit}$ measures the improved fit of locally diffused policies when the central government stops rewarding bottom-up innovation.[3] Finally, β_y captures the marginal effect of a one-unit increase in suitability on y (export, industrial output, or patent filings).[4]

To calculate $\Delta N_{topdown}$, we construct a counterfactual in which the share of top-down industrial policies after 2013 remains at its 2012 level, with total annual policy counts held constant. Under this scenario, 2,562 policies implemented as top-down since 2013 would instead have been bottom-up. Meanwhile, we calculate the total number of decentralized industrial policies that local governments adopted from each other through diffusion between 2013 and 2022, yielding $N_{localpolicy} = 9{,}376$.

Empirically, within a local government's policy portfolio, the average top-down industrial policy's suitability for local conditions is lower than that of the average bottom-up policy by 0.274 (0.201) according to the investment-based (supply-chain-based) suitability measure. Such a policy–locality mismatch quantifies one cost of centralization.

At the same time, decentralized competition among peer cities also induces policy–locality misalignments: for a given prefecture in a given year, having one additional economic neighbor governed by a political competitor lowers average suitability by 0.00743 (0.00769 under the supply-chain measure). This competitive penalty disappears after 2013, when the central government stopped rewarding bottom-up policy innovation. Given an average of 2.78 political competitors per city, top-down design therefore mitigates competition-induced misalignment by approximately 0.020 points for investment suitability and 0.021

3. α_{cost} is obtained from Table 3, while $\alpha_{benefit}$ is based on estimates reported in Table 4.

4. This is obtained by linking the triple-difference estimates of how suitability affects industrial growth rates (Table A.10) to the corresponding industries in the relevant years.

points for supply-chain suitability.[5]

Since Section 5.3 estimated the impact of local suitability on sales, exports, and patents, we can now translate suitability changes into economic impacts. Leveraging our investment suitability estimates, we calculate that the *yearly* cost of post-2013 centralization is 580 billion RMB in industrial output loss, 437 billion RMB in export loss, and 10,486 fewer patent filings, while the *yearly* benefit is 121 billion RMB in output gain, 91 billion RMB in export gain, and 2,194 additional patent filings.[6] Under both measures, the costs of centralization exceed the benefits by more than 400%.

For comparison, in the promoted industries, the average annual output during our sample period is 621 billion RMB, average annual exports amount to 711 billion RMB, and the average annual number of patent filings is 410,520. Benchmarking the net-cost figures against these baseline average figures indicates that the post-2013 centralization in policymaking has led to significant aggregate impacts on China's industrial policy growth.[7]

Costs and benefits of centralization: investment suitability

	N	α	β	$\bar{y}$	Yearly Cost/Benefit
Export					
Cost	2,562	0.274	0.205	2.294 billion RMB	437 billion RMB
Benefit	9,376	0.021	0.205	2.294 billion RMB	91 billion RMB
Sales					
Cost	2,562	0.274	0.319	1.962 billion RMB	580 billion RMB
Benefit	9,376	0.021	0.319	1.962 billion RMB	121 billion RMB
Patents					
Cost	2,562	0.274	0.133	84.7 patents	10,486 patents
Benefit	9,376	0.021	0.133	84.7 patents	2,194 patents

5. This number is based on the "top 30 political competitors in the top 30 economic neighbors" specification; using alternative top 40 or top 50 definitions yields similar results for this calculation.

6. According to supply-chain suitability measures, the annual costs are 377 billion RMB in industrial output loss, 28 billion RMB in export loss, and 11,871 fewer patent filings, and the benefits are 76 billion RMB in output gain, 5 billion RMB in export gain, and 2,396 additional patent filings.

7. According to our estimates, over the 13-year period from 2013 to 2025, centralization results in a cumulative net loss in industrial output of up to 5,967 billion RMB, equivalent to 4.4% of China's 2024 GDP.