

Staging Institutional Reform Across Space: Court Modernisation and the Geography of Legal Services in Contemporary China

Peiran Lu

lupr2003@bfsu.edu.cn

Beijing Foreign Studies University

Article

Keywords: China, judicial reform, institutional change, legal services, spatial inequality, comparative social change

Posted Date: September 30th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10370838/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Staging Institutional Reform Across Space: Court Modernisation and the Geography of Legal Services in Contemporary China

Keywords: China; judicial reform; institutional change; legal services; spatial inequality; comparative social change

Abstract

China's 2020 Civil Procedure Case-Separation Reform ran in twenty pilot cities before becoming national law in 2022. Using a panel of 296 cities for 2017–2023, this paper shows that the pilots — selected for their strength — were already diverging from other cities before the reform, so a difference-in-differences design cannot identify its effect. Identification rests on synthetic control: for the fifteen pilots whose pre-reform trajectory can be matched from non-pilot cities, the pre-existing divergence is eliminated and a post-reform gain remains, concentrated in the legal dimension and significantly larger than a placebo distribution ($p < 0.001$). Across nine spatial specifications the spillover to neighbouring cities is never negative: the pilots pull ahead without displacing their neighbours. Their advantage narrows once the reform becomes national. Staged legal modernisation produces a transitional spatial inequality, and the selectivity of pilot assignment is itself a finding about how reform travels across Chinese territory.

1. Introduction

In 2020, the first year of a national pilot, the courts of Weiyang District in Xi'an closed 28,428 civil cases — more than two-fifths above the previous year — in an average of forty-nine days, nearly a month faster than before (Supreme People's Court 2021). Weiyang is one of twenty prefecture-level cities that, between January 2020 and the end of 2021, ran China's Civil Procedure Case-Separation Reform ahead of the rest of the country; across all twenty, summary procedures disposed of more than two million cases in the first year alone, at an average of forty-eight days against a statutory ceiling of ninety. For the firms and lawyers who live with the cost and delay of commercial litigation, an improvement on that scale, available in some cities and not others, is the kind of difference that can tip a decision about where to register or where to practise. This paper asks what such a geographically uneven improvement does to the places that have it and, in particular, to the places next door.

The institutions that govern economic life are not spread evenly across a national territory. Courts that enforce contracts and resolve commercial disputes cluster where commerce is dense, and when a state sets out to reform them it rarely does so everywhere at once. In China the modernisation of such institutions has characteristically begun in a handful of chosen places and spread outward only later — the special economic zones of the 1980s, the pilot free-trade zones

of the 2010s, successive rounds of administrative-licensing reform — a pattern Heilmann (2008) called experimentation under hierarchy. That pattern has been studied mostly for what it reveals about how Beijing governs (He 2024; Luo, Wang, and Yang 2025). Less has been said about a question the staged design makes unavoidable: while a reform is still confined to its pilot sites, what becomes of the places that have not yet received it — above all, the places next door? This paper pursues that question through one well-bounded episode of legal modernisation.

The episode is China's Civil Procedure Case-Separation Reform (CPSR). On 28 December 2019 the Standing Committee of the National People's Congress authorised the Supreme People's Court to depart from designated provisions of the 2017 Civil Procedure Law for a two-year trial; on 15 January 2020 the Court issued its implementing measures and named twenty prefecture-level cities as pilots. The reform raised the small-claim ceiling, widened the summary procedure, allowed single-judge benches to hear second-instance appeals on settled facts, put online filing and remote hearings on a statutory footing, and expanded the judicial confirmation of mediated settlements — throughout, matching the weight of procedure to the complexity of the case. On 24 December 2021 its main elements were written into the amended Civil Procedure Law, in force from 1 January 2022. The CPSR is, in short, a clean instance of the staged pattern: a self-contained improvement, introduced in known places at a known moment and generalised soon after.

Using a balanced panel of 296 cities spanning 2017 to 2023, this paper traces the spatial pattern that accompanied the trial — and confronts, rather than assumes away, the central obstacle to reading it. The twenty pilots were chosen centrally for their existing capacity, and an event study shows them already pulling ahead of the rest of the country before 2020, most sharply in the legal environment. A conventional difference-in-differences design therefore cannot separate the reform from the trajectory that selected the pilots, and we do not treat its estimates as causal. Identification rests instead on synthetic control (Abadie, Diamond, and Hainmueller 2010): for the fifteen pilots whose pre-reform path can be reproduced from non-pilot cities, matching absorbs the pre-existing divergence and isolates a post-reform gain concentrated in the legal dimension, significantly larger than a placebo distribution of non-pilots ($p < 0.001$). Two further results complete the picture: across nine spatial-weight specifications the spillover of pilot status is never negative, ruling out displacement from nearby cities; and after the reform becomes national law the pilots' advantage begins to narrow, consistent with a premium resting partly on exclusive early access.

A further reading is worth stating at the outset, because the pre-reform divergence is a finding and not only an obstacle. That the pilots were already the fastest-improving jurisdictions before the reform is the local, measurable trace of a general feature of Chinese governance: policy pilots are allocated to the strongest localities, not a representative sample (He 2024; Luo, Wang, and Yang 2025). The selectivity that complicates identification is itself evidence about how reform is staged across Chinese territory — the centre concentrates trials where capacity already exists, so

the geography of a pilot reform tracks, and tends to widen, a pre-existing distribution of institutional quality.

A more familiar account of Chinese local-government behaviour does not fit here. The literature on yardstick competition among localities (Cai and Treisman 2005; Li and Zhou 2005) holds that officials compete for advancement by improving the indicators they control; but the CPSR pilots were chosen in Beijing, so the reform is not itself a product of local rivalry, and improving the business environment is something nearly every Chinese local government pursues, which makes it a poor candidate for explaining why this reform in particular tracks the pattern observed here. The pattern is better read through the efficiency of dispute resolution and its bearing on where activity settles.

The closest precedent is recent work on how Chinese judicial reform reshapes economic activity. Liu, Lu, Peng, and Wang (forthcoming) exploit the staggered post-2014 reform that moved control over court finances and personnel from local to provincial governments, and find that weakening local protectionism drew non-local investment into reformed places and knit regional markets together; related work links improvements in judicial independence to entrepreneurship and corporate innovation (Lai et al. 2023; Yang, Wang, and Lai 2023). This literature has established that the quality of courts shapes where economic activity goes, but it has examined that almost entirely through the lens of integration — reform reducing the frictions that segment a national market. It leaves open the question studied here: what a reform itself rolled out unevenly across space does to the distribution of institutional quality while the rollout lasts. The 2014 reform altered who controls the courts; the 2020 reform altered the procedure they run, first in twenty named cities. That difference in design is what brings the spatial question to the surface.

Two contributions follow. The first is empirical: a city-level spatial analysis of a Chinese civil-procedure reform that treats the selectivity of pilot assignment as both identification problem and finding, recovers a legal-dimension gain by synthetic control for the pilots that admit a counterfactual, and shows the gain is not won by drawing activity from neighbouring cities. The second is conceptual: the transitional geography of staged reform — reformed places pulling ahead while the reform is theirs alone, on a distribution already tilted toward them by the logic of pilot selection — is a general property of how institutional modernisation travels across space when it arrives unevenly, not a Chinese peculiarity. Section 7 draws that comparison out. Before that, Section 2 lays out the court system and the reform; Section 3 develops the interpretation; Sections 4 and 5 present the data, specification, and results; and Section 6 weighs the interpretation against its limits.

2. Institutional Background

2.1 The Chinese court system and civil jurisdiction

The Chinese court system has four tiers: the Supreme People's Court (SPC) at the apex, with supervisory and interpretive authority over all lower courts (Constitution, Articles 128–133; Organic Law of the People's Courts 2018), then the provincial Higher, prefectural Intermediate, and county-level Basic People's Courts. The “prefecture level” used throughout is the tier where the Intermediate People's Court sits and where most commercial disputes of moderate value are first heard or heard on appeal. Specialised courts — the Intellectual Property Courts, the Shanghai Financial Court, the Internet Courts, and the maritime courts — complement the territorial structure with subject-matter jurisdiction (Chen 2020; Finder 2022).

Civil cases are dispatched among courts under the jurisdiction rules of the Civil Procedure Law. An action against a legal person is heard, by default, at the defendant's domicile — normally its registered seat (Articles 22–23) — though a contract dispute may also be brought at the place of performance and a tort action at the place of the wrong (Articles 24, 28), and within a limited list of permitted fora contracting parties may stipulate jurisdiction by agreement (Article 35). These rules generate flows of cases that follow the geography of economic activity, so the procedural quality of a prefecture's courts is one of the considerations bearing on where a firm registers and how it drafts its jurisdiction clauses.

2.2 Civil procedure before the reform

The 2017 Civil Procedure Law provided three first-instance tracks at the Basic People's Court: the ordinary procedure (Articles 134–146), the default, with a three-judge panel and a six-month window; the summary procedure (Articles 157–163), for simple cases with clear facts, using a single judge and a three-month window; and the small-claim procedure, a single-judge, one-trial variant for monetary claims below 30% of the average annual local wage.

By the late 2010s the courts faced a documented mismatch between caseload and judicial supply: the annual civil caseload roughly doubled between 2013 and 2019 while the post-2014 quota reform capped the judicial corps (Finder 2017; Liebman, Roberts, Stern, and Wang 2020). The SPC's Five-Year Reform Plan (2019–2023) accordingly made the procedural separation of complex and simple cases a priority — the context in which the CPSR was designed.

2.3 The Civil Procedure Case-Separation Reform

The CPSR was authorised by the NPCSC's Decision of 28 December 2019, which let the SPC deviate from designated provisions of the Civil Procedure Law for a two-year pilot; the SPC's Implementing Measures of 15 January 2020 (Fa [2020] No. 11) set the operational rules. Five components are central, each an efficiency measure.

Small-claim procedure. The monetary ceiling was raised from 30% to 50% of the average annual local wage; cases under it are tried by a single judge in one trial, with limited written-judgment requirements and no appeal on the merits.

Summary procedure. Its scope was clarified and broadened, with explicit criteria for assigning cases to the lighter track.

Single-judge second-instance review. Where the facts were established at first instance, the Intermediate People's Court could hear the appeal with a single judge rather than the usual three.

Online litigation. The Measures gave a legal basis for online filing, evidence submission, videoconference hearings, and online delivery of judgments.

Judicial confirmation of mediated agreements. The expedited procedure for confirming mediated settlements was expanded, providing a low-cost route to an enforceable outcome.

The pilot ran from 15 January 2020 through the end of 2021, and the SPC's mid-term report to the NPCSC documented substantial first-year gains in disposition speed, pre-litigation mediation, and the use of judicial confirmation across the pilot courts (Supreme People's Court 2021). On 24 December 2021, on the SPC's recommendation and following NPCSC review, the principal pilot components were incorporated into a national amendment of the Civil Procedure Law, effective 1 January 2022. The 2022 observations therefore correspond to the first year of nationwide implementation rather than to the pilot phase, a point taken up in interpreting the later event-study coefficients.

2.4 Selection of pilot cities

The SPC's Measures designated twenty prefecture-level cities as pilot jurisdictions, plus the three specialised Intellectual Property Courts, the Shanghai Financial Court, and the three Internet Courts. The twenty pilot cities — on which the empirical analysis focuses — are Beijing, Shanghai, Nanjing, Suzhou, Hangzhou, Ningbo, Hefei, Fuzhou, Xiamen, Jinan, Zhengzhou, Luoyang, Wuhan, Guangzhou, Shenzhen, Chengdu, Guiyang, Kunming, Xi'an, and Yinchuan.

SPC explanatory documents give three selection criteria: existing commercial caseload and complexity; administrative and technological capacity for online filing and remote hearings; and balance across the eastern, central, and western macro-regions. Two aspects matter here. The cities did not put themselves forward — the choice was made centrally — so whatever spatial pattern emerges is not the residue of competition among them for pilot status. And the choice was far from random: when the sample opens, the twenty pilots already held a disproportionate share of the country's registered lawyers, law firms, and concluded commercial cases. The reform thus deepened an uneven distribution rather than disturbing a level one, and the endogeneity this implies is confronted directly in Section 5.6.

2.5 Comparative perspective

Differentiated case management is a recognisable family of civil-procedure reforms, and a brief survey locates the Chinese case and sets up Section 7. The Woolf reforms in England and Wales introduced a three-track system with a low-cost small-claims track, since extended online (Woolf 1996; Genn 2010; Sorabji 2014; HMCTS 2023); Singapore's Simplified Process (Rules of Court 2021) applies a managed, cost-capped procedure to mid-value disputes; the European Small Claims Procedure (Regulation 861/2007) and Order for Payment offer written, form-based routes for low-value claims (European Commission 2013 2022; Hess 2010); France's 2019 reform routed low-value consumer and tenancy disputes to a dedicated protection judge (Cadet 2020); and Germany's *Mahnverfahren* remains the reference instrument for expedited recovery of uncontested claims (Stürner 2007). CEPEJ reports that such reforms cut disposition times by roughly 10–30% for the targeted categories (CEPEJ 2020 2022) — a range that brackets the SPC's reported magnitudes.

The CPSR shares with these reforms the sorting of cases by complexity and the integration of digital infrastructure (Susskind 2019); it differs in being deployed first as a geographically concentrated pilot. That staged design is what makes the spatial question meaningful: where the comparators applied nationwide from the outset, the asymmetry studied here would not arise.

2.6 What changed in the spatial organisation of the courts

One question about a reform of this kind is whether it redrew the spatial organisation of the courts. It did not: the four-tier hierarchy, the territorial assignment of courts, and the jurisdiction rules of Articles 22–35 were untouched. The reform changed procedure within existing courts, so the spatial pattern documented here arises from a change in the relative procedural efficiency of courts that stayed in place.

3. Interpretive Framework

The reading offered here joins two bodies of work that rarely meet. One is the study of experimentation under hierarchy in China (Florini, Lai, and Tan 2012; Heilmann 2008), which traces how the state introduces change through geographically bounded pilots and scales up those it judges successful. Recent work refines this in ways that bear on the case: Luo, Wang, and Yang (2025) document a marked recentralisation toward Beijing since 2013, of which the CPSR — designed and directed by the Supreme People's Court — is an instance; and studies of pilot selection find it favours the most capable jurisdictions (He 2024), so the pilots were not a cross-section of the country but its strongest performers — a point the empirical design must confront. What this literature dwells on is the political logic of piloting; what it leaves aside is what the pilot phase does to localities while it lasts. The other body of work treats the map of economic activity as a product of the institutions that organise it, and finds that institutional change tends to widen spatial disparities before it narrows them. The reform examined here falls in the gap between them.

Civil procedure is one of the conditions under which firms do business. Quicker disposal of simple disputes, lower litigation costs, online filing, and a fast route to an enforceable settlement all cut what ordinary commercial disputes consume. When those conditions improve in some cities and not others, the reformed cities become, at the margin, better places to register a firm — registration largely fixes the default forum under Articles 22 and 23 — and to base the legal capacity that serves commercial clients. Whether the improvement then redistributes activity toward them or raises institutional quality in place is the empirical question this paper takes up.

The expectation that follows is directional but agnostic about its spatial form: the pilots should gain on the dimensions the reform touches — above all the legal environment — while whether that gain is mirrored by a relative decline in their neighbours or leaves them broadly unaffected is left for the data. The claim is bounded accordingly: the panel can show whether the reform is associated with an improvement in the pilots' business environment and with what spatial signature, not that efficiency, rather than something else, produced it.

4. Data and Empirical Strategy

4.1 Sample

The empirical analysis uses a balanced panel of 296 prefecture-level cities of mainland China, observed annually from 2017 to 2023. The pooled difference-in-differences estimates reported as a descriptive benchmark use 2018–2022 (296 cities x 5 years, 1,480 observations); the synthetic-control analysis that carries the identification (Section 5.6) uses the full 2017–2023 window, matching on the pre-reform years 2017–2019 and tracing the gap through the second post-rollout year. Cities are identified by their administrative code (Ministry of Civil Affairs); twenty are CPSR pilots (Section 2.4) and the remaining 276 are non-pilots. We refer to the non-pilots throughout as the comparison group rather than the control group: selection into the pilot was deliberate and patterned, the two groups are not exchangeable at baseline, and — as Section 5.1 shows — they were already diverging before the reform, which is why identification does not rest on a common-trend assumption between the two groups but on the within-pilot counterfactual that synthetic control supplies.

Figure 2 maps the twenty pilots. Their distribution matters for how the spatial spillover should be read: the pilots crowd the eastern coast — the Yangtze River Delta, the Pearl River Delta, the Bohai Rim — and thin out toward the interior and west, where Chengdu, Guiyang, Kunming, and Yinchuan stand largely alone. The non-pilot cities near a pilot are therefore an eastern set rather than a national one, and in the coastal clusters a pilot's nearest neighbours are often other pilots — which is why the spillover is estimated under several definitions of neighbourhood (Section 5.5).




Figure 2. The twenty CPSR pilot prefecture-level cities. Points mark the pilots designated under SPC Fa [2020] No. 11; provincial boundaries are shown for reference, and shaded ellipses mark the three coastal agglomerations in which several pilots adjoin one another — the Yangtze River Delta (Shanghai, Suzhou, Hangzhou, Ningbo, Nanjing), the Pearl River Delta (Guangzhou, Shenzhen), and the Bohai Rim (Beijing, Jinan). Locations only; coordinates are public city-centre longitude and latitude, boundaries provincial.

4.2 Variables

The dependent variable is the Business Environment Index (BEI) of the city, drawn from the City-level Business Environment Composite Index of the Peking University–Wuhan University Joint Task Force. The index aggregates six sub-indices — market environment, government services, legal environment, human capital and culture, financing, and innovation — on a 0–100 scale, using a hybrid methodology that integrates administrative data and enterprise surveys (Peking University and Wuhan University 2021). The analysis uses both the composite and, in the decomposition of Section 5.3, each of the six sub-indices in turn.

The index recommends itself on two counts: it is the measure empirical work on China’s local economic institutions has converged on, and its legal-environment sub-index is built from precisely what a procedural-efficiency reform should move — disposition time, litigation cost, enforcement reliability. That it resolves into named sub-indices makes the mechanism testable: the effect can be located in the dimension the reform should touch. Section 4.3 takes up how far a movement in the index reflects real change.

The policy variable is the standard interaction $DID_{it} = Treat_i \times Post_t$, where $Treat_i = 1$ if city i is in the pilot set and 0 otherwise, and $Post_t = 1$ for $t \geq 2020$ and 0 otherwise.

The identifying variation is within-city over time, net of common annual shocks: the two-way fixed-effects design (Section 4.4) absorbs all time-invariant differences between cities and all shocks common to a given year, so the policy term is identified from the differential movement of pilots and non-pilots after 2020. Table 1 lists the core variables.

Table 1. Core variables, measurement, and sources.

Variable	Symbol	Measurement at the city level	Source
Business Environment Composite Index	BEI	Annual composite on 0–100 scale; six sub-indices aggregated	Peking University–Wuhan University Joint Task Force (2021)
— six sub-indices		Market; government services; legal environment; human capital and culture; financing; innovation (each 0–100)	Peking University–Wuhan University Joint Task Force (2021)
Pilot $\times$ Post	DID	1 if city is a pilot AND year ≥ 2020 ; 0 otherwise	NPCSC (2019); SPC Fa [2020] No. 11

4.3 The Business Environment Index and what it measures

Because the dependent variable is a composite index, a fair question is whether a movement in it reflects a real change in the conditions of dispute resolution or the mechanical response of one component to a change in procedure. The index combines administrative records (including SPC statistics on disposition times and procedural costs) with enterprise-survey responses on firms' experience of dispute resolution. Movement in the administrative component alone would be partly mechanical; movement that also appears in the survey component reflects experience on the ground. The two cannot be fully separated here, so the index's response is best read as an upper bound on the substantive change. Constructing a narrower index from primary court-performance statistics would sharpen this, and is identified as a direction for further work.

4.4 Specification

The baseline difference-in-differences specification is a two-way fixed-effects regression:

$$BEI_{it} = \alpha_i + \varphi_t + \beta_1 DID_{it} + \varepsilon_{it}$$

where α_i is a city fixed effect, φ_t a year fixed effect, ε_{it} the idiosyncratic error clustered at the city level, and β_1 the coefficient of interest. The same equation is estimated with the composite index and with each of the six sub-indices as the dependent variable.

The fixed-effects design rests on city-specific characteristics correlated with both pilot status and the business environment — a city's legal tradition, administrative capacity, and pre-existing concentration of commercial litigation. When such characteristics are present and correlated with the regressors, fixed effects are preferred to random effects (Hausman 1978; Wooldridge 2010);

city and year fixed effects absorb time-invariant city heterogeneity and common annual shocks respectively, so β_1 is identified from the within-city change in pilots relative to non-pilots after 2020. Standard errors are clustered at the city level throughout.

This baseline ignores the spatial structure of the data. When the outcome is spatially autocorrelated and the policy variable has spatial structure, a non-spatial estimator can misattribute spatial variation to the policy term, and cannot speak to whether the pilots' gain came at their neighbours' expense (Anselin 1988; Gibbons and Overman 2012). Both motivate the spatial specification of Section 4.7.

4.5 Functional form

The index and its sub-indices enter the difference-in-differences specifications in levels, on their native 0–100 scale, so the coefficients in Table 3 are in index points and directly comparable across the composite and the six sub-indices. The spatial models of Section 5.5 use the log of the composite index, which is right-skewed across cities, so the spatial-autoregressive structure operates on a roughly symmetric variable; the sign of the spillover — positive throughout, never negative — is unchanged in levels.

4.6 Global Moran's I

The standard measure of global spatial autocorrelation is Moran's I (Moran 1950; Cliff and Ord 1981), defined for a variable x observed at N spatial units as:

$$I = (N)/(W) \cdot (\sum_i \sum_j w_{ij} (x_i - \bar{x})(x_j - \bar{x})) / (\sum_i (x_i - \bar{x})^2)$$

where w_{ij} is the spatial weight on the pair (i, j) , $W = \sum_i \sum_j w_{ij}$ the total weight, and $\bar{x}$ the cross-sectional mean. The statistic runs from about -1 (perfect dispersion) to $+1$ (perfect clustering); under the null of no spatial autocorrelation its expectation is $-1/(N-1)$, near zero for $N = 296$. Inference is by permutation test with 999 randomisations.

The variable is the (log) BEI of city i in year t , computed separately by year. The weights w_{ij} come from the geodesic distance between prefecture-polygon centroids, zeroed beyond 450 km, a band that leaves every city at least one neighbour (the largest nearest-neighbour distance in the sample is 370 km). We row-standardise, and as a robustness check re-estimate under wider distance bands (600 km and 800 km).

4.7 The spatial specification

To test whether the pilots' improvement is accompanied by displacement of their neighbours, the difference-in-differences equation is augmented with two spatial terms — the spatial lag of the outcome, which captures the clustering documented by Moran's I, and the spatial lag of the policy variable, which captures the effect of having pilot neighbours:

$$\log BEI_{it} = \rho \sum_j w_{ij} \log BEI_{jt} + \beta_1 DID_{it} + \theta_1 \sum_j w_{ij} DID_{jt} + \alpha_i + \varphi_t + \varepsilon_{it}$$

where ρ is the spatial autoregressive coefficient on the lagged dependent variable, θ_l the coefficient on the spatially-lagged policy term, and the remaining symbols follow the non-spatial specification. The model is estimated by maximum likelihood with city and year fixed effects (Elhorst 2014).

The coefficient θ_l is the one that bears on displacement: it captures the association between a city's (log) business environment and its having pilot neighbours. A negative θ_l would be the signature of a pilot drawing activity from the cities around it; a coefficient at or above zero would indicate no such drawing-down. It identifies a spatial association, not a causal effect (Gibbons and Overman 2012), and because it can depend on how neighbours are defined, Section 5.5 estimates it under nine alternative weight matrices.

4.8 Synthetic control

Because the pilots were already diverging before the reform (Section 5.1), the common-trend assumption fails and the difference-in-differences estimates serve only as a descriptive benchmark. The identification rests instead on the synthetic-control method (Abadie, Diamond, and Hainmueller 2010), which replaces the assumption of parallel trends with an explicitly constructed counterfactual. For each pilot i , the method selects non-negative weights w_j on the non-pilot cities j , summing to one, that minimise the distance between the pilot and the weighted donor pool over the pre-reform years 2017–2019:

$$\min_{w_j} \sum_{t=2017}^{2019} (BEI_{it} - \sum_j w_j BEI_{jt})^2$$

The resulting weighted average — the synthetic control — reproduces the pilot's pre-reform path, and the reform's effect is read off the gap between the pilot and its synthetic control after 2020. Because the match is to the pre-reform *trajectory* and not merely the pre-reform level, a pilot already on a rising path is matched to non-pilots on the same path, so the pre-existing divergence is absorbed into the counterfactual rather than attributed to the reform. Match quality is summarised by the pre-reform root-mean-square prediction error (RMSPE); pilots whose path cannot be reproduced ($\text{RMSPE} \geq 3$) are excluded from the causal estimate. Inference is by placebo: applying the same procedure to every non-pilot yields the distribution against which the pilots' gaps are judged (Section 5.6).

5. Results

5.1 Pre-reform divergence and the limits of difference-in-differences

A difference-in-differences design identifies a causal effect only if the pilots and non-pilots would have moved in parallel absent the reform. Figure 1 puts this assumption to the test through an event study, interacting pilot status with year indicators for 2018 through 2022, with 2017 — the earliest year in the panel — as the reference, so that both pre-reform years (2018 and 2019) can be examined for divergence. The two panels show the composite index and the legal-environment sub-index.

Event-study estimates (2017 reference): pilots diverge before the reform




Figure 1. Event-study estimates for the composite index (panel a) and the legal-environment sub-index (panel b), with 2017 as the reference year. Points are coefficients on Pilot $\times$ Year from a two-way fixed-effects regression; bars are 95% confidence intervals (city-clustered). The shaded region marks the pre-reform years; the dotted line marks 2020.

The pre-reform coefficients do not support a common trend. On the composite index (panel a) the 2018 coefficient is +2.35 and significant at the 1% level, indicating that the pilots were already pulling ahead of the non-pilots a full two years before the reform. The legal sub-index (panel b) is more troubling still: the pilots sit +2.25 above the non-pilots in 2018 and +5.26 above them in 2019, both significant, so the legal environment of the pilot cities was on a steepening upward path well before the CPSR began. This is the opposite of what the design needs. Rather than a clean break at 2020, the event study shows a pre-existing divergence that the reform date does not interrupt.

The pattern is a substantive feature of how the pilots were chosen: selected centrally for capacity and caseload, they were among the strongest, most rapidly improving jurisdictions before the reform, in keeping with the finding that Chinese pilots are allocated to the most capable localities (He 2024; Luo, Wang, and Yang 2025). The pre-trend in Figure 1 is the empirical signature of that selection: the pilots were on a faster trajectory, and a naïve difference-in-differences would credit the reform with improvement that was already under way. We therefore do not treat the pooled two-way fixed-effects estimates as causal. They are reported below as a descriptive benchmark, after which Section 5.6 turns to a synthetic-control design that confronts the selection problem directly by matching each pilot's pre-reform path.

5.2 Spatial autocorrelation

Table 2 reports the Global Moran's I of the (log) composite index for each year. It is positive and significant at the 1% level throughout, rising from 0.209 in 2018 to 0.313 in 2022 — positive and strengthening spatial clustering of the business environment across cities. Section 5.5 takes up what this implies for reading the policy effect; here it establishes that a spatial specification is

warranted, since with autocorrelation of this magnitude a non-spatial estimator would misattribute part of the spatial variation to the policy term.

Table 2. Global Moran’s I of the (log) Business Environment Index, 296 cities, 2018–2022.

Year	Global Moran’s I	p-value
2018	0.209***	< 0.01
2019	0.209***	< 0.01
2020	0.214***	< 0.01
2021	0.333***	< 0.01
2022	0.313***	< 0.01

Notes: Distance-based row-standardised spatial weights, 450 km band (Section 4.6). *p*-values from permutation tests with 999 randomisations. *** denotes significance at the 1% level. Results under wider distance bands (600 km and 800 km) are qualitatively identical, with Moran’s I positive and significant at the 1% level in every year.

5.3 Descriptive estimates and the channel of association

Table 3 reports the two-way fixed-effects estimates. Following Section 5.1, these are read as descriptive associations rather than causal effects: the pilots’ pre-reform divergence means the coefficients absorb both the reform and the pre-existing trajectory that selected the pilots. What they establish is not the magnitude of the reform’s effect but the *pattern* of the post-2020 divergence — which dimensions of the business environment the pilots pulled ahead on — a pattern the synthetic-control analysis of Section 5.6 then re-examines under a credible counterfactual.

On the composite index the coefficient is 2.232 (s.e. 0.821). The informative content is in the decomposition beneath it. The divergence is not spread evenly across the six sub-indices but concentrated in two. The legal-environment sub-index, the dimension the reform most directly touches, rises by 4.955 (s.e. 1.707), more than twice the composite figure; the innovation sub-index rises by 6.192 (s.e. 1.946); the human-capital-and-culture sub-index by a weaker 1.990; and the market-environment, government-services, and financing sub-indices show no significant movement. Whatever the pilots’ post-2020 advance reflects, it is located in the legal and innovation dimensions and not in the business environment at large.

That the divergence is concentrated in the legal dimension is consistent with the procedural reading of Section 3: the CPSR changed civil procedure, and the legal-environment sub-index is where such a change would register. The innovation dimension is plausibly connected through the same channel — improvements in judicial quality and the protection of legal rights have been linked to corporate innovation and entrepreneurial activity (Lai et al. 2023; Yang, Wang, and Lai 2023), since faster, more predictable enforcement lowers the risk attached to the contracts and licensing that innovative firms depend on. But both dimensions also diverge before the reform — the legal sub-index on the steepening path of Figure 1b, the innovation sub-index in a reversal of its own, with the pilots significantly below the non-pilots in 2018 before overtaking them — so

the descriptive coefficients cannot separate the reform's contribution from the selection trajectory. Establishing whether any of this divergence is attributable to the reform requires the counterfactual that Section 5.6 supplies; the decomposition here only narrows the question to the legal and innovation channels.

Table 3. Two-way fixed-effects estimates, business-environment index and its six sub-indices, 2018–2022. City and year fixed effects; standard errors clustered by city. Read as descriptive associations, not causal effects (Section 5.1).

Dependent variable	Coefficient	Std. error
Composite index	2.232***	(0.821)
<i>Sub-indices:</i>		
Market environment	0.539	(1.039)
Government services	1.799	(1.183)
Legal environment	4.955***	(1.707)
Human capital and culture	1.990**	(0.900)
Financing	−0.766	(1.236)
Innovation	6.192***	(1.946)

Notes: Each row is a separate regression of the named index on the Pilot $\times$ Post term (Post = 1 from 2020), with city and year fixed effects; $N = 1,480$ city-year observations (296 cities $\times$ 5 years), of which 20 cities are pilots. Standard errors clustered at the city level in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Following Section 5.1, coefficients are descriptive, not causal.

5.4 A descriptive view of the divergence

Figure 3 shows the year-by-year divergence behind the regression estimates, plotting the group-mean index for the twenty pilots and the 276 non-pilots, each indexed to its own 2019 value. The divergence is visible but, consistent with Section 5.1, it begins before the reform: the pilots are already separating from the non-pilots in the pre-reform years rather than tracking them to 2020 and breaking away after. The figure is included to display the raw pattern the formal analysis must contend with — the gap the synthetic control of Section 5.6 reproduces in its pre-reform years and nets out — not as evidence of a clean post-reform break.




Figure 3. Business-environment index by group, pilot versus non-pilot cities, 2018–2022, each series indexed to its 2019 value (2019 = 100). Panel (a): composite index; panel (b): legal-environment sub-index; dotted line marks 2020. Group means are unweighted. Source: City-level Business Environment Composite Index, Peking University–Wuhan University Joint Task Force.

5.5 The spatial structure: clustering without displacement

A distinct question is whether the pilots’ gain came at their neighbours’ expense. If the reform drew mobile firms and legal services across the pilot boundary, a non-pilot city’s trajectory would be depressed by proximity to a pilot. The strong, rising Moran’s I of Section 5.2 makes the question pressing: the business environment is spatially clustered, so pilots and their neighbours are not independent observations, and a displacement effect would show up as a negative association between a city’s outcome and the presence of pilot neighbours.

To test this, the specification is augmented with the spatial lag of the policy term, $W \cdot DID$, whose coefficient measures the association between a city’s (log) business environment and its having pilot neighbours. The estimate gives no sign of displacement: under every weighting scheme the coefficient is positive and, in most, significant. To confirm this is not an artefact of one definition of neighbourhood — the robustness the spatial-econometrics literature insists on (Gibbons and Overman 2012; LeSage and Pace 2009) — the model is re-estimated under nine schemes: distance bands at four thresholds, k -nearest-neighbour matrices at four values of k , and an inverse-distance matrix. Table 4 collects the spillover across all nine. It is positive throughout, ranging from +0.08 under the nearest-neighbour matrices to larger values under the wider distance bands, and never negative.

The sign is the substantive result: there is no evidence that the pilots advanced by drawing their neighbours down. The positive coefficients are most conservatively read as the absence of displacement rather than as a precisely estimated beneficial spillover, since their magnitude rises with the breadth of the weight matrix in a way that partly reflects the underlying spatial autocorrelation of the index rather than a causal transmission from pilot to neighbour; the

spatially-lagged outcome is itself strongly positive throughout. What the table rules out — robustly, across every specification — is the negative spillover that a displacement story requires. The spatial inequality the reform is associated with is one in which the pilots pull ahead, not one in which the places next door fall behind.

Table 4. Spatial spillover coefficient ($W \cdot DID$) under alternative spatial weight matrices. Spatial-lag-of-X panel with two-way fixed effects, (log) composite index as dependent variable.

Spatial weight matrix	$W \cdot DID$
Distance band, 450 km	0.318***
Distance band, 600 km	0.596***
Distance band, 800 km	0.732***
Distance band, 1000 km	0.970***
k -nearest neighbours, $k = 5$	0.079**
k -nearest neighbours, $k = 8$	0.132***
k -nearest neighbours, $k = 10$	0.166***
k -nearest neighbours, $k = 15$	0.236***
Inverse distance (≤ 800 km)	0.448***

Notes: Each row is a separate estimation under the named row-standardised weight matrix; $N = 1,480$. The coefficient on $W \cdot DID$ is positive in every specification and never negative — the result that bears on displacement. The magnitude rises with the breadth of the weight matrix, partly reflecting the spatial autocorrelation of the index (Section 5.2), and is not interpreted as a precise causal spillover. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.6 Identification through synthetic control

The descriptive estimates of Section 5.3 cannot be read causally because the pilots were diverging from the non-pilots before the reform. To recover an effect from this setting requires a counterfactual that reproduces each pilot’s pre-reform trajectory rather than merely its pre-reform level — which is what the synthetic-control method supplies (Abadie, Diamond, and Hainmueller 2010). For each pilot it constructs a weighted average of non-pilot cities that tracks the pilot’s composite index across the pre-reform years 2017–2019, and then reads the reform’s effect off the gap between the pilot and this synthetic counterpart after 2020. Because the synthetic control is built to match the pre-reform path, a pilot that was already on a steepening trajectory is matched to non-pilots on the same trajectory, so the pre-existing divergence is absorbed into the counterfactual instead of being misattributed to the reform.

The method works for the pilots whose pre-reform path can be reproduced from the donor pool. For fifteen of the twenty — all except Beijing, Shanghai, Guangzhou, Shenzhen, and Chengdu, whose pre-reform levels are too extreme for any combination of non-pilots to match — the pre-reform fit is close (root-mean-square prediction error below 3). For these fifteen the synthetic control reduces the raw pre-reform gap of roughly ten index points to 0.15 on average: the pre-trend that defeated the difference-in-differences design is, for this group, essentially eliminated. Figure 4 shows the year-by-year gap: through 2017–2019 it sits at zero, confirming the match,

and from 2020 it opens to an average of +5.85 index points — the effect that remains after the pilots’ own trajectory is netted out.




Figure 4. Synthetic-control gap on the composite index for the fifteen well-matched pilots, by year. Bold line: mean pilot–synthetic gap; grey lines: individual pilots; dashed line: raw unmatched gap, for contrast. The synthetic control holds the pre-reform gap (2017–2019) at approximately zero and isolates a post-reform gap that emerges only from 2020.

To gauge whether a post-reform gap of this size could arise by chance, the same procedure is run as a placebo on every non-pilot city, building a synthetic control for it from the other non-pilots. The resulting placebo gaps centre on +1.0 index points. The pilots’ mean gap of +5.85 sits firmly in the right tail of this placebo distribution, and the difference is significant on a rank test (Mann–Whitney $p < 0.001$): the post-reform divergence of the pilots is larger than the method recovers for cities that received no reform. Two features of the result temper it, and both are reported rather than set aside. The synthetic control speaks only for the fifteen matchable pilots, not the five strongest, whose exclusion is itself a consequence of their extreme selection — the cities most likely to have benefited are the ones for which no credible counterfactual exists. And the method controls for selection on the pre-reform trajectory, not for an unobserved shock that coincided with the reform; that residual possibility is taken up in Section 6.

5.7 The post-rollout path

If the pilots’ post-reform gain reflects, even in part, the reform rather than their continuing trajectory, its behaviour after the reform became national law bears on the interpretation. The CPSR’s main provisions entered the Civil Procedure Law from 1 January 2022, so by 2023 the non-pilots had operated under the same rules for a full year. Extending the synthetic-control gap to 2023 asks whether the pilots’ advantage narrows once their procedural exclusivity ends.

It narrows but does not close. The mean synthetic gap for the fifteen matched pilots runs from +5.71 in 2020 and +4.29 in 2021, through the pilot-exclusive period, to +7.56 in 2022 and +5.49 in 2023, the second year of national operation (Figure 4). The gap is thus smaller in 2023 than at its 2022 peak, consistent with a partial erosion of the pilots' advantage as the reform generalises, but it remains well above both zero and the placebo benchmark. The standard event study shows a sharper convergence — the 2023 coefficient falls close to zero — but that specification is contaminated by the pre-reform divergence of Section 5.1; the synthetic-control gap is the more reliable guide. A single post-rollout year cannot distinguish a partial, ongoing convergence from a new and lower equilibrium gap, and we do not press the point further. What the post-rollout evidence supports is the weaker statement that the pilots' advantage began to narrow once the reform ceased to be theirs alone — consistent with, though not proof of, an advantage resting partly on exclusive early access.

6. Limitations and an Agenda for Identification

The analysis confronts its central threat — pilots selected for strength, already diverging before the reform — rather than assuming it away; the synthetic-control design of Section 5.6 is the response. But that design has limits of its own: the effect is recovered only for pilots that admit a credible counterfactual, a residual confounder remains, the mechanism beneath the index is not directly observed, and the post-rollout evidence is a single year. Each is set out below with the design that would address it.

6.1 The effect is identified only for the matchable pilots

Synthetic control recovers an effect for fifteen of the twenty pilots — those whose pre-reform path a combination of non-pilots can reproduce. The five it cannot match are Beijing, Shanghai, Guangzhou, Shenzhen, and Chengdu, and their exclusion is not incidental: they are among the strongest cities in the country, so far above the donor pool that no weighted average of non-pilots tracks their pre-reform trajectory. The result is an identification gap precisely where the reform's effect might be largest. This is the selection problem in its most concrete form — the cities most transformed by being chosen are the ones that admit no counterfactual — and only a credible instrument for pilot status would close it; we do not have one.

6.2 A residual confounder the pre-reform match cannot rule out

Synthetic control absorbs selection on the pre-reform trajectory, but not a shock that struck the pilots and their synthetic controls differently after 2020. The window overlaps the COVID-19 pandemic, whose disruption was spatially patterned and fell heavily on the same eastern urban belt that holds many of the pilots. Because the synthetic control is built from non-pilots — many of them also eastern, urban, and pandemic-exposed — it absorbs the component of the shock common to both, but a differential effect that tracked pilot status specifically could still be read as a reform effect. Two features make a pure pandemic story unlikely: the gain is concentrated in the legal and innovation dimensions (Section 5.3) rather than those most exposed to disruption,

and the spillover evidence (Section 5.5) rules out the displacement it would predict. The window nonetheless remains a caution on the magnitudes.

6.3 The mechanism is not directly observed

The interpretation of Section 3 is that quicker, more transparent dispute resolution raised the quality of the pilots' legal environment, with effects reaching the activity that depends on reliable enforcement. The decomposition of Section 5.3 locates the post-reform divergence in the legal and innovation dimensions, which is consistent with this reading, but the index cannot trace the behaviour beneath it — whether the gain reflects firms relocating, lawyers and law firms concentrating, or commercial caseloads shifting toward the reformed courts. If the legal-environment gain were accompanied by such movement, it would show up in city-level series the index does not resolve: counts of newly registered firms, registered lawyers and law firms, and concluded commercial cases, available in principle from market-regulation, justice-ministry, and judicial-yearbook sources. Treating each as an outcome would show whether real activity moved or whether the legal environment improved in place — a distinction the present data cannot adjudicate.

6.4 The post-rollout evidence is a single year

The post-rollout narrowing of the pilots' advantage (Section 5.7) rests on a single year, 2023, the first full year under national law. One year cannot distinguish an ongoing convergence from a new, lower equilibrium gap; the narrowing could continue, halt, or reverse. Extending the panel several years past 2022 and tracing the synthetic-control gap year by year would settle whether the advantage was genuinely a transitional premium tied to exclusivity or a more durable lead that the reform merely accelerated — the last of the extensions this study leaves open.

7. Conclusion and Comparative Relevance

China's 2020 Civil Procedure Case-Separation Reform left a spatial trace, but reading it required first reckoning with how the pilots were chosen. Selected centrally for their strength, the twenty cities were already diverging from the rest of the country — most sharply in the legal environment — before the reform began, so the question was never simply whether the pilots improved but whether any of their improvement could be attributed to the reform rather than to the trajectory that selected them. Matching each pilot to a synthetic counterpart, we find that for the fifteen pilots that admit a credible counterfactual a post-reform gain remains after the pre-existing divergence is netted out, concentrated in the legal dimension and significantly larger than the same method recovers for unreformed cities. That gain is not won at the neighbours' expense: across every spatial specification the spillover of pilot status is never negative. After the reform becomes national law the advantage begins to narrow, consistent with a premium resting partly on exclusive early access, though a single post-rollout year cannot establish that it closes.

Two readings of the case reach beyond it. The first concerns method: where pilots are chosen for capacity, the pre-reform divergence is not noise but the central fact, and confronting it — through synthetic control, and the candid admission that the strongest pilots admit no counterfactual — is what separates a credible estimate from one that credits the reform with improvement already under way. The selectivity that complicates identification is itself a finding: the geography of a Chinese pilot reform tracks, and tends to widen, a distribution of institutional quality already tilted toward the strongest places (He 2024; Luo, Wang, and Yang 2025), because the centre concentrates trials where capacity exists.

The second reading is comparative. Differentiated case management of the kind the CPSR implements is not unique to China: the reforms of England and Wales, Singapore, the European Union, France, and Germany (Section 2.5) share its aim of fitting procedure to the complexity of the dispute, and several arrived through staged rollouts. Setting the Chinese case beside them isolates geographical staging as the feature that generates a transitional spatial inequality — reformed places pulling ahead while the reform is theirs alone — which might be called a *staged-reform spatial premium*. The evidence here suggests this premium is better understood as a temporary, localised lift than a zero-sum transfer from the neighbours, the spatial counterpart to the market integration that reforms reaching everywhere at once produce (Liu et al. forthcoming). The lesson for sequencing follows: where pilots are concentrated in already-strong places, a staged rollout widens an existing spatial inequality for as long as it lasts, so the gains from a longer trial must be weighed against the distributional cost of the gap it sustains — and against the possibility, which only longer data will resolve, that the gap outlives the exclusivity that opened it.

References

1. Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, 105(490), 493–505. doi.org/10.1198/jasa.2009.ap08746
2. Anselin, L. (1988). *Spatial econometrics: Methods and models*. Kluwer Academic.
3. Cadiet, L. (2020). Réforme de la justice et organisation judiciaire. *Justices*, 18, 17–35.
4. Cai, H., & Treisman, D. (2005). Does competition for capital discipline governments? Decentralization, globalization, and public policy. *American Economic Review*, 95(3), 817–830. doi.org/10.1257/0002828054201306
5. CEPEJ. (2020). *European judicial systems — CEPEJ evaluation report, 2020 evaluation cycle (2018 data)*. Council of Europe Publishing.
6. CEPEJ. (2022). *European judicial systems — CEPEJ evaluation report, 2022 evaluation cycle (2020 data)*. Council of Europe Publishing.
7. Chen, A. H. Y. (2020). *An introduction to the legal system of the People's Republic of China* (5th ed.). LexisNexis.
8. Cliff, A. D., & Ord, J. K. (1981). *Spatial processes: Models and applications*. Pion.
9. Elhorst, J. P. (2014). *Spatial econometrics: From cross-sectional data to spatial panels*. Springer.
10. European Commission. (2013). *Report on the application of Regulation (EC) No 861/2007 establishing a European*

Small Claims Procedure. COM(2013) 795 final. 11. European Commission. (2022). *Evaluation of the European Small Claims Procedure*. SWD(2022) 184. 12. Finder, S. (2017). The Supreme People's Court and the Chinese judicial reform agenda. *Supreme People's Court Monitor*. 13. Finder, S. (2022). The Supreme People's Court of the People's Republic of China. In M. Sellers & S. Kirste (Eds.), *Encyclopedia of the philosophy of law and social philosophy*. Springer. doi.org/10.1007/978-94-007-6730-0_1023-1 14. Florini, A. M., Lai, H., & Tan, Y. (2012). *China experiments: From local innovations to national reform*. Brookings Institution Press. 15. Genn, H. (2010). *Judging civil justice* (Hamlyn Lectures). Cambridge University Press. 16. Gibbons, S., & Overman, H. G. (2012). Mostly pointless spatial econometrics? *Journal of Regional Science*, 52(2), 172–191. doi.org/10.1111/j.1467-9787.2012.00760.x 17. Hausman, J. A. (1978). Specification tests in econometrics. *Econometrica*, 46(6), 1251–1271. doi.org/10.2307/1913827 18. He, A. J. (2024). Unpack the black box of pilot sampling in policy experimentation: A qualitative comparative analysis of China's public hospital reform. *Governance*, 37(3), 803–824. doi.org/10.1111/gove.12804 19. Heilmann, S. (2008). Policy experimentation in China's economic rise. *Studies in Comparative International Development*, 43(1), 1–26. doi.org/10.1007/s12116-007-9014-4 20. Hess, B. (2010). *Europäisches Zivilprozessrecht*. C. F. Müller. 21. HMCTS. (2023). *HM Courts and Tribunals Service annual report and accounts 2022–23*. Her Majesty's Courts and Tribunals Service. 22. Lai, S., Yang, L., Wang, Q., & Anderson, H. D. (2023). Judicial independence and corporate innovation: Evidence from the establishment of circuit courts. *Journal of Corporate Finance*, 80, 102421. doi.org/10.1016/j.jcorpfin.2023.102421 23. LeSage, J. P., & Pace, R. K. (2009). *Introduction to spatial econometrics*. Chapman and Hall/CRC. 24. Li, H., & Zhou, L.-A. (2005). Political turnover and economic performance: The incentive role of personnel control in China. *Journal of Public Economics*, 89(9–10), 1743–1762. doi.org/10.1016/j.jpubeco.2004.06.009 25. Liebman, B. L., Roberts, M. E., Stern, R. E., & Wang, A. Z. (2020). Mass digitization of Chinese court decisions. *Journal of Law and Courts*, 8(2), 177–201. doi.org/10.1086/707855 26. Liu, E., Lu, Y., Peng, W., & Wang, S. (forthcoming). Court capture, local protectionism, and economic integration: Evidence from China. *Review of Economics and Statistics*. 27. Luo, K., Wang, S., & Yang, D. Y. (2025). *Laboratories of autocracy: Landscape of central–local dynamics in China's policy universe* (NBER Working Paper). National Bureau of Economic Research. 28. Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, 37(1/2), 17–23. doi.org/10.2307/2332142 29. NPCSC. (2019). *Decision authorizing the Supreme People's Court to carry out the pilot programme of the reform of separation of cases into simple and complicated ones under civil procedure*. Standing Committee of the National People's Congress. 30. Peking University–Wuhan University Business Environment Research Joint Task Force. (2021). *China city business environment research report (2020)*. Peking University. 31. Sorabji, J. (2014). *English civil justice after the Woolf and Jackson reforms: A critical analysis*. Cambridge University Press. 32. Stürner, R. (2007). The principles of transnational civil procedure. *Rechts Zeitschrift*, 69(2), 201–253. 33. Supreme People's Court of the P.R.C. (2018–2022). *Annual judicial statistical bulletin*. SPC. 34. Supreme People's Court of the P.R.C. (2020).

Measures for the implementation of the pilot programme of the reform of separation between complicated cases and simple ones in civil procedure (Fa [2020] No. 11). SPC. 35. Supreme People's Court. (2021). *Report on the pilot reform of the diversion of civil procedure between simple and complex cases* [Mid-term report to the Standing Committee of the National People's Congress, 27 February 2021]. Supreme People's Court of the People's Republic of China. 36. Susskind, R. (2019). *Online courts and the future of justice*. Oxford University Press. 37. Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data* (2nd ed.). MIT Press. 38. Woolf, H. K. (1996). *Access to justice: Final report to the Lord Chancellor on the civil justice system in England and Wales*. HMSO. 39. Yang, L., Wang, Q. S., & Lai, S. (2023). Fair and square: The impact of judicial independence on entrepreneurship. *Finance Research Letters*, 55(B), 103992. doi.org/10.1016/j.frl.2023.103992

Data Availability Statement

All data and code required to reproduce every result reported in this article are supplied to the editors and referees as supplementary files. The materials comprise: (i) the raw city-level Business Environment Index panel for 296 prefecture-level cities, 2017-2023, as obtained from the Peking University-Wuhan University Joint Task Force; (ii) the cleaned analysis panel derived from it, together with the pilot-city indicator coded from the NPCSC Decision of 28 December 2019 and SPC Fa [2020] No. 11; (iii) the city-centre coordinates used to construct the spatial weight matrices, compiled from public sources; (iv) a codebook giving variable definitions, coding rules, and every transformation applied; and (v) the complete analysis scripts, which reproduce all tables, figures, and reported statistics from the raw file in a single run. The Business Environment Index is a third-party dataset compiled by the Joint Task Force; onward public redistribution is subject to the compiler's terms, and the file is therefore supplied for the purposes of editorial assessment and peer review. All legal instruments cited are publicly available, and the Supreme People's Court Annual Judicial Statistical Bulletins are at www.court.gov.cn.

Ethical Approval

This article does not contain any studies with human participants performed by any of the authors.

Informed Consent

This article does not contain any studies with human participants performed by any of the authors.

Conflicts of Interest

The authors declare no conflicts of interest.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [replicationpackage.zip](#)