

Do ownership reform pilots accelerate academic technology transfer? Evidence from China's job-related invention scheme

Xinzhe Li

zzzhe6@163.com

Nanjing University of Science and Technology

Zhaoqiang Zhong

Nanjing University of Science and Technology

Research Article

Keywords: Academic entrepreneurship, Technology transfer, Property rights, Professor's privilege, Patent policy, China

Posted Date: August 25th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10589040/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Should academic inventions belong to the university or to the scientist? Every large economy that has answered this question has moved toward institutional ownership, and the evidence behind that consensus comes almost entirely from countries that took rights away from individual scientists. China has run the opposite experiment. Between 2016 and 2020 it designated selected universities and research institutes to assign scientists ownership shares or long-term use rights in inventions made in public employment, the largest reallocation of academic invention rights yet attempted. We evaluate it on the full population of Chinese patents, linking 1.58 million applications from 134 academic institutions to a roster of 586,317 inventors and 9.5 million candidate matches on firm-filed applications, and comparing 16 designated universities with 88 comparable universities that were not designated. Designation did not raise university-industry co-patenting. Three post-designation years of data exclude increases above two percent of the control mean, and the composition of filings is equally unmoved. What designation did shift is where academics patent. In China's most research-intensive universities, firm-side patenting by pre-designation academic inventors rose by about six percent relative to controls, a result that holds under stricter name matching, alternative windows and permutation inference, though it weakens against propensity-score-matched controls. Ownership reform moved individual behaviour at the margin while leaving the institutional machinery of technology transfer untouched. Paired with the Norwegian evidence that withdrawing such rights halves academic entrepreneurship, our findings show that ownership rules are not a portable lever. They work through the transfer infrastructure that surrounds them, and granting rights into a system built for institutional control does not reproduce the effects of taking them away from one built for individual control.

JEL codes O31, O32, O34, O38

1. Introduction

Who should own the inventions that scientists produce in public research organisations? Forty years of policy experimentation have produced two competing templates. The Bayh-Dole template vests title in the employing institution and relies on technology transfer offices to license inventions to firms (Mowery et al., 2001; Grimaldi et al., 2011). The professor's privilege template vests title in the individual scientist, trusting personal ownership to supply the incentive that bureaucratic licensing lacks. The empirical record on the choice between them is thin and asymmetric. Norway's abolition of the professor's privilege in 2003 reduced both start-up formation and patenting by university researchers by roughly half, which suggests that withdrawing rights from individuals is costly (Hvide and Jones, 2018). Whether the reverse operation, extending ownership to individuals inside an institution-owned system, produces symmetric gains is largely untested at scale. The question matters because several governments, China foremost among them, have committed substantial policy capital to that operation.

China offers the most ambitious test to date. Between 2016 and 2020 the country rolled out, in staggered fashion, the largest reallocation of academic invention ownership yet attempted. Sichuan

province began piloting mixed ownership of job-related inventions in 2016, splitting title between the institution and the inventing team. In May 2020 nine central ministries issued a national pilot scheme, and in October 2020 designated 40 universities, research institutes and hospitals in which scientists could receive ownership shares or use rights of at least ten years in their job-related achievements. The diagnosis behind the reform is familiar from the international literature. Chinese universities patent heavily but commercialise little, and the gap has been attributed to the absence of individual ownership stakes (Chen et al., 2016; Fisch et al., 2016). If that diagnosis is correct, designated institutions should show measurable increases in commercialisation activity relative to comparable institutions that were not designated.

We test that proposition on the full register of Chinese patent filings from the China National Intellectual Property Administration. We build institution-level panels for the 2020 national cohort, the Sichuan early adopters, and 88 non-designated universities from the same national tier, and we measure four families of outcomes. The first is patent output. The second is output composition, in particular the share of invention patents, the most examination-intensive category. The third is co-application between academic institutions and firms. The fourth, new to this setting at scale, is firm-side patenting by academic inventors, built from a roster of 586,317 inventors active in sample institutions before the reform whom we track across 9.5 million candidate name-by-patent matches on patents filed by firms. If the reform pushes scientists toward the market, their names should appear on company patents at a higher rate in designated institutions.

One feature of the setting fixes the interpretation. Designation confers eligibility rather than automatic transfer, since rights reach a scientist only after disclosure, institutional review and a recorded agreement, and completed assignments are not published at institution level. Our estimates are therefore intention-to-treat effects of designation, which is the policy variable governments actually control.

The results separate two channels sharply. Institutional collaboration did not respond. Co-application between designated universities and firms tracked the control group throughout, and the estimates are tight enough to exclude increases above two percent of the control mean, while the composition of filings is unchanged. Individual behaviour did respond. Firm-side patenting by pre-designation academic inventors rose by 0.32 percentage points against a base of 5.7 percent, about six percent in relative terms, concentrated entirely in the sixteen double first-class universities that anchor our main sample. The event study shows the shift arriving with designation and persisting, the estimate survives restricting matches to uncommon inventor names, alternative estimation windows and permutation inference, and it weakens only when the control pool is replaced by propensity-score-matched universities. The pattern is coherent. Designation changed what individual scientists do at the margin without changing what their institutions do.

Two features of the design make this reading credible. The outcome set spans both the institutional channel and the individual channel, so the contrast between them is internal to a single research design

rather than a comparison across studies. And the Sichuan cohort, designated in 2016, extends the observation window to seven post-designation years and shows the same absence of an institutional response.

The contribution is threefold. First, we provide the first large-sample causal evidence on extending ownership toward individual academic scientists, the operation that four decades of policy debate has argued about and that the professor's privilege literature has only ever observed in reverse. The asymmetry we document is the paper's central theoretical result. Removing individual rights in Norway halved academic entrepreneurship; granting them in China left institutional transfer untouched and moved individual patenting only at the margin. Ownership is not a lever that works the same way in both directions, and its effect depends on the transfer infrastructure it is inserted into. Second, we build a measure of firm-side patenting by academic inventors that runs over an entire national patent register, tracking half a million named scientists across nine and a half million candidate matches, and we show that it detects a policy response the institutional aggregates miss entirely. This measurement approach is portable to any patent system with inventor names. Third, we settle an empirical question that has been argued from anecdote in Chinese policy debate for a decade, by showing that the co-patenting channel through which reform was expected to work did not move, and quantifying how small any such movement could have been.

The setting also suits causal study. Ownership reform arrives as a dated, named-institution intervention rather than a continuous funding flow, and the Chinese pilots supply treatment lists fixed by ministerial document, a deep pool of never-designated institutions from the same administrative tier, and outcome data covering the universe of patenting rather than a survey sample. That combination lets us state what the data exclude in policy-relevant units, a step empirical work often skips (Abadie, 2020).

The remainder of the paper proceeds as follows. Section 2 describes the reform and develops the conceptual framework. Section 3 presents the data and measurement. Section 4 sets out the empirical strategy. Section 5 reports results. Section 6 discusses interpretation and limitations. Section 7 concludes.

2. Policy background and conceptual framework

2.1 The ownership question in academic technology transfer

Property rights theory holds that ownership of an asset should rest with the party whose non-contractible investment matters most for the asset's value (Grossman and Hart, 1986; Hart and Moore, 1990). Applied to academic inventions, the argument cuts both ways. The scientist supplies the tacit knowledge without which most early-stage inventions cannot be developed, which argues for individual ownership. The institution supplies laboratories, legal capacity and reputation, and a fragmented ownership landscape raises transaction costs for licensees, which argues for institutional ownership

(Kenney and Patton, 2009). The United States settled the question in favour of institutions in 1980, and the growth of university licensing after Bayh-Dole was long read as vindication, although later scholarship attributes much of that growth to secular trends already under way (Mowery et al., 2001; Thursby and Thursby, 2011). Recent work exploiting variation in who holds title continues to find that the allocation of ownership shapes commercialisation outcomes (Wang and Qian, 2023).

European experience pushed in the opposite direction. Between 2000 and 2007 Germany, Austria, Denmark, Finland and Norway abolished their professor's privilege regimes and transferred ownership to universities, hoping to replicate the American model. The evaluations that followed are sobering. Hvide and Jones (2018) document a decline of approximately fifty percent in both start-up formation and patenting by Norwegian university researchers after the 2003 abolition. Czarnitzki et al. (2016) find that the German reform reduced invention disclosures channelled through the new institutional system relative to what the privilege regime had produced. These studies identify the effect of removing individual ownership. They cannot, by construction, identify the effect of extending it within a system whose complementary institutions grew up around institutional ownership.

A second strand of literature suggests why extending rights might disappoint. Scientists engage with industry predominantly through consulting, contract research and collaborative projects rather than by licensing their patents, and these channels are governed by norms and relationships more than by title (Perkmann et al., 2013; Perkmann et al., 2021). Where royalty-sharing already gives inventors a substantial claim on licensing income, the marginal incentive effect of an ownership share may be small (Lach and Schankerman, 2008), and the institutional environment conditions whether any incentive translates into activity (Siegel and Wright, 2015; Fini et al., 2018; Åstebro et al., 2012). On this view ownership reallocation is neither necessary nor sufficient, and its effect depends on whether ownership is the binding constraint.

The Chinese case carries its own literature. University technology transfer in China grew up under strong state steering, limited transfer-office capacity and patent-count evaluation, and reviews emphasise that its bottlenecks have been organisational rather than legal (Chen et al., 2016). Industry-science linkages deepened after the reforms of the early 2000s (Motohashi and Yun, 2007), while academic patenting responded strongly to subsidies and evaluation incentives, so filing volumes overstate the stock of commercially valuable inventions (Hu and Jefferson, 2009; Dang and Motohashi, 2015; Fisch et al., 2016). Studying the provincial experiments, Liu et al. (2025) report that mixed ownership is associated with changes in the volume and channel of university commercialisation. Any evaluation here must separate whether more patents get filed, where subsidy dynamics dominate, from whether existing inventions move toward firms, where ownership incentives should operate.

Direct quasi-experimental evidence on the incentive margin remains scarce. Lach and Schankerman (2008) show that universities offering inventors larger royalty shares generate more licensing income, but royalty shares vary across institutions rather than being assigned, and the margin they study is income rather than control. Reviews of transfer-support policies conclude that credible causal evidence

is thin across the board (Kochenkova et al., 2016; Sandström et al., 2018). The Chinese pilots offer an opportunity to observe the extending operation at scale with a defined designation date and a comparable non-designated pool.

2.2 China's job-related invention reform

Job-related scientific achievements produced in Chinese public universities and research institutes are owned by the employing institution and are administered as state-owned assets. Before the reforms, commercialising an academic invention required approval chains designed for state asset management, and administrators who approved transactions later judged undervalued could face liability for loss of state assets. The 2015 revision of the Law on Promoting the Transformation of Scientific and Technological Achievements delegated disposal rights to institutions and mandated minimum inventor rewards, but title remained institutional and the liability regime remained intact. Patent output of Chinese universities nonetheless grew rapidly through this period, growth that, as reviewed above, owed more to subsidy and evaluation incentives than to commercial demand (Wei et al., 2017).

The ownership experiments began in Sichuan. In January 2016 Southwest Jiaotong University issued regulations granting inventing teams a seventy percent ownership share in job-related inventions, an initiative widely described in Chinese policy discourse as the science and technology counterpart of the household responsibility reforms. In December 2016 the province formalised a pilot of mixed ownership covering twenty universities and research institutes (Sichuan Department of Science and Technology and Sichuan Intellectual Property Office, 2016), extended to 45 institutions in 2018 and province-wide thereafter. In May 2020 the reform went national. Nine ministries jointly issued a pilot scheme for assigning ownership or long-term use rights in job-related achievements to research personnel (Ministry of Science and Technology et al., 2020a), and in October 2020 named 40 institutions (Ministry of Science and Technology, 2020b), among them seven universities administered by the Ministry of Education, four by the Ministry of Industry and Information Technology, four Chinese Academy of Sciences affiliates, and provincial universities, research academies and one hospital. The pilot term was three years. Twenty of the 40 belong to the double first-class tier that anchors our control group.

Two features of the national scheme matter for interpretation and are easy to overstate in the other direction. First, what the scheme offers is normally co-ownership rather than transfer of full title. Institutions may assign the inventing team an ownership share, with seventy percent shares common among early implementers following the Southwest Jiaotong template, or grant exclusive use rights of no less than ten years while retaining title. Both instruments move residual control toward the inventor, which is the margin property rights theory identifies as decisive for non-contractible development effort (Grossman and Hart, 1986; Hart and Moore, 1990), and both go beyond the income claim that the 2015 law already guaranteed. Neither amounts to the individual title that the professor's privilege conferred.

Second, the scheme is a package rather than an isolated ownership provision. Alongside rights assignment it directs pilot institutions to exercise autonomy in managing achievements as state assets, including discretion over whether formal valuation is required, to establish due-diligence exemptions that

protect administrators who act in good faith, to manage the process end to end, and to strengthen professional technology transfer capacity (Ministry of Science and Technology et al., 2020a). The reform therefore cannot identify the effect of ownership reallocation holding complements fixed. What designation identifies is the effect of entering the pilot package as institutions actually implemented it, and we frame our estimates accordingly.

Implementation follows a common sequence that is central to reading our results. A researcher discloses an achievement and applies for rights, the institution reviews and agrees a split or a use-rights term, the agreement is recorded where shares change, and only then can the researcher transact. Each step preserves an institutional veto, and neither the date of the first completed assignment nor the annual count of assignments is published at institution level, so we observe eligibility rather than take-up.

Figure 1 summarises the rollout and the composition of the designated group. Three features matter for the research design. The timing is staggered but concentrated, one early provincial cohort and one large national cohort, which allows each to be evaluated against never-designated institutions without the contaminated comparisons that arise when earlier-treated units serve as controls for later ones. The national list was published on 19 October 2020, so 2021 is the first application year fully exposed to designation, and we date treatment accordingly. And designation ran through recommending authorities, provincial science departments and central ministries, which favoured institutions with strong commercialisation records. From 2022 the model spread beyond the original pilots, with Zhejiang and Anhui launching provincial schemes (Zhejiang Provincial Department of Science and Technology et al., 2022), Anhui extending to 106 institutions and Hubei introducing a filing-based scheme in 2024 (Anhui Provincial Department of Science and Technology, 2024; Hubei Provincial Department of Science and Technology, 2024), Guangxi launching a management reform pilot in the same year (Guangxi Science and Technology Department et al., 2024), and Liaoning announcing full implementation in 2025 (Liaoning Provincial Department of Science and Technology, 2025), a diffusion that postdates our estimation window.

2.3 Empirical expectations

The framework yields a testable split. If ownership is the binding constraint on academic commercialisation, designation should raise both university-industry collaboration and the rate at which academics patent through firms. If the binding constraints lie downstream, in the capacity of transfer intermediaries, in what firms will pay, or in evaluation systems that reward publication over commercialisation, then designation should move neither. The informative case is the intermediate one. Control rights are held by individuals, not institutions, so a reform that binds only for the minority of scientists who want to commercialise personally would move the individual channel while leaving institutional collaboration flat. That is the pattern we find, and Section 2.2 identifies why it is plausible here, since Chinese inventors already held income rights and what designation added was control.

The mapping from these hypotheses to observable outcomes needs care in one respect. Co-application between a university and a firm is a visible trace of collaboration, but its sign under successful reform is ambiguous. If scientists take ownership shares and commercialise through vehicles of their own, patents may be filed by those vehicles without the university appearing as co-applicant, so the co-application share can fall while commercialisation rises. We therefore treat co-application as an indicator of collaboration intensity rather than as a directional test, and we rely on firm-side patenting by pre-reform academic inventors, which captures activity precisely when it moves outside the university's own filings, as the more diagnostic outcome. A response on that margin locates the effect at the level the reform targets, which is the individual scientist rather than the university.

3. Data and measurement

3.1 Patent data and institutional sample

Our source is the full register of the China National Intellectual Property Administration covering invention, utility model and design applications published from 1985 through November 2025. Records carry the application number and date, applicant names and types, applicant city, inventor names, main International Patent Classification code, grant details where applicable, and the current rights holder. Because the register is organised by publication year while the reform runs on application timing, we pool publication-year files from 2010 through 2025 and rebuild panels on application year.

Coverage differs by application year and we quantify it rather than relying on a nominal lag. Applications enter the register when they publish, invention applications at eighteen months or earlier on request and utility models on grant, so an application year is observed through a shrinking number of publication years as the extraction date approaches. Conditional on publication, 99.5 percent or more of invention applications and essentially all utility models from 2016 to 2024 appear within two calendar years of filing, and our extract runs to November 2025, covering nearly the full two-year window for the 2023 cohort and only part of it for 2024. Matched counts rise monotonically through 2023 and fall in 2024, the signature of censoring beginning in 2024. We end estimation in 2023, report a window ending in 2022 throughout, and rebuild the outcomes from applications published within two years of filing as a further check; all three agree. The audit also corrects an interpretation the raw series invites, since aggregate firm-side matches keep rising through 2023, any decline in a fixed-cohort measure after 2021 reflects the ageing of a cohort defined on pre-2020 activity rather than incomplete publication.

The designated group comprises the 40 institutions named in the October 2020 national list, verified against the original Ministry of Science and Technology document, and the Sichuan early adopters, analysed as a separate cohort. Four institutions appear in both programmes, Sichuan University, Southwest Jiaotong University, Chengdu University of Technology and Chengdu University of Traditional Chinese Medicine. To keep timing clean we assign these to the Sichuan cohort and exclude them from the national cohort, leaving 36 nationally designated institutions. The Sichuan cohort in our sample contains eight institutions rather than the twenty or forty-five covered by the provincial programme,

because the province published implementation documents without a complete institution list and we include only units verifiable individually from primary or contemporaneous official sources. The same gap means we cannot certify that every remaining control is untreated by the provincial scheme, which is why we drop the two Chengdu universities and why we treat the Sichuan estimates as descriptive rather than as a second identification. Appendix A lists every designated, excluded and control institution with its source.

Our main specification compares like with like on two dimensions. The designated group contains 25 universities, ten research institutes and one hospital, and no comparable population of never-designated elite research institutes exists, the Chinese Academy of Sciences having placed four affiliates in the pilot and circulated guidance to the remainder. Institution fixed effects absorb stable level differences across institution types but not differential trends among them. The same logic applies to institutional tier, since only 16 of the 25 designated universities belong to the double first-class tier from which every control is drawn. We therefore estimate the main results on those 16 designated universities against 88 non-designated double first-class universities, a sample matched on both type and tier, and report the nine other designated universities, the ten research institutes and the hospital as extensions.

The control pool begins as 90 double first-class universities not designated in any national or Sichuan pilot. Two of them, the University of Electronic Science and Technology of China and Southwestern University of Finance and Economics, are located in Chengdu and could plausibly have been covered by the Sichuan provincial programme, whose full institution list was never published. Because we cannot verify their non-designated status, we drop both, leaving 88 controls. Control institutions remain elite, patent-intensive and often located in the same cities as designated ones, which makes this a demanding comparison.

Contamination of the control group is a live concern. Control universities operate in the same cities and technology markets as designated ones, and some entered provincial schemes later. Within our window the provincial expansions that reached control-tier institutions began in late 2022 in two provinces and could affect filings from 2023. We identify one control university, Anhui University, that entered a provincial scheme in 2022, and we report estimates excluding it and estimates ending in 2022. We cannot rule out that other control institutions adopted comparable rules at their own initiative without a public designation, which would bias our estimates toward zero. We regard this as the most important unresolved threat to the design and return to it in Section 6.

Two data-handling choices deserve explanation. Applications can appear twice, once as a published application and again as a granted patent under the same number, so we deduplicate on application number and retain the grant record where both exist. And because the register carries no reliable licensing history, we do not use transfer-of-title measures. Comparing applicant with current holder could in principle proxy assignment, but reassignments accumulate over years, so the difference rate falls from six percent for 2015 filings to under one percent for 2023 filings in both groups alike, a censoring artifact that would contaminate any post-designation comparison. Name matching uses a hand-built dictionary

of institutional name variants covering former names, campuses, affiliated hospitals and asset management subsidiaries. The extract contains 1.58 million distinct applications after deduplication.

3.2 Outcome measures

We construct four outcome families at the institution-year level. Patent output is the count of applications. Composition is the share of invention patents in total applications. We treat this as a composition measure rather than a quality measure, since it records the filing category chosen rather than the value of the underlying invention, although the examination asymmetry between invention patents and utility models makes it informative about filing intent (Dang and Motohashi, 2015). University-industry co-application uses the applicant-type field, which codes each application's applicant mix. We classify an application as a co-application when its applicants include both a firm and an academic institution, and measure both the count and the share of an institution's applications so classified.

The fourth measure requires more construction. We first build a roster of every inventor named on any application of the 134 sample institutions with application years 2010 through 2019, recording for each name the set of four-digit IPC subclasses in which the person patented and the institution accounting for the most appearances. The roster contains 586,317 names. We then scan the entire register for applications whose applicants include no academic institution but do include at least one firm, and record appearances of roster names where the application's main IPC subclass belongs to the name's academic subclass set. The firm-applicant requirement matters because applications filed by individuals alone, which account for about five percent of candidate matches, need not represent commercial activity. The scan yields 9.5 million candidate name-by-patent matches under this restriction. The institution-year outcome is the number of distinct pre-2020 academic inventors from an institution appearing on at least one such patent in a given application year, scaled by the institution's roster size. We require at least twenty pre-designation inventors for inclusion, which drops one control university and no designated university.

We call this outcome firm-side patenting by pre-reform academic inventors rather than mobility, because the data cannot establish that anyone moved. An appearance is consistent with founding a company, joining one, consulting for one, or a continuing collaboration, and the measure aggregates these deliberately. Breadth suits our purpose, since the reform should encourage all of them, but it forecloses claims about labour mobility as such.

Name-based tracking of Chinese inventors is error-prone and we treat its accuracy as a limitation rather than a solved problem. Restricting matches to the technology subclasses in which a person patented academically removes the bulk of coincidental matches, and assigning each name to its modal institution suppresses cross-institution affiliations. Because common names dominate the residual risk, we re-estimate on names with at most ten pre-period appearances and again on names with at most five; both variants appear in the results tables alongside the full measure. False matching adds noise to the dependent variable and would bias the estimates if false-match rates evolved differently across

groups, so we rely on the restricted-name variants to show how the estimate behaves as identification of individuals tightens. We also provide a data-based check. For a sample of 556,833 matched name-patent pairs from the 2022 publication year we ask whether the firm-side patent carries at least one co-inventor who also appears on the person's pre-period academic patents, which is strong corroboration that the match is the same individual. Team continuity holds for 14.9 percent of matches overall, 22.7 percent for names with at most ten academic appearances and 13.9 percent for more common names. We call this the corroborated-match share. It is not a precision estimate, since a scientist can patent at a firm with an entirely new team and would then be a true match without corroboration, but the gap between rare and common names is the pattern false matching would produce and is why we report the restricted variants alongside the full measure.

We also considered citation-based quality measures, but forward citations are populated for under a third of records in this register and accumulate too slowly to inform post-2020 cohorts.

Table 1 reports descriptive statistics for the main sample. Designated universities are larger than controls before designation, with 1,095 against 720 mean annual applications, and file a higher share of invention patents, but they start from almost identical co-application shares, 6.8 against 6.9 percent, and similar rates of firm-side patenting, 4.1 against 3.7 percent of their rosters. Level differences of this kind are absorbed by institution fixed effects; differential trends are what the event studies in Section 5.2 examine.

Table 1
Descriptive statistics.

Applications per institution-year	Pilot pre	Pilot post	Control pre	Control post
	1,095	2,249	720	1,146
Invention patent share	0.803	0.830	0.746	0.782
University-industry co-application share	0.068	0.083	0.069	0.096
Firm-side patenting rate	0.041	0.065	0.037	0.057
Universities	16	16	88	88
Distinct applications (all years, full sample)	1,579,087			
Academic inventor roster (2010–2019)	586,317			
Firm-side inventor matches	9,523,552			
Notes. Pre covers application years 2013 to 2020, post 2021 to 2023, reflecting publication of the national list in October 2020. Designated refers to the 16 double first-class universities of the 2020 national cohort that form the main sample, after excluding four Sichuan early adopters, nine other designated universities, ten research institutes and one hospital, all reported as extensions. Control refers to 88 double first-class universities never designated during the window, after dropping two Chengdu universities whose provincial status could not be verified. Firm-side statistics use the subsample with at least twenty pre-period inventors, 16 designated and 87 control universities. All outcomes are counts or shares.				

Figure 2 plots the raw group means. Both groups roughly double patent output over the window and both show rising co-application and firm-side patenting, with no visible divergence at designation. The invention share rises in both groups after 2020, a movement we attribute to the phase-out of filing subsidies and the accompanying campaign against low-quality applications rather than to the reform.

4. Empirical strategy

For institution i and application year t we estimate

$$y_{it} = \alpha_i + \lambda_t + \beta (\text{Designated}_i \times \text{Post}_t) + \varepsilon_{it},$$

where α and λ are institution and application-year fixed effects, Designated equals one for universities on the October 2020 national list, and Post equals one from 2021, the first application year fully exposed to designation. The estimation sample is 16 designated and 88 control universities over 2013 to 2023. A balanced panel would contain 1,144 institution-years and we observe 1,140, because Shanghai University of Finance and Economics and Zhongnan University of Economics and Law report no patents in one and three early years respectively; both are control universities and the gaps precede designation. Firm-side outcomes cover 1,126 institution-years across 103 universities. A balanced panel would

contain 1,133, and the seven missing cells all belong to Shanghai University of Finance and Economics, whose roster supports the measure in only four years; the roster threshold of twenty pre-designation inventors removes one further control university and no designated university. Standard errors are clustered by institution. Count outcomes enter as $\log(1 + y)$, which we prefer to dropping zeros because 13 institution-years have no co-applications; share outcomes enter unweighted so that each institution contributes equally. The coefficient β is the intention-to-treat effect of designation.

Because the national cohort is designated at a single date and compared against never-designated controls, the negative-weighting problems of staggered designs do not arise (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfoeuille, 2020), and the Sichuan cohort is estimated separately against the same pool rather than pooled (Callaway and Sant'Anna, 2021). Event-study versions interact designation with event-time indicators from eight years before to two years after, with the year before designation as reference (Sun and Abraham, 2021; Roth et al., 2023).

Identification requires parallel trends conditional on fixed effects, and we are careful about what our tests establish. Failing to reject a joint test of pre-treatment coefficients is not evidence that trends are parallel, particularly with 16 treated clusters, so we report the pre-treatment coefficients and their standard errors in full and let readers judge the width of the region they exclude. We complement this with a matched subsample. Propensity scores come from a logit of designation on pre-designation log patent stock, invention share and co-application share, estimated on the 104 universities of the main sample. We match one to two on the nearest neighbour without replacement within a caliper of 0.2 standard deviations of the linearised score. Because matched sets vary in size, we report unweighted regressions and regressions in which each designated university receives weight one and each of its matched controls receives weight one divided by the number of controls matched to that designated university; the two agree. Appendix B reports the propensity model, common support, balance and matched-set composition.

We also report permutation placebo tests, reassigning designation to random sets of universities and comparing the observed estimate with the resulting distribution. These are label-sensitivity checks, not randomisation inference and not a validation of the clustered standard errors, since designation was not random and no known assignment mechanism justifies exchangeability. Official documents state that recommending authorities favoured institutions with stronger reform capacity and commercialisation records. For the same reason we do not claim that selection signs the bias. Selection on strong performers could bias estimates upward, but it could equally work against finding an effect if designated institutions were already near a ceiling, if their pre-designation performance was unusually high and mean-reverting, or if authorities selected units precisely because transfer bottlenecks were salient there.

Because several of our conclusions rest on failing to reject, we report for every estimate the 95 percent confidence interval and read its upper bound as the largest positive effect the data exclude (Abadie, 2020). We do not conduct formal equivalence tests, since the policy documents set no numerical target against which an equivalence margin could be pre-specified, and we therefore avoid the language of

equivalence and simply state the excluded region in relative terms. Readers who wish to apply their own threshold can compare it with the bounds we report. We also report the minimum detectable effect at eighty percent power, computed as 2.8 times the standard error, which describes what the design could have found. With 16 treated clusters, cluster-robust standard errors are fragile, and we therefore also report permutation distributions as a descriptive cross-check on the sensitivity of estimates to which universities carry the designation label.

5. Results

5.1 Main estimates

Table 2 reports the baseline estimates. The institutional channel did not move. Co-application falls by 1.2 percentage points against a standard error of 0.7 and the invention share by 1.1 against 2.1, and log co-applications rise by 0.086 against 0.092. None of these approaches significance, and as Section 2.3 explains the negative sign carries no directional information, since inventor-led commercialisation can itself remove the university from the applicant list. Patent volume rises by 0.169 log points and reaches the five percent level under clustered inference, but its permutation p-value is 0.17 and its interval runs from 0.4 to 33 percent, so filing volume remains genuinely uninformative here. The individual channel did move. Firm-side patenting by pre-designation academic inventors rises by 0.32 percentage points on a base of 5.7 percent, about six percent in relative terms, with a p-value of 0.011. The estimate is not an artefact of common Chinese surnames. Restricting matches to names with at most ten academic appearances gives 0.22 percentage points and at most five gives 0.17, both significant, so tightening the identification of individuals shrinks the estimate proportionally without eliminating it.

Table 2

Effects of designation on designated universities, baseline difference-in-differences estimates.

Outcome	ITT	SE	p	95% CI	Base	MDE
Log applications	0.169	0.084	0.045	[0.004, 0.333]	6.60	0.235
Invention patent share	-0.011	0.021	0.60	[-0.051, 0.029]	0.782	0.057
Co-application share	-0.012	0.007	0.090	[-0.027, 0.002]	0.096	0.020
Log co-applications	0.086	0.092	0.35	[-0.095, 0.267]	4.06	0.258
Firm-side patenting rate	0.0032	0.0012	0.011	[0.0007, 0.0056]	0.057	0.003
Firm-side rate, names ≤ 10	0.0022	0.0010	0.027	[0.0002, 0.0041]	0.031	0.003
Firm-side rate, names ≤ 5	0.0017	0.0008	0.036	[0.0001, 0.0032]	0.021	0.002

Notes. Each row is a separate regression of the outcome on the designation-by-post interaction with institution and application-year fixed effects, application years 2013 to 2023, post from 2021. The sample has 1,140 institution-years covering 16 designated and 88 control double first-class universities, and 1,126 institution-years covering 103 universities for firm-side outcomes. Standard errors are clustered by institution. Base is the mean for control universities in post years, in logs where the outcome is logged. MDE is the minimum detectable effect at eighty percent power. R^2 ranges from 0.66 for the co-application share to 0.95 for the firm-side rate.

The intervals tell the substantive story. On co-application the upper bound is 0.19 percentage points, two percent of the control mean, so the channel through which the reform was expected to work can be ruled out as a source of any economically meaningful gain. On firm-side patenting the interval runs from one to ten percent of the base and excludes zero. These two facts together are the paper's core result, and they point in opposite directions.

5.2 Dynamics

Table 3 traces the estimates through event time, beginning at event year minus eight so that the full sample period from 2013 appears. Joint tests do not reject flat pre-designation paths, with p-values of 0.21 for co-application and 0.18 for firm-side patenting. Co-application shows no step at designation and no drift after it, with post coefficients of minus 1.4, minus 1.9 and minus 2.0 percentage points against pre-designation values of similar magnitude. Firm-side patenting steps up at designation and stays up, at 0.21, 0.14 and 0.15 percentage points. The one caveat visible in the table is that firm-side coefficients drift upward over the earlier pre-designation years before flattening in the three years before designation; a continuing trend of that kind could generate positive post coefficients on its own, and it is the reason we test the result against matched controls in Section 5.3.

Table 3
Event-study estimates for headline outcomes.

Event year	Co-application share	Firm-side patenting rate
-8	-0.0161 (0.0149)	-0.0032 (0.0018)
-7	-0.0194 (0.0136)	-0.0027 (0.0017)
-6	-0.0065 (0.0145)	-0.0023 (0.0018)
-5	-0.0020 (0.0114)	-0.0022 (0.0017)
-4	0.0004 (0.0124)	-0.0016 (0.0014)
-3	-0.0015 (0.0104)	0.0000 (0.0014)
-2	0.0032 (0.0103)	-0.0001 (0.0009)
-1	reference	reference
0	-0.0139 (0.0072)	0.0021 (0.0008)
+ 1	-0.0189 (0.0094)	0.0014 (0.0009)
+ 2	-0.0199 (0.0124)	0.0015 (0.0011)
Pre-trend F (p)	1.41 (0.21)	1.50 (0.18)
Notes. Coefficients on designation-by-event-time interactions, event year zero is 2021, year minus one omitted, institution and year fixed effects included, standard errors clustered by institution in parentheses. The final row reports the joint F-test of all pre-treatment coefficients. Failing to reject does not establish parallel trends.		

Figure 3 shows the institution-level counterpart. Designated universities sit interleaved with controls around the diagonal for both outcomes, with no visible tendency to lie above it, which is consistent with the increases in Figure 2 being common to both groups.

The aggregate movements visible in Fig. 2 deserve comment because they are large. Filing volumes in both groups surged after 2020 and invention shares rose, developments that reflect system-wide forces, among them the phase-out of filing subsidies announced in 2021 and the accompanying campaign against low-quality applications, which culled utility models and mechanically raised invention shares. These are common shocks absorbed by the year effects, and designated and control universities respond to them identically, which is itself informative given how large the movements are.

5.3 Robustness

Table 4 collects the robustness results, and they divide along a clear line. The co-application estimate is stable and never significant, at minus 0.011 with a window ending in 2022, minus 0.012 dropping 2020, minus 0.012 excluding Anhui University, minus 0.011 when outcomes are rebuilt from applications published within two years of filing, and minus 0.005 on the matched sample. The firm-side estimate under the same two-year maturity rule is 0.0032, identical to the baseline. The firm-side estimate is

equally stable, at 0.0032 with a window ending in 2022, 0.0034 dropping 2020, 0.0031 excluding Anhui University and 0.0032 under the two-year maturity rule, with a permutation placebo p-value of 0.043. It moves on one dimension only. Matching the 16 designated universities to 30 propensity-score-matched controls reduces it to 0.0009. The matched comparison discards three quarters of the control pool to balance pre-designation patent stock, and with 30 controls it has little power, so we read the divergence as a limit on precision rather than as a refutation. Five of six perturbations leave the estimate intact; the sixth leaves it positive but no longer distinguishable from zero.

Three extensions appear at the foot of the table. Adding the nine designated universities outside the double first-class tier gives a firm-side estimate of 0.0024 and a co-application estimate of minus 0.007, so widening the sample attenuates the firm-side estimate without reversing it. The ten designated research institutes and the hospital, compared with university controls, show a co-application estimate of minus 0.027; we report this as exploratory because the counterfactual is a different institution type. The Sichuan cohort, designated in December 2016 and observed for seven post-designation years, shows no detectable effect on co-application and an output estimate significant only at the ten percent level, and rests on eight clusters with an institution list we could not fully verify.

Table 4
Robustness checks and extensions.

Panel A. Main sample specifications	Estimate	SE	p
Window ending 2022, co-application share	-0.011	0.007	0.11
Window ending 2022, firm-side rate	0.0032	0.0011	0.004
Window ending 2022, log applications	0.170	0.081	0.035
Dropping 2020, co-application share	-0.012	0.008	0.14
Dropping 2020, firm-side rate	0.0034	0.0014	0.015
Excluding Anhui University, co-application share	-0.012	0.007	0.089
Excluding Anhui University, firm-side rate	0.0031	0.0012	0.013
Two-year maturity window, co-application share	-0.011	0.007	0.13
Two-year maturity window, firm-side rate	0.0032	0.0013	0.011
Two-year maturity window, log applications	0.173	0.089	0.053
Matched sample, co-application share	-0.005	0.008	0.55
Matched sample, firm-side rate	0.0009	0.0015	0.57
Matched sample weighted, firm-side rate	0.0009	0.0015	0.57
Permutation placebo p, co-application share			0.30
Permutation placebo p, firm-side rate			0.043
Permutation placebo p, log applications			0.17

Panel B. Extensions beyond the main sample	Estimate	SE	p
All 25 designated universities, co-application share	-0.007	0.007	0.31
All 25 designated universities, firm-side rate	0.0024	0.0010	0.024
Research institutes and hospital, co-application share	-0.027	0.015	0.07
Sichuan cohort, co-application share	0.008	0.023	0.71
Sichuan cohort, log applications	0.324	0.178	0.07
Notes. Matched-sample rows use one-to-two nearest-neighbour propensity score matching without replacement on pre-designation log patent stock, invention share and co-application share, giving 16 designated and 30 control universities; the weighted row gives each designated university weight one and each matched control weight one divided by the number of controls in its matched set. Appendix B reports diagnostics. Two-year maturity rows rebuild outcomes from applications published within two calendar years of filing, a window that is nearly complete for the 2023 cohort; the rows ending in 2022 remove the residual maturity concern entirely. Rows for research institutes and the hospital compare the eleven non-university designated units with the same university control pool and are exploratory, since no never-designated elite institute pool exists. Permutation rows report two-sided p-values from 2,000 random reassignments of the designation label on pre-post collapsed data and are label-sensitivity checks rather than randomisation inference. Sichuan rows compare eight verified early adopters with the never-designated pool from 2011, designation dated December 2016 with post from 2017, and rest on eight clusters.			

Figure 4 visualises the permutation placebo distributions. For co-application and volume the observed values fall inside the placebo distributions, with p-values of 0.30 and 0.17. For firm-side patenting the observed value sits in the upper tail with a p-value of 0.043.

5.4 Heterogeneity

Table 5 reports the six subgroup comparisons we specify as a family, three for each headline outcome, splitting designated units into double first-class universities, other universities, and research institutes with the hospital. The firm-side response is confined to the double first-class group, at 0.32 percentage points against 0.09 for other designated universities and minus 0.08 for the non-university units. Co-application is negative for the double first-class group and for the non-university units and slightly positive for other universities, with none reaching significance.

Multiple testing tempers the one significant estimate. Holm adjustment across the six tests raises the p-value on the double first-class firm-side estimate from 0.011 to 0.067, and every other adjusted p-value exceeds 0.34. The concentration of the response in the most research-intensive universities is therefore a pattern we report with the multiplicity correction attached, and it is the pattern the complementarity argument in Section 6 predicts.

Table 5
Subgroup estimates with Holm-adjusted p-values.

Subgroup	ITT	SE	p	Holm p	Designated n
Firm-side rate, double first-class	0.0032	0.0012	0.011	0.067	16
Co-application, research institutes and hospital	-0.0266	0.0146	0.068	0.340	11
Co-application, double first-class	-0.0123	0.0073	0.090	0.359	16
Firm-side rate, other designated universities	0.0009	0.0012	0.444	1.000	9
Co-application, other designated universities	0.0031	0.0077	0.686	1.000	9
Firm-side rate, research institutes and hospital	-0.0008	0.0019	0.693	1.000	10
Notes. Rows are ordered by raw p-value. Each regression retains the 88 control universities and the indicated designated subgroup, with institution and application-year fixed effects and standard errors clustered by institution. The family of six tests is defined before estimation as the three institution-type splits applied to the two headline outcomes. Holm-adjusted p-values are computed from unrounded raw p-values using the step-down procedure.					

5.5 How small is small

Confidence bounds acquire meaning through benchmarks, and one is available here. When individual ownership was withdrawn in Norway, patenting and firm formation by university researchers fell by roughly half (Hvide and Jones, 2018). Our design would detect gains far smaller than that. The minimum detectable effect at eighty percent power is about 21 percent of the base for the co-application share and about six percent for the firm-side rate, and the realised upper bound excludes increases above about two percent of the base for co-application. Any response approaching the Norwegian magnitude would have been visible on either margin. We stress that the comparison is heuristic, since the two settings differ in the direction of the change, in the rights actually conferred, in implementation intensity, in outcome definitions and in observation window.

In absolute terms the upper bound on the co-application share, 0.19 percentage points, corresponds to about two additional joint applications per year at the post-period control mean of 1,146 filings. The firm-side point estimate of 0.32 percentage points corresponds to about 23 additional pre-designation inventors appearing on company patents, computed as 0.0032 multiplied by the mean roster of 7,224 names. These magnitudes describe the patent-visible response to designation. They say nothing about licensing income, contract research or company formation that leaves no patent trace within our window.

6. Discussion

6.1 Why the institutional channel did not move

The two results have to be explained together. Designation gave universities the authority to hand ownership shares to inventing teams, and in the three years that followed those universities collaborated with firms exactly as much as universities that received no such authority, while their scientists appeared on company patents somewhat more often. Three accounts can produce that pattern, and we think the evidence favours one of them.

The first is that take-up was too limited to register. Rights reach a scientist only after disclosure, review and a recorded agreement, and completed assignments are not published, so a small number of transactions would leave institutional aggregates flat. We regard this as an incomplete explanation rather than a wrong one. If designation had produced only a handful of assignments, it should have moved nothing at all, yet the individual channel moved and did so immediately in 2021, precisely where a small number of high-value transactions would show up first.

The second is that the reform relaxed a constraint that was not binding for institutions. This is the account we favour. Chinese universities already had the legal authority to license and to take equity, and inventors already held a statutory claim to at least half of net licensing income. What designation added was control rights for individuals and administrative cover for the officials who transfer them. Neither addresses the constraints that actually govern institutional collaboration, which run through the capacity of transfer offices, the terms firms are willing to offer, and evaluation systems that reward publications and titles. Ownership was the wrong margin for the institutional channel, and the flat co-application estimate is what a well-identified test of the wrong margin looks like.

The third account explains why the individual channel behaved differently. Control rights matter most to the minority of academics who want to commercialise through a firm rather than through their department, and for them the binding constraint really was the inability to move an invention out of institutional custody. That group is small, which is why the aggregate effect is six percent rather than sixty, and it is concentrated where the surrounding infrastructure already exists, which is why the entire effect sits in the double first-class universities. On this reading the reform worked, narrowly, on exactly the population property rights theory says it should reach, and did nothing for everyone else.

Two observations support this interpretation against the alternatives. The 2015 law revision, which delegated disposal rights and raised inventor reward floors nationwide, applied to designated and control universities alike and is absorbed by year effects, so what we estimate is the incremental effect of control rights over and above income rights. And the institutional null is the most robust result in the paper, holding across every specification we ran, while the individual result is confined to the group with the strongest transfer infrastructure. A reform that failed uniformly would not produce that contrast.

6.2 Relation to the professor's privilege literature

Placed beside Hvide and Jones (2018), our results establish an asymmetry that the ownership debate has assumed away. Norway withdrew individual ownership from a system built around it and academic entrepreneurship halved. China extended individual ownership into a system built around institutional control and university-industry collaboration did not move at all, while individual patenting shifted by six percent in the best-equipped universities. The direction of a legal change therefore does not determine the size, or even the existence, of its behavioural effect. Ownership rules operate through the transfer infrastructure that surrounds them, and that infrastructure is what the two countries differ in. In Norway the privilege sat atop functioning channels connecting professors to investors and firms. In China the new rights were introduced into an environment where valuation services, liability safe harbours, personnel mobility rules and early-stage capital are still developing, and where the rights themselves reach individuals only after institutional processes that preserve a veto.

The comparison is across countries and across directions of change, so it identifies a pattern rather than a parameter. But the pattern is the substantive point, and it has a clear implication for a policy literature that has treated ownership allocation as the central design question. Where the complementary machinery exists, ownership rules bind and moving them has large effects. Where it does not, ownership is close to inert, and reforms that reassign it will disappoint regardless of how generous the assignment is. This is the sense in which effective transfer reform is a bundle rather than a lever (Siegel and Wright, 2015; Kochenkova et al., 2016; Aghion et al., 2008).

The reading generates testable predictions. Provincial schemes that followed increasingly attach due-diligence exemptions, transformation funds and valuation services, and if complements bind, those variants should outperform the original designation. Effects may also surface in registers patents do not reach, since business registration data would capture company formation directly and administrative series on assignment counts and licensing income would capture rent redistribution. Recent work on the provincial mixed-ownership experiments reports that mixed ownership raises inventor-led commercialisation (Liu et al., 2025). The two sets of results fit together once the designs are compared. That study examines the provincial mixed-ownership model, which assigns shares directly and has run since 2016, and measures commercialisation from university statistical reporting. We examine designation into the national scheme and measure patent outcomes. The margin on which our estimates move, individual firm-side patenting, is the one closest to the inventor-led commercialisation they identify, so both designs locate the response with individuals rather than institutions. Pairing administrative assignment counts with the designation variation used here would let the two designs be compared directly, and is the natural next step.

6.3 Limitations

Four limitations bound the interpretation, and we state each once. Our estimand is designation rather than rights received, so the paper measures what a government controls, not what an individual scientist experiences; converting these intention-to-treat estimates into dose-response estimates requires the assignment counts that pilot institutions already report to their supervising ministries but do not publish. Outcomes are confined to patents, so licensing income, contract research and company formation that

leaves no patent trace are invisible, and the institutional null should be read as covering co-patenting rather than every form of collaboration. The firm-side measure identifies individuals by name and technology class, which is why we report the restricted-name variants; we have not validated it against labelled records, and doing so with addresses, personnel directories or business registration is the obvious next step for anyone building on it. And the firm-side estimate loses significance against propensity-score-matched controls, which places its magnitude, though not its sign, beyond what three post-designation years can settle.

7. Conclusion

China conducted the experiment the technology transfer literature had been missing. It offered ownership of academic inventions to the scientists who made them, at national scale, with dated designation that permits clean comparison, and with a control group of equally elite universities that received nothing. Three years on, the institutional machinery of technology transfer is unchanged. Designated universities collaborate with firms no more than their peers, and the data are precise enough to rule out anything beyond a two percent gain in co-patenting. What changed is where individual academics patent. In China's most research-intensive universities, scientists who had been patenting through their institution began appearing on company patents about six percent more often, a shift that arrives with designation and survives every test we could apply short of matching on pre-designation scale.

The result is not that ownership reform fails. It is that ownership reform reaches individuals and stops there. That distinction matters for policy because the two channels are managed by different instruments. Governments seeking more university-industry collaboration will not obtain it by reassigning title, since the constraints on collaboration lie in transfer capacity, in what firms will pay for early-stage academic inventions, and in academic evaluation systems that still reward publication over commercialisation. Governments seeking more academic entrepreneurship may well obtain it by reassigning title, but only where the surrounding infrastructure lets a scientist act on the rights they receive. The Chinese pilots supplied the rights nationally and the infrastructure unevenly, and the geography of our effect follows the infrastructure rather than the rights.

Three extensions follow directly. The assignment counts that pilot institutions report to their ministries would convert our intention-to-treat estimates into dose-response estimates and settle how much of the institutional null reflects limited take-up. Business registration data would capture the company formation our patent measures can only glimpse. And the provincial schemes that began in 2022, which bundle ownership with due-diligence exemptions, transformation funds and valuation services, provide the variation needed to test the complementarity argument this paper advances, since bundled variants should outperform the unbundled designation we study. The measurement apparatus developed here transfers directly to all three.

Declarations

Author Contribution

Li wrote the main manuscript text and prepared figures , Zhong analysed the data, All authors reviewed the manuscript.

Data Availability

The findings of this study are based exclusively on publicly available secondary data: the China National Intellectual Property Administration (CNIPA) patent register (1985–2025), publicly disclosed listed-firm annual reports, and published academic literature. No new primary datasets were generated. The derived institution-year panels and analysis code will be deposited in a public repository upon acceptance.

Conflict of interest

The authors declare no conflict of interest.

References

1. Abadie, A. (2020). Statistical nonsignificance in empirical economics. *American Economic Review: Insights*, 2(2), 193–208.
2. Aghion, P., Dewatripont, M., & Stein, J. C. (2008). Academic freedom, private-sector focus, and the process of innovation. *RAND Journal of Economics*, 39(3), 617–635.
3. Åstebro, T., Bazzazian, N., & Braguinsky, S. (2012). Startups by recent university graduates and their faculty: Implications for university entrepreneurship policy. *Research Policy*, 41(4), 663–677.
4. Bloom, N., Van Reenen, J., & Williams, H. (2019). A toolkit of policies to promote innovation. *Journal of Economic Perspectives*, 33(3), 163–184.
5. Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
6. Chen, A., Patton, D., & Kenney, M. (2016). University technology transfer in China: A literature review and taxonomy. *Journal of Technology Transfer*, 41(5), 891–929.
7. Czarnitzki, D., Doherr, T., Hussinger, K., Schliessler, P., & Toole, A. A. (2016). Knowledge creates markets: The influence of entrepreneurial support and patent rights on academic entrepreneurship. *European Economic Review*, 86, 131–146.
8. Dang, J., & Motohashi, K. (2015). Patent statistics: A good indicator for innovation in China? Patent subsidy program impacts on patent quality. *China Economic Review*, 35, 137–155.
9. de Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.

10. D'Este, P., & Perkmann, M. (2011). Why do academics engage with industry? The entrepreneurial university and individual motivations. *Journal of Technology Transfer*, 36(3), 316–339.
11. Fini, R., Rasmussen, E., Siegel, D., & Wiklund, J. (2018). Rethinking the commercialization of public science: From entrepreneurial outcomes to societal impacts. *Academy of Management Perspectives*, 32(1), 4–20.
12. Fisch, C. O., Block, J. H., & Sandner, P. G. (2016). Chinese university patents: Quantity, quality, and the role of subsidy programs. *Journal of Technology Transfer*, 41(1), 60–84.
13. Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
14. Grimaldi, R., Kenney, M., Siegel, D. S., & Wright, M. (2011). 30 years after Bayh-Dole: Reassessing academic entrepreneurship. *Research Policy*, 40(8), 1045–1057.
15. Grossman, S. J., & Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4), 691–719.
16. Hart, O., & Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6), 1119–1158.
17. Hu, A. G., & Jefferson, G. H. (2009). A great wall of patents: What is behind China's recent patent explosion? *Journal of Development Economics*, 90(1), 57–68.
18. Hvide, H. K., & Jones, B. F. (2018). University innovation and the professor's privilege. *American Economic Review*, 108(7), 1860–1898.
19. Kenney, M., & Patton, D. (2009). Reconsidering the Bayh-Dole Act and the current university invention ownership model. *Research Policy*, 38(9), 1407–1422.
20. Kochenkova, A., Grimaldi, R., & Munari, F. (2016). Public policy measures in support of knowledge transfer activities: A review of academic literature. *Journal of Technology Transfer*, 41(3), 407–429.
21. Lach, S., & Schankerman, M. (2008). Incentives and invention in universities. *RAND Journal of Economics*, 39(2), 403–433.
22. Liu, X., Li, X., Yang, W., & Zhang, B. (2025). How does patent mixed ownership impact university technology commercialization? Evidence from China. *Technovation*, 145, 103259.
23. Ministry of Science and Technology, National Development and Reform Commission, Ministry of Education, Ministry of Industry and Information Technology, Ministry of Finance, Ministry of Human Resources and Social Security, Ministry of Commerce, National Intellectual Property Administration, & Chinese Academy of Sciences (2020a). Implementation plan for the pilot programme granting research personnel ownership or long-term use rights of job-related scientific and technological achievements. Guo Ke Fa Qu [2020] No. 128, issued 9 May and released 18 May. Beijing.
24. Ministry of Science and Technology (2020b). List of pilot units for granting research personnel ownership or long-term use rights of job-related scientific and technological achievements. Guo Ke Fa Qu [2020] No 273, 12 October. Beijing.

25. Motohashi, K., & Yun, X. (2007). China's innovation system reform and growing industry and science linkages. *Research Policy*, 36(8), 1251–1260.
26. Mowery, D. C., Nelson, R. R., Sampat, B. N., & Ziedonis, A. A. (2001). The growth of patenting and licensing by U.S. universities: An assessment of the effects of the Bayh-Dole Act of 1980. *Research Policy*, 30(1), 99–119.
27. Perkmann, M., Salandra, R., Tartari, V., McKelvey, M., & Hughes, A. (2021). Academic engagement: A review of the literature 2011–2019. *Research Policy*, 50(1), 104114.
28. Perkmann, M., Tartari, V., McKelvey, M., Autio, E., Broström, A., D'Este, P., Fini, R., Geuna, A., Grimaldi, R., Hughes, A., Krabel, S., Kitson, M., Llerena, P., Lissoni, F., Salter, A., & Sobrero, M. (2013). Academic engagement and commercialisation: A review of the literature on university-industry relations. *Research Policy*, 42(2), 423–442.
29. Roth, J., Sant'Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2), 2218–2244.
30. Sandström, C., Wennberg, K., Wallin, M. W., & Zherlygina, Y. (2018). Public policy for academic entrepreneurship initiatives: A review and critical discussion. *Journal of Technology Transfer*, 43(5), 1232–1256.
31. Siegel, D. S., & Wright, M. (2015). Academic entrepreneurship: Time for a rethink? *British Journal of Management*, 26(4), 582–595.
32. Sichuan Department of Science and Technology, & Sichuan Intellectual Property Office (2016). Implementation plan for the pilot reform of ownership of job-related scientific and technological achievements in Sichuan Province. *Chuan Ke Zheng* [2016] 5, December. Chengdu.
33. Zhejiang Provincial Department of Science and Technology and five other departments (2022). Implementation plan for expanding the pilot programme granting research personnel ownership or long-term use rights of job-related scientific and technological achievements. *Zhe Ke Fa Cheng* [2022] No 29, issued 26 August, effective 29 September. Hangzhou.
34. Anhui Provincial Department of Science and Technology (2024). Notice on expanding the pilot reform granting ownership or long-term use rights of job-related scientific and technological achievements. Hefei.
35. Hubei Provincial Department of Science and Technology (2024). Notice on filing arrangements for the reform granting ownership of job-related scientific and technological achievements. 7 March. Wuhan.
36. Guangxi Science and Technology Department and six other departments (2024). Implementation plan for the pilot reform of the management of job-related scientific and technological achievements in Guangxi. *Gui Ke Fa* [2024] No 61, 22 May. Nanning.
37. Liaoning Provincial Department of Science and Technology (2025). Notice on the comprehensive implementation of the reform granting ownership of job-related scientific and technological achievements to research personnel. 27 March. Shenyang.

38. Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
39. Thursby, J. G., & Thursby, M. C. (2011). Has the Bayh-Dole Act compromised basic research? *Research Policy*, 40(8), 1077–1083.
40. Wang, J., & Qian, Y. (2023). The impact of university patent ownership on commercialization. NBER Working Paper 31021, March 2023, revised December 2024. National Bureau of Economic Research.
41. Wei, S. J., Xie, Z., & Zhang, X. (2017). From Made in China to Innovated in China: Necessity, prospect, and challenges. *Journal of Economic Perspectives*, 31(1), 49–70.

Appendix

Appendix A. Institution lists

Designated universities in the national university cohort (25), of which the 16 double first-class universities form the main sample. Beijing University of Technology, Shenyang University of Chemical Technology, University of Science and Technology Liaoning, Shanghai University, University of Shanghai for Science and Technology, Shanghai Maritime University, Nanjing Tech University, Soochow University, Zhejiang University of Technology, Hangzhou Dianzi University, Hubei University of Technology, Jinan University, Guangdong University of Technology, Hainan University, Fudan University, Shanghai Jiao Tong University, Nanjing University, Zhejiang University, Xi'an Jiaotong University, Harbin Institute of Technology, Beihang University, Beijing Institute of Technology, Northwestern Polytechnical University, Peking Union Medical College, and the University of Science and Technology of China. Source is the October 2020 institution list (Ministry of Science and Technology, 2020b).

Designated research institutes and hospital reported separately (11). Beijing Academy of Science and Technology, Jiangsu Industrial Technology Research Institute, Zhejiang Academy of Agricultural Sciences, Guangdong Academy of Sciences, Institute of Animal Science of the Chinese Academy of Agricultural Sciences, China National Rice Research Institute, Institute of Science and Technology of the National Health Commission, Lanzhou Institute of Chemical Physics of the Chinese Academy of Sciences, National Center for Nanoscience and Technology, Shanghai Institute of Microsystem and Information Technology of the Chinese Academy of Sciences, and Beijing Jishuitan Hospital.

Designated institutions excluded from the national cohort because they also belong to the Sichuan programme (4). Sichuan University, Southwest Jiaotong University, Chengdu University of Technology, and Chengdu University of Traditional Chinese Medicine.

Sichuan cohort analysed separately (8). The four institutions above plus Sichuan Agricultural University, Chengdu University, the Sichuan Academy of Chinese Medicine Sciences, and the Sichuan Institute of Mechanical Research and Design. Sources are the 2016 provincial implementation document and contemporaneous official reporting; the province did not publish a complete institution list, so this is a verified subset rather than the full programme.

Appendix B. Matching diagnostics

Propensity scores come from a logit of designation on pre-designation log patent stock, invention share and co-application share. The model reported here is estimated on the 104 universities of the same-tier main sample, with coefficients of 0.504 on log stock, 1.836 on invention share and minus 6.506 on co-application share and a pseudo R^2 of 0.043. The wider 25-university extension uses a separate model estimated on its own 113 universities, and matching is carried out within each sample rather than transferred between them. Matching is one to two on the nearest neighbour without replacement within a caliper equal to 0.2 standard deviations of the linearised score. In the main same-tier sample the 16 designated universities draw 30 distinct matched controls; in the wider 25-university sample they draw 47. Because matched-set sizes vary, we report unweighted regressions and regressions weighted by matched-set size, which give the same estimate of 0.0009 with a standard error of 0.0015. Designation propensities range from 0.11 to 0.41 among designated universities and from 0.01 to 0.40 among controls, so common support is wide.


Balance improves on the covariate that drives selection. The standardised mean difference in log patent stock falls from 0.42 before matching to 0.02 after, with matched means of 6.60 against 6.58 and a variance ratio of 0.92. The invention share moves from 0.17 to 0.01, with matched means of 0.764 against 0.763 and a variance ratio of 1.30. The co-application share moves from minus 0.10 to 0.03, with matched means of 0.061 against 0.060 and a variance ratio of 0.97. We report both unweighted matched regressions and weighted regressions. In the weighted version each designated university receives weight one and each matched control receives weight one divided by the number of controls matched to that same designated university, so every matched set contributes equally. The two give the same firm-side estimate of 0.0009 with a standard error of 0.0015. Standard errors remain clustered by institution.

Appendix C. Sample composition

Control universities (88). All double first-class universities not designated in any national or Sichuan pilot during the window, after excluding the University of Electronic Science and Technology of China and Southwestern University of Finance and Economics, both located in Chengdu, whose status under the Sichuan programme could not be verified. The full list, the institutional name dictionary and the construction code will be provided as supplementary material.

Figures

(a) Sequence of reform measures



(b) Pilot institutions by recommending authority




Figure 1

Panel a sets out the measures of the national framework, the Sichuan programme and the later provincial expansions, with arrows marking the sequence of steps within each track and the bar beneath indicating the application years used in estimation. Boxes are arranged in chronological order rather than drawn to scale. SWJTU denotes Southwest Jiaotong University. Panel b decomposes the 40 institutions of the October 2020 national list by recommending authority. Provincial recommendations span nine provinces and municipalities, while ministry recommendations comprise seven universities of the Ministry of Education, four of the industry and information technology ministry, two research institutes each from the agriculture ministry and the health commission, and four Chinese Academy of Sciences affiliates. Sources are the 2020 national pilot scheme and the October 2020 institution list of the Ministry of Science and Technology, the Sichuan provincial implementation document of 2016, and the provincial documents of Zhejiang, Anhui, Hubei, Guangxi and Liaoning cited in the text.




Figure 2

Unadjusted institution-mean outcomes by group and application year. The dotted line marks the October 2020 publication of the national pilot list, after which 2021 is the first fully exposed application year. Panel d scales distinct pre-reform academic inventors appearing on firm-filed patents by pre-reform roster size. Counts for 2024 onward are censored and excluded from estimation.




Figure 3

Each marker is one institution, plotting its 2015 to 2019 mean against its 2021 to 2023 mean. Markers on the diagonal are unchanged. Designated universities are interleaved with controls in both panels. The figure is descriptive and shows no visible separation of the two groups, which is weaker evidence than the regression estimates but useful for seeing the dispersion behind them.




Figure 4

Histograms of placebo estimates from 2,000 random reassignments of the designation label across the 104 sample universities, holding the number of designated universities fixed, estimated on pre-post collapsed panels. Vertical lines mark the observed estimates, with two-sided p-values of 0.30, 0.043 and

0.17. Because designation was not random, these are label-sensitivity checks rather than randomisation inference.