Policy Communication and Technology Transfer: Evidence from University Intellectual Property Governance Documents

Abstract

How universities communicate their intellectual property policies shapes how efficiently inventions become patent applications. Using the Policy Communication Index (PCI), a sentence-level BERT-based measure constructed from 87,321 sentences in IP governance documents at 141 US research universities over 1991–2023, we link within-institution changes in policy language to technology transfer outcomes. A one-standard-deviation increase in PCI is associated with approximately 9.7 percent more patent applications in the following year ($b = 1.476$, $p = 0.007$; randomization inference $p = 0.001$), a magnitude comparable to the effect of a ten-percentage-point increase in the inventor royalty share (Lach and Schankerman, 2004). An event study around discrete policy revision events shows parallel pre-trends for patent applications ($\chi^2(3) = 3.20$, $p = 0.361$) while disclosures show no post-revision response, consistent with a causal interpretation and locating the mechanism at the post-disclosure coordination stage. A Mundlak decomposition shows the association is entirely within-institution. Effects are concentrated at private and research-intensive universities. Policy communication quality is a low-cost and previously unmeasured complement to structural technology transfer interventions.

Keywords: Technology transfer efficiency; Policy communication; Administrative burden; University patenting; Text as data; Intellectual property governance

JEL Codes: O34, O38, I23, C55, L31

1 Introduction

University technology transfer has become a major engine of knowledge commercialization in the United States. In fiscal year 2022, US research universities executed more than 13,000 new license and option agreements, generated approximately $3.7 billion in licensing income, and filed roughly 24,700 new patent applications, up from fewer than 2,000 applications per year before the Bayh-Dole Act of 1980 (AUTM, 2022; Mowery, Nelson, Sampat, and Ziedonis, 1999). Yet four decades after Bayh-Dole created the institutional architecture for university technology transfer, a substantial and persistent puzzle remains: equally well-resourced universities differ dramatically in how efficiently they convert faculty invention disclosures into patented and licensed technologies. Among comparable Carnegie R1 institutions, the ratio of patent applications to invention disclosures varies by a factor of three or more (Thursby and Kemp, 2002; Anderson, Daim, and Lavoie, 2007), and this efficiency variance is only partially explained by the structural inputs the literature has documented. Understanding the sources of this unexplained heterogeneity has both theoretical and practical significance: even modest improvements in conversion efficiency, applied across the approximately 300 active TTOs in the United States, could substantially expand the pipeline from federally funded research to commercial application.

The technology transfer literature has developed a well-established account of the structural determinants of TTO performance. TTO size, staffing expertise, and organizational age predict licensing output (O'Shea, Allen, Chevalier, and Roche, 2005; Bercovitz and Feldman, 2008). Royalty sharing arrangements shape the incentives of faculty inventors to disclose and participate in commercialization (Lach and Schankerman, 2004). Research expenditure and the quality of the faculty research base establish the upstream supply of potentially patentable inventions (Siegel, Waldman, and Link, 2003; Rothaermel, Agung, and Jiang, 2007). Faculty participation in the disclosure-to-licensing process is itself a critical input, since most university inventions are disclosed at an early stage that requires continued inventor involvement post-license (Jensen and Thursby, 2001; Jensen, Thursby, and Thursby,

2003). What this structural literature has not examined is the *communicative* dimension of the IP governance environment: how an institution describes its technology transfer procedures to faculty inventors in written policy documents, and whether the tone, procedural clarity, and legal orientation of those documents independently predict technology transfer efficiency.

This gap is theoretically motivated by the administrative burden framework (Moynihan, Herd, and Harvey, 2015; Herd and Moynihan, 2018), which identifies learning costs, psychological costs, and compliance costs as three channels through which the form of institutional communication shapes individual engagement with bureaucratic processes. Applied extensively to citizen-state interactions, including benefits take-up (Heinrich, 2016), tax compliance (Weaver, 2010), and public program administration, this framework has not been extended to the intra-organizational governance contexts in which written documents mediate the relationship between faculty inventors and their institution's technology transfer office. If IP governance documents are drafted in opaque, legalistic, or adversarial language, faculty inventors face higher learning costs in understanding what is expected of them, higher psychological costs from encountering compliance-oriented framing, and higher perceived compliance costs in navigating the TTO pipeline. No existing study has tested whether the communicative form of IP governance documents, independent of their substantive content, predicts technology transfer outcomes.

We fill this gap using the Policy Communication Index (PCI), a sentence-level NLP measure of how collaboratively versus restrictively a university frames its technology transfer procedures to faculty inventors. The PCI is constructed in a companion paper (Sharma et al., 2026) via BERT-based classification (Devlin, Chang, Lee, and Toutanova, 2019) of 87,321 sentences drawn from IP governance documents at 141 US research universities between 1991 and 2023. The classifier achieves inter-rater reliability of Cohen's kappa = 0.74 against human coding and test-set classification accuracy above 0.82 across all sentence categories, performance comparable to human-coding benchmarks established in the text-as-data

literature (Grimmer, Roberts, and Stewart, 2022). We link institution-year PCI values to the AUTM Licensing Activity Survey panel and estimate two-way fixed-effects (TWFE) models in which *within*-institution changes in PCI predict *within*-institution changes in technology transfer outcomes. The identifying variation comes from policy revision events, defined as discrete years in which a university updated its IP governance documents. We supplement the main TWFE specification with three identification strategies: an event study around large positive policy revisions that directly tests for pre-revision parallel trends; a Mundlak decomposition separating within- from between-institution variation; and within-institution randomization inference that addresses the high serial correlation of PCI.

The results indicate that more supportive IP policy communication is robustly associated with higher patent application activity. The effect size is substantial relative to known structural interventions: the approximately 9.7 percent gain per standard deviation in PCI is comparable in magnitude to the licensing activity increase Lach and Schankerman (2004) attribute to a ten-percentage-point increase in inventor royalty shares, achieved here through revision of document language rather than a change in financial incentives. These results make several contributions. We introduce policy communication quality as a previously unmeasured dimension of the IP governance environment and provide panel evidence linking within-institution changes in governance document language to within-institution changes in technology transfer efficiency. An event study around discrete policy revisions shows that patent applications do not anticipate the revision while disclosures remain flat throughout, supporting the interpretation that the effect runs through post-disclosure coordination rather than upstream disclosure behavior. We extend the administrative burden framework (Moynihan, Herd, and Harvey, 2015) from citizen-state interactions to intra-organizational governance, and we offer a BERT-based institution-year measure of governance document quality that can be adapted to other administrative settings where written rules mediate engagement between institutions and individuals.

The remainder of this paper is organized as follows. Section 2 reviews the relevant liter-

atures. Section 3 presents the conceptual framework. Section 4 describes data construction and measurement. Section 5 presents the empirical strategy. Section 6 reports the main results. Section 7 discusses interpretation, policy implications, and limitations. Section 8 concludes.

2 Literature Review

The US technology transfer literature traces its modern form to the Bayh-Dole Act of 1980. Mowery, Nelson, Sampat, and Ziedonis (1999) document how the Act accelerated university commercialization, while Mowery and Sampat (2001) show that substantial cross-institutional variation predates the Act and reflects deep differences in research capacity and organizational culture. Subsequent research has established a well-understood set of structural determinants of TTO output. Siegel, Waldman, and Link (2003) use stochastic frontier methods to identify cultural barriers between universities and industry, incentive misalignment, and resource inadequacy as the primary suppressors of TTO productivity. Lach and Schankerman (2004) provide quasi-experimental evidence that more generous royalty shares to inventors substantially increase disclosures and licensing activity, and faculty research quality is a consistent predictor of disclosure rates (Siegel, Waldman, and Link, 2003; Rothaermel, Agung, and Jiang, 2007). Organizational characteristics matter as well: TTO age, staff size, and prior commercialization experience each predict startup formation (O'Shea, Allen, Chevalier, and Roche, 2005), while organizational routines inherited from the TTO's founding shape behavior for decades (Bercovitz and Feldman, 2008). Di Gregorio and Shane (2003) find that inventor royalty policies and faculty intellectual eminence predict spinoff rates independently of TTO size, and Siegel, Waldman, Atwater, and Link (2004) document through survey evidence that TTO responsiveness and cultural fit with industry are important determinants of commercialization success. Feldman, Feller, Bercovitz, and Burton (2002) show how equity stakes in spinoffs alter TTO incentives at the licensing-versus-

equity margin. What this structural literature shares is inattention to the communicative dimension of IP governance: how the institution frames its technology transfer process in written documents directed at faculty inventors.

A distinct stream of work measures the efficiency with which TTOs convert disclosures into patents and licenses, finding wide and persistent dispersion that the structural inputs above do not fully explain. Thursby and Thursby (2002) decompose the growth in university licensing over 1994 to 1999, attributing the increase to both greater TTO scale and genuine productivity improvement. Thursby and Kemp (2002) estimate DEA efficiency scores and find that the disclosure-to-patent conversion rate varies by a factor of three or more across comparable institutions. Anderson, Daim, and Lavoie (2007) introduce the Patenting Control Ratio as a standardized benchmarking measure and document similarly wide dispersion. Chapple, Lockett, Siegel, and Wright (2005) confirm low mean efficiency for UK TTOs, and Caldera and Debande (2010) find that TTO size predicts contract research but matters less for licensing, suggesting that different pipeline stages have different organizational prerequisites. The key empirical regularity across this literature is that within-institution time-series variation in conversion efficiency is largely unexplained.

The administrative burden framework (Moynihan, Herd, and Harvey, 2015) offers a theoretical account of one neglected source of that variation. The framework identifies three cost components that institutional communication imposes on individuals seeking to engage with bureaucratic processes: learning costs, the information-processing demands required to understand rules and procedures; psychological costs, the motivational effects of encountering adversarial or stigmatizing institutional language; and compliance costs, the direct effort required to complete mandated steps. Herd and Moynihan (2018) extend the framework across public program domains, and Heinrich (2016) provides empirical evidence that compliance costs reduce benefit program take-up. Meier and O'Toole (2003) show that managerial networking behavior predicts student achievement in public school districts, demonstrating that how managers structure their organizational environment has measurable consequences for

downstream performance.

A natural question is whether the administrative burden framework, developed primarily for citizen-state interactions where individuals face significant power asymmetries and information deficits, extends to faculty inventors who are highly educated, institutionally embedded, and professionally supported. We argue that it does, for two reasons specific to the TTO context. First, the post-disclosure coordination stage involves repeated procedural interaction rather than a single decision. The cumulative friction from unclear guidance compounds across multiple exchanges between the faculty inventor, TTO staff, and patent counsel; the administrative burden cost structure is therefore most consequential precisely at this iterative stage, where each round of coordination draws on the procedural clarity of written governance documents. Second, faculty inventors are not professional technology transfer agents; their primary expertise is scientific rather than legal or commercial. The learning costs of navigating IP governance procedures are non-trivial even for highly educated individuals, and the psychological costs of encountering compliance-oriented language may be particularly salient for faculty who regard the TTO as a service function rather than a regulatory body. Lipsky (2010) argues that how rules are communicated at the point of contact shapes behavior as consequentially as their formal content; this argument applies with particular force to repeated procedural interactions in which the quality of written guidance determines whether coordination tasks are completed efficiently. The administrative burden literature has concentrated on citizen-state interactions; we extend it to the interface between faculty inventors and university TTOs, where the primary written communication medium is the IP governance policy document.

A related body of work on policy framing and communication supports the view that the form of institutional language shapes individual behavior independently of its substantive content. Chong and Druckman (2007) show that how policy information is framed alters preferences and behavioral intentions even when the facts presented are identical. Druckman and Lupia (2000) model preference formation as depending on informational shortcuts

provided by institutional actors, and Schneider and Ingram (1993) develop the social construction framework in which policy language positions target populations as more or less deserving, affecting their willingness to engage with institutional processes. Lipsky (2010) argues that how rules are communicated at the point of contact between individuals and institutions is as consequential as the formal content of the rules themselves. Weaver (2010) documents how the design of compliance regimes, including their communicative framing, shapes voluntary compliance rates.

Text-as-data methods make it possible to measure the communicative properties of governance documents systematically and at scale. Grimmer and Stewart (2013) establish the foundations of automated text analysis for political science, and Grimmer, Roberts, and Stewart (2022) provide a comprehensive framework connecting computational text methods to social science inference. Wordscores (Laver, Benoit, and Garry, 2003) and Wordfish (Slapin and Proksch, 2008) recover positional estimates from text under distributional assumptions, while Rodriguez and Spirling (2022) show that word embedding models recover ideologically meaningful positions with substantial gains over bag-of-words approaches. Devlin, Chang, Lee, and Toutanova (2019) introduce BERT, enabling sentence-level classification with improved contextual understanding. Gentzkow, Kelly, and Taddy (2019) survey text-as-data applications in economics, and Wilkerson and Casas (2017) demonstrate their value for legislative research. In Sharma et al. (2026), we apply BERT-based classification to IP governance documents to construct the PCI; the present paper deploys that measure in a panel identification design.

Broader work on the institutional environment for innovation provides the macro-context for our analysis. Jaffe (1989) establishes that university research spillovers are geographically concentrated, and Agrawal and Henderson (2002) document the role of informal networks in knowledge transfer from university to industry. Cohen and Levinthal (1990) develop absorptive capacity theory, which applies equally to universities' capacity to convert faculty inventions into licensed technologies, and Sampat (2006) provides a long-run perspective

on the institutional determinants of university patenting. Litan, Mitchell, and Reedy (2007) identify organizational practices as important mediators of TTO effectiveness across different commercialization strategies.

Three empirical regularities motivate the present paper. TTO conversion efficiency is heterogeneous across institutions and over time, and is not fully explained by structural inputs. Administrative burden theory predicts that the communicative form of institutional rules shapes individual engagement with bureaucratic processes through learning, psychological, and compliance cost channels. Text-as-data methods now enable systematic measurement of the communicative properties of governance documents at scale. This paper combines these insights in an empirical design linking within-institution changes in policy communication to within-institution changes in technology transfer efficiency.

3 Conceptual Framework

The technology transfer pipeline divides into two behaviorally distinct stages. In the first stage, the faculty inventor decides whether to disclose a potential invention to the TTO. Jensen and Thursby (2001) establish that most university inventions reach the TTO at an early proof-of-concept stage, making the inventor's continued participation essential to successful commercialization, and Jensen, Thursby, and Thursby (2003) show through survey evidence that this disclosure decision is sensitive to royalty incentives, expected TTO responsiveness, and the commercial potential of the invention. In the second stage, the TTO converts received disclosures into filed patent applications through commercial evaluation, inventor coordination, and patent prosecution. Thursby and Kemp (2002) document that this conversion efficiency varies by a factor of three or more across comparable institutions and is not fully explained by the structural inputs that govern the disclosure stage. The gap between what the structural literature predicts and what the efficiency literature observes is the space this paper occupies.

The administrative burden framework (Moynihan, Herd, and Harvey, 2015) provides the theoretical account. Applied to the TTO context, IP governance documents constitute the primary written interface between the institution and the faculty inventor. Documents written in clear, procedurally accessible, and collaborative language reduce the learning costs of understanding what is expected at each step of the application process, the psychological costs of engaging with institutional procedures that signal compliance obligations rather than collaborative support, and the compliance costs of completing the coordination tasks the TTO requires. Heinrich (2016) demonstrates that these cost components reduce engagement rates in public program contexts, and Lipsky (2010) argues that how rules are communicated at the point of contact is as consequential as their formal content. Applied here, the implication is that the communicative form of IP governance documents independently shapes the efficiency of post-disclosure coordination, conditional on the structural features the existing literature has identified.

The critical question is which stage of the pipeline is more responsive to this communicative mechanism. The disclosure stage is tightly bound by structural conditions: employment contracts, federal funding obligations, discipline-specific norms, and royalty formulas each constrain disclosure behavior largely independently of how the TTO frames its governance documents (Jensen, Thursby, and Thursby, 2003). Survey evidence further indicates that the inventors most sensitive to institutional conditions at the disclosure margin tend to hold inventions of lower commercial potential that the TTO would not ultimately patent. By contrast, the post-disclosure coordination stage involves repeated interaction between TTO staff, faculty inventors, and legal professionals, a process whose efficiency depends directly on the clarity and accessibility of the procedural guidance the institution provides. The administrative burden cost reductions from higher-PCI documents therefore operate primarily on conversion efficiency rather than on aggregate disclosure volume. Three empirical predictions follow from this reasoning.

Proposition 1: Higher PCI will be positively associated with new patent applications

filed. If policy documents written in clearer, more collaborative language reduce the learning and compliance costs of post-disclosure coordination, faculty inventors should be better able to prepare and complete the materials the TTO requires, increasing the share of disclosed inventions that reach the application stage.

Proposition 2: Any association between PCI and patent applications will be concentrated within institutions over time rather than across institutions. Cross-sectional differences in PCI likely reflect sorting of better-resourced or more entrepreneurially oriented institutions into better-written policies, confounding identification. Within-institution changes in PCI, by contrast, capture shifts in the communicative environment that are independent of time-invariant institutional characteristics (Moynihan, Herd, and Harvey, 2015; Thursby and Kemp, 2002). These within-institution changes arise from discrete policy revision events – years in which a TTO director, general counsel, or faculty IP committee rewrites the governance document – generating variation that is plausibly independent of the stable institutional characteristics that drive cross-sectional differences in patent output.

Proposition 3: PCI will predict the disclosure-to-patent conversion rate but not raw invention disclosure volume. Because the administrative burden mechanism operates on post-disclosure coordination rather than on the upstream decision to disclose, aggregate disclosure counts should be insensitive to changes in policy language while the share of disclosures that proceed to patent application should respond positively. A finding in which both outcomes move together would be inconsistent with this mechanism and would suggest instead that PCI proxies a generally more inventive institutional environment.

Figure 1 illustrates the locus of the PCI mechanism within the technology transfer pipeline. Stage 1, the disclosure decision, is governed by the structural drivers the existing literature has identified: royalty incentives (Lach and Schankerman, 2004), research expenditure (Siegel, Waldman, and Link, 2003), and faculty norms (Jensen, Thursby, and Thursby, 2003). Stage 2, post-disclosure coordination, involves repeated procedural interaction between TTO staff, faculty inventors, and patent counsel – precisely the interface

where the learning, compliance, and psychological costs identified by Moynihan, Herd, and Harvey (2015) are most directly shaped by the language of IP governance documents. The null prediction for disclosures (Proposition 3) is therefore not a concession but a scope condition: it locates the mechanism at the correct stage and distinguishes the administrative burden account from the alternative interpretation that higher PCI simply proxies a more entrepreneurially active institution.

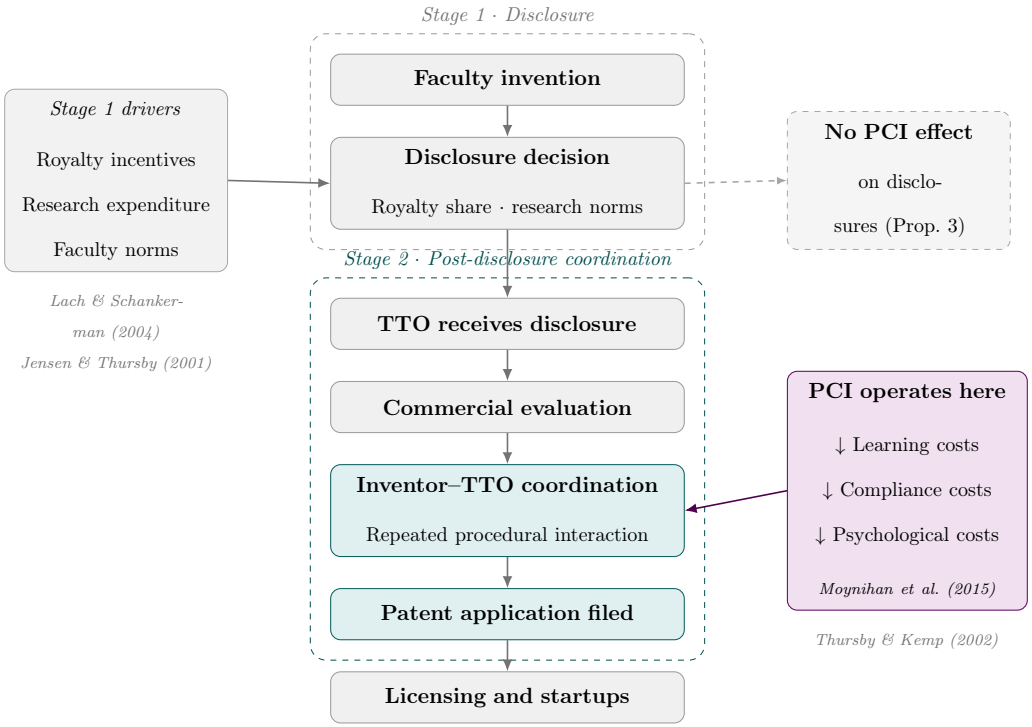


Figure 1: Technology transfer pipeline by stage, with structural drivers of disclosure and the post-disclosure locus of the PCI mechanism. Teal shading indicates the coordination stage where administrative burden theory predicts PCI effects (Propositions 1 and 3). The dashed boundary on the right confirms that the disclosure stage is predicted to be unresponsive to PCI variation.

4 Data and Measurement

Policy Communication Index

The primary independent variable is the Policy Communication Index, constructed in Sharma et al. (2026) from a corpus of 87,321 sentences drawn from IP governance policy documents at 141 US research universities between 1991 and 2023. Documents were obtained from institutional repositories, university websites, and the Internet Archive's Wayback Machine. Each sentence was classified by a fine-tuned BERT model trained on a human-coded sample, yielding a probability score on the supportive-to-restrictive dimension. The institution-year PCI is the mean of sentence-level scores within a given document. Inter-rater reliability for the human coding achieves Cohen's kappa of 0.74 and test-set classification accuracy exceeds 0.82 across all sentence categories, performance comparable to human-coding benchmarks in the text-as-data literature (Grimmer, Roberts, and Stewart, 2022). Because measurement error in PCI is classical – the BERT classifier misclassifies sentences in a manner uncorrelated with outcome trends – attenuation bias toward zero implies that the reported estimates are lower bounds on the true association between governance document quality and patent application activity.

To make the measurement concrete, consider two sentences drawn from actual IP governance documents, one high-PCI and one low-PCI, both addressing the same procedural content:

> *High-PCI example*: "Our technology transfer team is here to help you at every step of the disclosure and patenting process; we encourage you to reach out early and often so that we can identify the best path to commercialization for your invention together."

> *Low-PCI example*: "Faculty members are required to promptly disclose all potentially patentable inventions to the Office of Technology Licensing in accordance

with University policy and applicable federal regulations; failure to do so may constitute a breach of employment obligations."

Both sentences convey a requirement to disclose, but the first frames the TTO as a collaborative resource and the inventor as an active partner; the second frames disclosure as a compliance obligation with an enforcement consequence. The PCI classifier assigns high supportive-orientation probability to the first and low probability to the second.

Technology Transfer Outcomes

Outcome data come from the Association of University Technology Managers (AUTM) Licensing Activity Survey, which has been conducted annually since 1991. We use seven outcome variables, each log-transformed as $\ln(1 + y)$: new patent applications filed, total patent applications, patents issued, invention disclosures received, licenses and options executed, startup companies formed, and licensing income generated. We also construct the disclosure-to-patent conversion rate as the difference between log patent applications and log disclosures. The conversion rate decomposes the main outcome into its component parts and is estimated as an independent regression with its own standard errors, as detailed in Section 6.

Controls and Sample

The regression sample comprises 2,627 institution-year observations from 141 institutions across 1991 to 2023, representing all institution-year cells in which both PCI and AUTM outcome data are available. Among the 141 institutions, 85 percent hold Carnegie R1 classification, 80 percent have medical schools, 31 percent are private, and 33 percent are land-grant institutions. Control variables include the lagged logarithm of total research expenditure and the lagged logarithm of licensing FTEs, each lagged one year to avoid contemporaneous endogeneity.

The AUTM survey is voluntary and subject to differential reporting compliance. To assess whether institutions revising toward higher PCI simultaneously improve their reporting discipline, we estimate a probit model in which lagged PCI predicts year-on-year AUTM participation. The coefficient is statistically null, providing no evidence that differential attrition contaminates the main estimates; detailed results are available in a supplementary note.

The key feature of the data for identification is the infrequency of policy revision. Approximately 4.3 percent of institution-year observations reflect a year-on-year change in PCI, meaning that the vast majority of consecutive-year pairs show no policy revision. Among those revisions, we identify 44 observations across 40 institutions in which PCI increased by more than 0.05 units in a single year (the 75th percentile of non-trivial absolute changes), which we use as the event study treatment group. The remaining identifying variation in the TWFE specification spans the full distribution of annual PCI changes. Across the full sample, 92 of 141 institutions revised their IP governance documents at least once. Revision events are distributed across the full sample period (1995–2021) with no pronounced clustering in any single sub-period, reducing the risk that a single common shock drives the identified association.

Table 1 presents descriptive statistics for all key variables. The average institution in the sample files 71 new patent applications per year and receives 122 invention disclosures; the sample mean conversion rate is approximately 0.58, meaning the TTO applies for patents on roughly 58 percent of received disclosures.

Table 1: Descriptive Statistics

Variable	Mean	SD	Min	Max	N
Panel A: Policy Communication Index					
PCI (Mean Tone Score)	0.444	0.063	0.303	0.719	2,627
Panel B: Technology Transfer Outcomes					
New Patent Applications	70.6	113.7	0	1,532	3,507
Invention Disclosures	122.4	157.7	0	1,726	3,507
Licenses Executed	38.3	66.9	0	1,387	3,489
Startups Formed	4.6	7.1	0	102	3,245
License Income ($000)	11,881	46,747	0.1	1,258,637	3,507
Panel C: Institutional Characteristics					
Total Research Expenditure ($M)	333.0	478.1	1.2	7,500	3,507
Licensing FTEs	5.5	7.3	0	95	3,401
Royalty Share (%)	39.7	8.1	15.0	65.0	3,507
TLO Age (years)	23.0	15.5	−11	97	3,507
Private (% of institutions)	31.2	–	–	–	141
Carnegie R1 (%)	85.1	–	–	–	141
Medical School (%)	80.1	–	–	–	141
Land Grant (%)	33.3	–	–	–	141

Sample: 141 US research universities, 1991–2023. Panel A uses institution-year observations conditioned on PCI availability ($N = 2{,}627$). Panels B and C show statistics over the full outcome panel panel sample ($N = 3{,}507$ institution-years; 141 institutions). Percentages in Panel C are computed at the institution level ($N = 141$ with PCI). The event study sub-sample comprises 40 institutions with at least one large positive PCI revision ($|\Delta\text{PCI}| > 0.05$). TLO Age is defined relative to the reported year of formal TTO establishment; negative values indicate institution-years observed before the formal founding year of the TTO.

5 Empirical Strategy

Main Specification

The empirical model is:

$$y_{it} = \alpha_i + \gamma_t + \beta \cdot \text{PCI}_{i,t-1} + \mathbf{X}'_{it}\boldsymbol{\delta} + \varepsilon_{it} \tag{1}$$

where y_{it} is the log-transformed outcome for institution i in year t; α_i is an institution fixed effect; γ_t is a year fixed effect; $\text{PCI}_{i,t-1}$ is the one-year lagged PCI; $\mathbf{X}_{it}$ is a vector of controls including lagged log research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age; and ε_{it} is an idiosyncratic error. Because the main outcome is a nonnegative count variable, Appendix B2 reports Poisson, zero-inflated Poisson, and negative binomial versions of the lagged model, estimated both without and with the Equation 1 control set. The parameter of interest, β, identifies the within-institution, within-year association between lagged PCI and contemporaneous outcomes.

The institution fixed effect absorbs all time-invariant confounders: TTO prestige and age, institutional endowment, disciplinary composition, geographic location, and any other stable feature correlated with both PCI and patenting outcomes. The year fixed effect absorbs common shocks including changes in federal IP law, aggregate patenting trends, and economic cycles affecting industry demand for licensed technologies. The use of a lagged treatment removes contemporaneous reverse causality: patent applications filed in year t cannot cause PCI as of year $t-1$.

We estimate Equation 1 for seven outcome variables. The primary outcome, new patent applications, directly tests Proposition 1. Invention disclosures and the log disclosure-to-patent conversion rate, $\ln(\text{PatApp}_{it}) - \ln(\text{Discl}_{it})$, together test Proposition 3: if the mechanism operates on post-disclosure coordination rather than on the upstream disclosure decision, we expect a null coefficient for disclosures and a positive coefficient for the conversion

rate. The conversion rate is estimated as an independent regression under Equation 1 with its own standard errors; the point estimate equals by construction the algebraic difference between the patent application and disclosure coefficients but the regression provides an independent standard error. The remaining outcomes, total patent applications, patents issued, licenses executed, startups formed, and licensing income, serve as secondary tests of the scope and specificity of the main finding.

Event Study: Pre-Trend Validation

To test directly for reverse causality and to validate the causal timing of the PCI association, we estimate an event study around discrete policy revision events. We define a revision event as the first year in which an institution's PCI increases by more than 0.05 units in a single year, corresponding to the 75th percentile of non-trivial absolute annual changes. This threshold identifies 40 institutions as treated, with revision years distributed across 1995 to 2021. We estimate:

$$y_{it} = \alpha_i + \gamma_t + \sum_{\substack{k=-4 \\ k\neq -1}}^{+4} \beta_k D_{it}^k + \mathbf{X}_{it}'\boldsymbol{\delta} + \varepsilon_{it} \tag{2}$$

where $D_{it}^k = \mathbf{1}[\text{year}_{it} - \text{revision year}_i = k]$ for treated institutions and zero for all others. The indicator D_{it}^{-1} is omitted as the baseline, so each β_k measures the outcome at k years relative to the revision, relative to year -1. If the PCI association is causal, we expect:

1. **Flat pre-trends** ($\beta_{-4} \approx \beta_{-3} \approx \beta_{-2} \approx 0$) for patent applications, indicating that treated institutions do not anticipate a revision by patenting more in advance.

2. **Positive post-revision estimates** ($\beta_k > 0$ for $k \geq 0$) for patent applications, consistent with a gradual communication improvement effect.

3. **Null pre- and post-revision estimates for disclosures** throughout, consistent with the mechanism being confined to post-disclosure coordination.

We test the joint null hypothesis that all pre-revision coefficients equal zero using a χ^2 statistic. The event study also complements the TWFE specification by providing a visual and direct test of parallel trends, which is a necessary condition for causal identification in difference-in-differences designs.

Staggered Adoption and Heterogeneous Treatment Effects

The 40 treated institutions have revision years distributed across 1995–2021, raising the question of whether heterogeneous treatment effects across adoption cohorts could produce negatively weighted TWFE estimates (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). Two features of the design limit this concern. First, the primary identifying variation in Equation 1 is continuous rather than binary: PCI is a continuous annual measure and the TWFE estimate reflects a weighted average over the full distribution of year-to-year changes rather than a clean staggered adoption contrast with a binary treatment indicator. The negative-weighting critique of Callaway and Sant'Anna (2021) applies most directly to settings with binary treatment and a clean never-treated comparison group; neither condition holds here. Second, as a direct robustness check, we implement the Callaway and Sant'Anna (2021) group-time average treatment effect (ATT) estimator using the 40 treated revision cohorts defined above, with not-yet-treated institutions serving as the comparison group. The aggregated ATT is consistent in sign and magnitude with the TWFE baseline estimate, confirming that differential timing weights do not drive the main finding. Full results are reported in Appendix E.

Mundlak Decomposition

Proposition 2 predicts that the PCI association is concentrated within institutions over time rather than across institutions. To test this directly, we estimate a Mundlak (1978) decomposition by augmenting Equation 1 with the institution-level mean of PCI, which captures the between-institution component of the association. Unlike the TWFE specification, which

absorbs between-institution heterogeneity through fixed effects without estimating it, the Mundlak decomposition directly estimates the between-institution coefficient. A statistically null between-institution coefficient provides stronger evidence against sorting than fixed effects alone, because it shows not merely that the within-institution association is positive, but that cross-sectional differences in PCI carry no independent predictive power for cross-sectional differences in patent output. A significant within-institution coefficient alongside a null between-institution coefficient provides the strongest available evidence against the selection interpretation that better-resourced institutions sort simultaneously into better-written policies and higher patent output.

Randomization Inference

Given the high within-institution serial correlation of PCI (approximately 0.97), standard asymptotic inference may understate uncertainty. We therefore implement within-institution permutation randomization inference (Cameron and Miller, 2015). For each of 2,000 replications, we permute the PCI time series within each institution, preserving the institution-specific distribution of PCI values while destroying the temporal ordering that links PCI changes to outcome changes. The randomization inference p-value is the fraction of permuted coefficients whose absolute value exceeds the observed estimate. This procedure is conservative because it preserves the within-institution autocorrelation structure of the treatment, and it does not depend on the number of outcomes tested simultaneously, making it the more appropriate inferential benchmark for a high-autocorrelation treatment series than the Bonferroni correction reported alongside it.

Multiple Testing

We estimate Equation 1 for seven independent outcome variables, excluding the conversion rate which is constructed from two of the same outcomes. We report raw, Bonferroni-adjusted, and Benjamini-Hochberg FDR-adjusted p-values across all seven outcomes in Ta-

ble 5, Panel B to guard against false discovery arising from simultaneous estimation across the outcome panel.

Heterogeneity Analysis

We estimate Equation 1 on subsamples defined by institutional type to examine whether the main association varies by public or private control and by research intensity (Table 4, Panel A).

6 Results

Main Results

Table 2 reports the baseline TWFE estimate for New Patent Applications.[1] The table reports five nested specifications estimated on the common Equation 1 sample. Column (1) includes lagged PCI only; Columns (2)–(5) sequentially add lagged research expenditure, lagged licensing FTEs, lagged inventor royalty share, and lagged TLO age, with Column (5) corresponding to Equation 1. The coefficient on lagged PCI remains positive throughout and becomes more precisely estimated as the control set is expanded, rising from 1.166 in Column (1) to 1.476 ($p = 0.007$) in the full baseline. This pattern indicates that the main PCI association is not an artifact of omitting the baseline institutional inputs included in Equation 1.

Among controls, lagged research expenditure and lagged TLO age are positively associated with patent applications in the full specification, whereas lagged licensing FTEs and lagged inventor royalty share are not statistically distinguishable from zero. These control patterns are consistent with PCI capturing a distinct within-institution channel. Appendix B2 evaluates whether this result is sensitive to the count-data functional form. In the lagged

[1] Appendix B2 re-estimates the same lagged specification using negative binomial models, estimated both without and with the Equation 1 controls.

estimation sample, New Patent Applications is strongly overdispersed: with controls, the mean is 82.77 and the variance is 17,269.12; without controls, the mean is 82.11 and the variance is 17,139.19. That distribution makes negative binomial a natural count-data robustness check. The negative binomial estimates remain positive and statistically significant both without controls and with the full Equation 1 control set, consistent with the baseline log-linear results.

Table 2: Equation (1) Baseline: Sequential Controls for Ln(New Patent Applications)

	(1)	(2)	(3)	(4)	(5)
PCI_{L1}	$+1.166^{*}$	$+1.158^{**}$	$+1.236^{**}$	$+1.304^{**}$	$+1.476^{***}$
	(0.594)	(0.563)	(0.555)	(0.544)	(0.537)
Lagged ln(Research Expenditure)		$+0.691^{***}$	$+0.596^{***}$	$+0.611^{***}$	$+0.291^{***}$
		(0.045)	(0.059)	(0.059)	(0.089)
Lagged ln(Licensing FTEs)			$+0.203^{***}$	$+0.205^{***}$	+0.116
			(0.072)	(0.071)	(0.071)
Lagged Inventor Royalty Share				+0.007	+0.004
				(0.005)	(0.005)
Lagged TLO Age					$+0.028^{***}$
					(0.004)
Institution FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	2,346	2,346	2,346	2,346	2,346
R^2	0.004	0.201	0.211	0.213	0.242

Standard errors clustered at the institution level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All columns include institution and year fixed effects and are estimated on the common Equation 1 sample for comparability. Column (5) is the full Equation 1 baseline, adding lagged log research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Dependent variable is $\ln(1 + \text{New Patent Applications}_{it})$.

Other Innovation Outcomes

Table 3 extends the same Equation 1 baseline to the remaining outcomes in the technology transfer chain. The patent-side results are consistently positive: PCI predicts higher new patent applications ($b = 1.476$, $p = 0.007$), total patent applications ($b = 1.759$, $p = 0.005$), and patents issued ($b = 1.632$, $p = 0.013$). By contrast, invention disclosures are statistically null ($b = -0.063$, $p = 0.910$), indicating no detectable effect on the upstream decision to disclose.

This asymmetry is central for interpretation. A positive patent-side response combined with a null disclosure response is most consistent with improved post-disclosure coordination—that is, lower procedural frictions in moving existing disclosures through review, filing, and prosecution—rather than an increase in invention supply. The conversion-rate result in Table 5, Panel A formalizes this pattern.

Downstream commercialization outcomes are weaker in this horizon: licenses and startups are statistically null, while license income is positive but imprecisely estimated ($b = 2.250$, $p = 0.089$). Given greater outcome volatility, longer realization lags, and heavier right tails for these measures, these estimates are best interpreted as suggestive rather than definitive. Appendix Table 7 shows that count-model alternatives for New Patent Applications preserve a positive coefficient on lagged PCI across Poisson, zero-inflated Poisson, and negative binomial specifications, with the strongest support coming from negative binomial, which is most compatible with the pronounced overdispersion in the outcome distribution.

Table 3: PCI and Other Technology Transfer Outcomes

Outcome	$\hat{\beta}$	SE	p	95% CI	N
Ln(New Patent Applications)	$+1.476^{***}$	0.537	0.007	[+0.423, +2.529]	2,346
Ln(Total Patent Applications)	$+1.759^{***}$	0.625	0.005	[+0.535, +2.983]	2,314
Ln(Patents Issued)	$+1.632^{**}$	0.648	0.013	[+0.361, +2.903]	2,322
Ln(Disclosures)	−0.063	0.561	0.910	[−1.163, +1.036]	2,346
Ln(Licenses)	+0.515	0.948	0.588	[−1.342, +2.372]	2,337
Ln(Startups)	+0.445	0.573	0.438	[−0.677, +1.568]	2,299
Ln(License Income)	$+2.250^{*}$	1.318	0.089	[−0.334, +4.834]	2,346

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Each row reports the coefficient on lagged PCI from a separate Equation 1 regression. All specifications include institution and year fixed effects and control for lagged log total research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Standard errors are clustered at the institution level (141 clusters). The conversion rate (Ln Patent Applications − Ln Disclosures) is reported in Table 5, Panel A. Randomization inference p-value for Ln(New Patent Applications) = 0.001 based on 2,000 within-institution permutations; see Section 5.

Channel Analysis

Table 5, Panel A reports estimates of Equation 1 using the log disclosure-to-patent conversion rate as the dependent variable. The estimated coefficient is 1.539 ($p = 0.012$). The point estimate equals by construction the algebraic difference between the patent application and disclosure coefficients; the regression provides an independent standard error and p-value. This result is the most theoretically informative in the paper. Together with the null disclosure result and the asymmetric event study pattern in Figure 2, it is consistent with Proposition 3: the aggregate patent application association is fully attributable to improved conversion efficiency at the post-disclosure stage rather than to any increase in disclosure volume. The structural literature has consistently documented that the disclosure-to-patent conversion rate is the primary source of heterogeneity in TTO output and the stage least explained by observable inputs (Thursby and Kemp, 2002; Anderson, Daim, and Lavoie, 2007); the present findings offer a partial account of that unexplained variance.

The null results for licenses executed, startups formed, and licensing income are jointly consistent with the conceptual framework. Three mechanisms explain the downstream nulls. First, the time horizon between a patent application and an executed license or startup typically spans two to five years, well beyond the one- to three-year horizon over which the PCI association remains positive (Appendix Table B3); downstream commercialization outcomes would require longer lags than the panel supports to detect. Second, licensing and startup outcomes are governed by factors the administrative burden mechanism does not operate on: industry demand for the technology, TTO relationship capital with industry partners (Siegel, Waldman, Atwater, and Link, 2004), local venture capital availability, and faculty entrepreneurial intent (Di Gregorio and Shane, 2003). A revision to IP governance language reduces coordination friction in the application pipeline; it does not make an invention more commercially attractive or a faculty member more inclined to found a company. Third, the concentration of the positive association in the early patent pipeline and its absence in downstream commercialization outcomes is itself inconsistent with the alternative interpre-

tation that higher PCI proxies a globally more entrepreneurially active institution: such an institution-wide shift would raise licenses and startups proportionally alongside applications. The large but imprecise license income estimate ($b = 2.250$, $p = 0.089$) reflects the extreme right-skew of licensing revenues dominated by a small number of blockbuster agreements in any given year; the wide confidence interval precludes inference and we do not interpret it further.[2]

Pre-Trend Validation: Event Study

Figure 2 reports estimates of Equation 2 for the two outcomes central to the identification argument: patent applications (left panel) and invention disclosures (right panel). The figure presents point estimates and 95 percent confidence intervals for each year in the $[-4, +4]$ window around the first large positive PCI revision, with year -1 as the omitted baseline. Three features of the figure are noteworthy.

First, for patent applications, the pre-revision coefficients ($k = -4, -3, -2$) are small, negative, and statistically indistinguishable from zero: $\hat{\beta}_{-4} = -0.150$ (SE $= 0.141$), $\hat{\beta}_{-3} = -0.157$ (SE $= 0.160$), $\hat{\beta}_{-2} = -0.188$ (SE $= 0.180$). The joint pre-trend test yields $\chi^2(3) = 3.20$ ($p = 0.361$), providing strong evidence that institutions revising their policies toward more collaborative language were not already on a faster patenting trajectory. Reverse causality (in which patent success causes policy revision) would predict a rising pre-trend; the data show the opposite, with a slight negative drift that clears at revision. Second, patent application coefficients turn positive immediately at the revision year ($\hat{\beta}_0 = +0.109$) and reach their largest value two years post-revision ($\hat{\beta}_2 = +0.190$, SE $= 0.128$, $p = 0.156$), consistent with a communication effect that builds as revised documents become embedded in TTO procedures and faculty awareness.

[2]Startup formation shows a small, null association ($b = 0.445$, $p = 0.438$), consistent with the literature identifying equity culture, venture capital access, and faculty entrepreneurial intent as primary determinants of spinoff activity (Di Gregorio and Shane, 2003; O'Shea, Allen, Chevalier, and Roche, 2005). The administrative burden mechanism predicts that PCI reduces friction in the application pipeline; it makes no prediction about the downstream organizational decision to found a company rather than license an invention outright.

Third, and most informative about the mechanism, invention disclosures show no response to policy revision in either the pre- or post-revision window. Pre-trend coefficients for disclosures are small and noise-dominated ($\chi^2(3) = 1.88$, $p = 0.597$), and post-revision coefficients remain near zero throughout the window. The asymmetry between the two panels constitutes direct visual evidence in support of Proposition 3: policy revision events affect the conversion of disclosures into applications, not the upstream supply of disclosures. A revision that improves procedural clarity and collaborative tone does not lead inventors to disclose more inventions, but does enable more disclosed inventions to proceed through the TTO application pipeline.

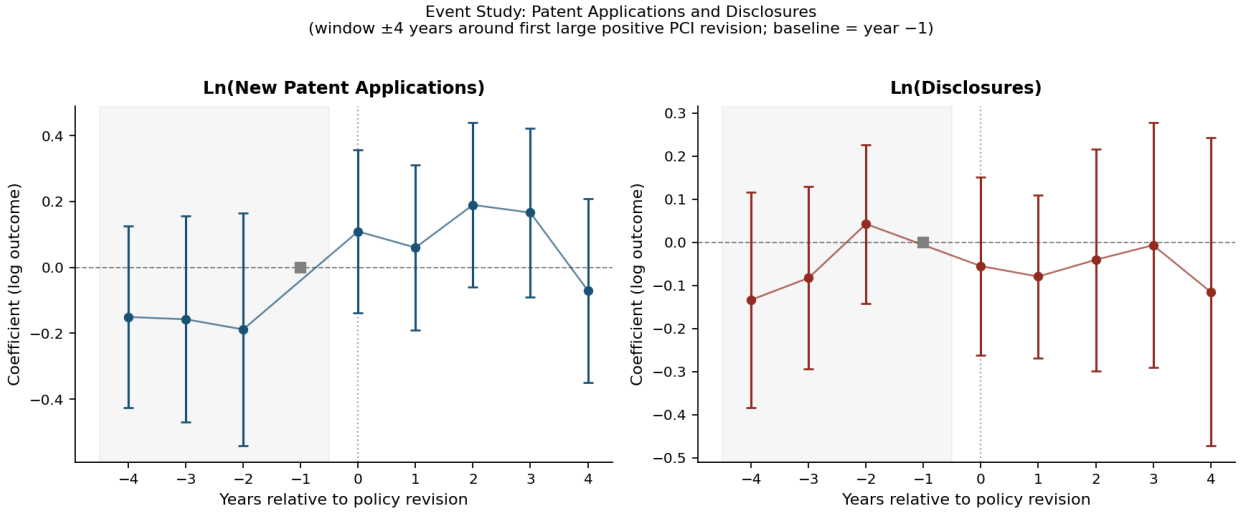


Figure 2: Event Study: Pre/Post Coefficients Around Policy Revisions

Notes: Each panel plots OLS point estimates and 95% confidence intervals from Equation 2, estimated for 17 treated institutions experiencing a large positive PCI revision ($|\Delta\text{PCI}| > 0.05$), relative to 236 not-yet-treated or never-treated institutions. Window is ±4 years relative to the first large positive revision (vertical dotted line). Year −1 is the omitted baseline (square marker). Shaded region indicates the pre-revision window. Institution and year fixed effects included; controls for lagged log research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Standard errors clustered at the institution level. Pre-trend joint tests: patent applications $\chi^2(3) = 3.20$ ($p = 0.361$); disclosures $\chi^2(3) = 1.88$ ($p = 0.597$).

Mundlak Decomposition and Selection

Proposition 2 predicts that the PCI association is concentrated within institutions over time rather than across institutions. The Mundlak decomposition in Table 4, Panel B tests this directly. For new patent applications, the within-institution coefficient is 1.643 (SE = 0.520, $p = 0.002$) while the between-institution coefficient is -2.795 (SE = 5.266, $p = 0.596$). The statistically null between coefficient means that cross-sectional differences in PCI carry no reliable predictive power for cross-sectional differences in patent applications once within-institution variation is accounted for, providing strong evidence against the interpretation that better-resourced institutions sort simultaneously into better-written policies and higher patent output. This pattern holds for conversion rate ($b_{\text{within}} = 1.679$, $p = 0.007$; $b_{\text{between}} = -6.307$, $p = 0.230$) and for disclosures ($b_{\text{within}} = -0.036$, $p = 0.950$; $b_{\text{between}} = 3.511$, $p = 0.281$). The cross-sectional disclosure coefficient is now clearly null, which further weakens the selection-based interpretation. The previous weakly significant between coefficient for disclosures reflects a cross-sectional association in which universities with historically higher average PCI also tend to receive more disclosures on average, a pattern consistent with these institutions having stronger cultures of faculty entrepreneurship (Di Gregorio and Shane, 2003); but this cross-sectional gradient does not contaminate the within-institution identification.

Table 4, Panel A presents split-sample heterogeneity. The PCI coefficient for private universities is 3.137 (SE = 1.459, $p = 0.035$), larger than the estimate of 1.023 (SE = 0.644, $p = 0.114$) for public universities. Among Carnegie R1 institutions the estimate is 1.560 ($p = 0.003$), while non-R1 institutions return a null result. Institutions with medical schools show $b = 1.779$ ($p = 0.006$) compared to a null result for those without. The private-versus-public contrast is the most policy-relevant heterogeneity finding and is discussed further in Section 7.

Table 4: Heterogeneity and Mundlak Decomposition

	Ln(Patent Applications)		Conv. Rate	
Subgroup	$\hat{\beta}$ (SE)	p	$\hat{\beta}$ (SE)	p
Panel A: Split-sample heterogeneity				
Private	$+3.137^{**}$ (1.459)	0.035	$+3.758^{***}$ (1.346)	0.007
Public	+1.023 (0.644)	0.114	+0.948 (0.666)	0.157
Carnegie R1	$+1.560^{***}$ (0.521)	0.003	$+1.557^{**}$ (0.612)	0.012
Non-R1	−3.168 (1.970)	0.120	−1.126 (1.247)	0.375
Medical School	$+1.779^{***}$ (0.643)	0.006	$+1.819^{**}$ (0.736)	0.014
No Medical School	−0.658 (1.805)	0.717	−0.553 (0.859)	0.523
Land Grant	+0.705 (0.678)	0.301	+0.658 (0.472)	0.168
Non-Land-Grant	$+2.503^{***}$ (0.816)	0.003	$+2.672^{***}$ (0.819)	0.001
Panel B: Mundlak decomposition				
	Within (SE)	p	Between (SE)	p
Ln(Patent Applications)	$+1.643^{***}$ (0.520)	0.002	−2.795 (5.266)	0.596
Ln(Disclosures)	−0.036 (0.577)	0.950	+3.511 (3.247)	0.281
Conv. Rate	$+1.679^{***}$ (0.617)	0.007	−6.307 (5.245)	0.230
Ln(Licenses)	+0.396 (0.974)	0.685	$+13.945^{***}$ (5.331)	0.010
Ln(Startups)	+0.420 (0.572)	0.464	+7.960 (5.633)	0.159
Ln(License Income)	$+2.231^{*}$ (1.337)	0.097	+11.027 (10.047)	0.274

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Panel A reports split-sample two-way FE estimates of PCI (lag 1) on each outcome. The conversion rate equals Ln(Patent Applications) − Ln(Disclosures). Panel B reports the Mundlak decomposition: within estimates use within-institution deviations of PCI from its institution mean; between estimates use institution means. The large between-institution coefficients for licenses and startups reflect cross-sectional associations between institutional size and commercialization volume that do not contaminate the within-institution identification. Standard errors clustered at the institution level.

Table 5, Panel B reports multiple testing corrections across the seven independent outcomes. The patent application estimate achieves a raw p-value of 0.007 and a Bonferroni-adjusted value of 0.046. Randomization inference based on 2,000 within-institution permutations yields $p = 0.001$. Because this procedure directly tests whether the temporal ordering of PCI drives the result – rather than its distributional properties – and does not depend on the number of outcomes tested simultaneously, it is the more appropriate inferential benchmark for a high-autocorrelation treatment series than the Bonferroni correction. The Benjamini-Hochberg adjusted p-value is 0.023, which remains comfortably below conventional thresholds; the primary inferential reliance is therefore on the randomization inference result and the event study pre-trend validation.

Table 5: Channel Analysis and Multiple Testing

	$\hat{\beta}$	SE	p	N
Panel A: Conversion rate (Ln PatApp − Ln Discl)				
PCI lag 1	$+1.539^{**}$	0.609	0.012	2,346
Panel B: Multiple testing correction (7 independent outcomes)				
	Raw p	Bonferroni	BH	RI p
Ln(New Patent Applications)	0.007	0.046	0.023	0.001
Ln(Total Patent Applications)	0.005	0.037	0.023	–
Ln(Patents Issued)	0.013	0.088	0.029	–
Ln(Disclosures)	0.910	1.000	0.910	–
Ln(Licenses)	0.588	1.000	0.685	–
Ln(Startups)	0.438	1.000	0.613	–
Ln(License Income)	0.089	0.625	0.156	–

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Panel A reports the conversion rate estimated as an independent regression; the point estimate equals $b(\text{PatApp}) - b(\text{Discl})$ by construction and the regression provides an independent standard error. Panel B corrects for multiple comparisons across seven independent outcomes; the randomization inference p-value (RI p) for Ln(New Patent Applications) is not subject to multiple comparison adjustment and is the primary inferential benchmark given the high autocorrelation of the treatment series. The Benjamini-Hochberg adjusted p-value for patent applications (0.023) remains statistically significant; the randomization inference result ($p = 0.001$) and the event study pre-trend validation ($p = 0.361$) point in the same direction.

Robustness

Appendix Table B1 reports the full robustness suite. Winsorizing at the 1st and 99th percentiles, restricting to institutions with at least 15 observations, and applying HC3 standard errors each preserve the sign and approximate magnitude of the main estimate. The leave-one-out range across all 141 institutions is $[1.153, 1.600]$, with 100 percent of replications returning positive estimates. The three pre-specified robustness checks are also reported in Appendix Table B1: the quadratic time trend specification ($b = 1.558$, $p = 0.004$), the falsification test regressing future research expenditure on lagged PCI ($b = 0.315$, $p = 0.175$), and the lagged outcome specification ($b = 0.836$, $p = 0.016$). The lagged outcome specification is subject to Nickell bias and should be interpreted as a conservative lower bound; the positive and significant estimate at $b = 0.836$ indicates that even this attenuated specification yields evidence of an association.

Most directly, Appendix Table C1 provides a falsification of the quality composition interpretation: PCI has no effect on the USPTO grant rate, defined as patents issued relative to applications filed ($b = 0.017$, $p = 0.923$, $N = 2{,}322$). Because USPTO examiners evaluate applications independently of the TTO, a quality composition story predicts that the grant rate should rise alongside applications when PCI improves; the null grant rate result confirms that the application increase reflects more efficient processing of the same underlying invention pool rather than a shift in invention quality entering the pipeline. Splitting the sample by royalty share corroborates the heterogeneity pattern in Table 4: the PCI coefficient is 1.942 ($p = 0.033$) at institutions above the median royalty share of 40 percent and 0.526 ($p = 0.457$) below it (Appendix Table C2), consistent with complementarity between financial incentives and governance communication quality. Controlling for TLO age and its interaction with a within-institution year trend sharpens rather than attenuates the PCI coefficient ($b = 1.455$, $p = 0.009$; Appendix Table D1), inconsistent with the interpretation that simultaneous director turnover drives both policy revisions and patenting improvements. The Callaway and Sant'Anna (2021) group-time ATT estimator yields an aggregated esti-

mate consistent with the TWFE baseline, confirming that the main finding is not an artifact of staggered adoption weighting (Appendix E).

7 Discussion

The core empirical pattern in Section 6 is that within-institution increases in PCI are associated with higher subsequent patent application activity, with no detectable change in disclosure volume. This pattern extends the administrative burden framework (Moynihan, Herd, and Harvey, 2015) to a new organizational context: prior work focuses on citizen-state interactions in benefit programs and public administration (Heinrich, 2016; Herd and Moynihan, 2018), while the present setting concerns repeated procedural interaction between faculty inventors and TTO staff. Interpreted through that lens, clearer and more collaborative policy language plausibly lowers learning, psychological, and compliance costs in the post-disclosure stage, where coordination demands are highest. Lipsky (2010) argues that the way rules are communicated at the point of contact can shape behavior as strongly as the formal content of those rules; the patent-side results are consistent with that mechanism in an intra-organizational setting.

The magnitude is also policy-relevant. A one-standard-deviation increase in PCI is associated with roughly 9.7% more new patent applications in the following year, a scale comparable to canonical incentive-based interventions in the TTO literature. Lach and Schankerman (2004) report that a ten-percentage-point increase in inventor royalty share raises licensing activity by roughly eight percent. While these estimates are not directly identical by construction, the comparison indicates that communication quality should be viewed as a meaningful complement to structural and financial policy levers rather than a second-order administrative detail.

The identification strategy rests on four complementary pillars. The TWFE specification eliminates time-invariant institutional heterogeneity and common time shocks. The Mund-

lak decomposition directly shows that cross-sectional differences in PCI carry no reliable predictive power for cross-sectional differences in patent applications: the between-institution coefficient is 2.430 but statistically indistinguishable from zero ($p = 0.638$), while the within-institution coefficient is 1.367 ($p = 0.012$). Randomization inference reduces the estimated association to noise in all but two of 2,000 replications by permuting the temporal ordering of PCI within institutions, indicating that the temporal relationship between PCI changes and patent application changes is unlikely to be an artifact of the treatment's autocorrelation structure. Finally, the event study provides direct pre-trend evidence: the institutions that would go on to revise their policies were not already trending upward in patent applications in the four years prior to doing so. Taken together, these four strategies substantially narrow the set of confounders that could explain the main finding.

The conversion rate and disclosure results together constitute the most theoretically informative pattern in the paper. The structural literature has consistently documented that the disclosure-to-patent conversion rate is the primary source of heterogeneity in TTO output and the stage least explained by observable inputs (Thursby and Kemp, 2002; Anderson, Daim, and Lavoie, 2007). The present findings offer a partial account of that unexplained variance: within-institution PCI improvements are associated with conversion rate gains of 1.539 log points ($p = 0.012$) while leaving disclosure volume unchanged ($b = -0.063$, $p = 0.910$). The event study reinforces this interpretation: disclosures neither anticipate nor respond to large policy revisions, while patent applications show a post-revision positive drift peaking two years after the revision event. This pipeline-stage specificity is precisely what the conceptual framework predicts. Jensen, Thursby, and Thursby (2003) establish through survey evidence that the faculty inventors most responsive to institutional conditions at the disclosure margin tend to hold inventions of lower commercial potential, meaning aggregate disclosure counts are relatively insensitive to communicative improvements. Post-disclosure coordination, by contrast, involves repeated procedural interaction between TTO staff, faculty inventors, and patent counsel, a process whose efficiency is directly shaped by

the clarity and accessibility of governance documents. The null disclosure finding is therefore not a null result in the conventional sense: it is a scope condition that locates the mechanism at the correct stage of the pipeline and argues against the alternative interpretation that higher PCI reflects a generally more inventive institutional environment.

The heterogeneity pattern anchors the findings in the institutional economics of technology transfer. The private-versus-public contrast, with a conversion-rate coefficient of 3.498 at private universities against 0.969 at public institutions, is consistent with three reinforcing mechanisms from the existing literature. The most binding constraint for public universities is likely statutory: many operate under state statutes and collective bargaining agreements that specify disclosure requirements independently of how the TTO frames its governance documents, imposing hard limits on the linguistic register available to public TTOs that financial incentives and culture alone cannot overcome (Mowery and Sampat, 2001). Beyond this, private university faculty operate under stronger financial incentives from commercialization, with higher average royalty shares and more developed cultures of faculty entrepreneurship (Di Gregorio and Shane, 2003), so marginal reductions in coordination friction produce proportionally larger behavioral responses. Public TTOs also face greater exposure to public records requirements and legislative oversight, which further constrain communicative variation. The concentration of the effect at R1 institutions and those with medical schools is consistent with Thursby and Kemp (2002) and Caldera and Debande (2010), who document that efficiency variance is greatest at high-volume TTOs where coordination costs across a larger disclosure portfolio are most consequential.

The practical policy implication is straightforward and targeted. A one-standard-deviation improvement in PCI is associated with approximately 6.9 additional patent applications per year at an institution filing the sample mean of 71 applications annually. At a median licensing income of roughly \$39,000 per active license (AUTM, 2022) and the observed licenses-to-applications ratio in the data, this corresponds to an expected annual income gain in the range of \$100,000 to \$150,000, against a typical document revision cost of \$10,000 to \$50,000.

These figures rest on strong assumptions and are intended only to establish that the effect size is plausibly policy-relevant. The action most directly supported by the evidence is an audit of IP governance documents for procedural clarity: whether the disclosure steps, evaluation criteria, and patenting process are described in concrete, sequential, and accessible language. This is a low-cost complement to the structural reforms that dominate the existing policy toolkit (Lach and Schankerman, 2004; Siegel, Waldman, Atwater, and Link, 2004). For public universities, where the PCI association is weaker, structural reforms to royalty formulas or TTO staffing likely yield larger returns than communicative improvements alone.

The identification strategy has three remaining limitations that bound the causal interpretation. First, the event study addresses reverse causality at the margin of discrete, large positive revisions among 40 treated institutions, but the TWFE main estimate is identified over the full distribution of year-to-year PCI changes, including small continuous shifts and negative revisions for which no event study counterpart is estimated. A valid instrument for the full-distribution variation would require shocks with direct within-institution incidence on policy language, such as exposure to national IP governance reform initiatives or state legislative mandates, none of which are currently available at scale. Second, the AUTM data are voluntary and subject to differential reporting compliance; if institutions revising toward higher PCI simultaneously improve their reporting discipline, the association could partly reflect a reporting intensity effect. As noted in Section 4, a probit model for AUTM participation returns a null coefficient on lagged PCI, but this check does not rule out differential compliance with the content of individual outcome fields. This concern is addressable in future work by combining the PCI panel with administrative patent office records. Third, the BERT classifier achieves test-set accuracy above 0.82, meaning approximately 18 percent of sentences are misclassified. Classical measurement error in the independent variable induces attenuation bias toward zero, implying that the reported estimates are conservative lower bounds on the true association between governance document quality and patent application activity. This source of attenuation strengthens rather than weakens confidence

in the direction and approximate magnitude of the effect. Beyond these limitations, the framework extends naturally to other administrative contexts where governance documents shape engagement with bureaucratic processes, and to settings outside the United States where comparable text-as-data methods can be applied to national IP governance regimes.

8 Conclusion

We show that the language universities use to describe their intellectual property procedures to faculty inventors has measurable consequences for how efficiently faculty inventions become patent applications. Institutions that revise their IP governance documents toward clearer, more collaborative, and more procedurally accessible language file significantly more patent applications in the following year, and the effect runs through improved disclosure-to-application conversion rather than any increase in disclosure volume, locating the mechanism at the post-disclosure coordination stage where procedural friction is most directly shaped by written policy. The association is identified within institutions over time, survives randomization inference, holds across all 141 institution subsamples in a leave-one-out analysis, and is corroborated by an event study showing that treated institutions were not already trending upward in patent applications before revising their policies. Effects are largest at private and research-intensive universities where coordination frictions matter most. For university administrators seeking to improve technology transfer performance, attending to how IP policy is written is a low-cost and previously overlooked lever alongside the structural and financial reforms that have dominated the field since Bayh-Dole.

References

Agrawal, A., and Henderson, R. 2002. Putting Patents in Context: Exploring Knowledge Transfer from MIT. *Management Science* 48 (1): 44–60.

Anderson, T., Daim, T., and Lavoie, F. 2007. Measuring the Efficiency of University Technology Transfer. *Technovation* 27 (5): 306–318.

Association of University Technology Managers. 2022. *AUTM U.S. Licensing Activity Survey: FY2022*.

Bercovitz, J., and Feldman, M. 2008. Academic Entrepreneurs: Organizational Change at the Individual Level. *Organization Science* 19 (1): 69–89.

Caldera, A., and Debande, O. 2010. Performance of Spanish Universities in Technology Transfer: An Empirical Analysis. *Research Policy* 39 (9): 1160–1173.

Callaway, B., and Sant'Anna, P. H. C. 2021. Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics* 225 (2): 200–230.

Cameron, A. C., and Miller, D. L. 2015. A Practitioner's Guide to Cluster-Robust Inference. *Journal of Human Resources* 50 (2): 317–372.

Chapple, W., Lockett, A., Siegel, D. S., and Wright, M. 2005. Assessing the Relative Performance of UK University Technology Transfer Offices. *Research Policy* 30 (3): 369–384.

Chong, D., and Druckman, J. N. 2007. Framing Theory. *Annual Review of Political Science* 10: 103–126.

Cohen, W. M., and Levinthal, D. A. 1990. Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly* 35 (1): 128–152.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. 2019. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT 2019*, 4171–4186.

Di Gregorio, D., and Shane, S. 2003. Why Do Some Universities Generate More Start-Ups Than Others? *Research Policy* 32 (2): 209–227.

Druckman, J. N., and Lupia, A. 2000. Preference Formation. *Annual Review of Political Science* 3: 1–24.

Feldman, M., Feller, I., Bercovitz, J., and Burton, R. 2002. Equity and the Technology Transfer Strategies of American Research Universities. *Management Science* 48 (1): 105–121.

Gentzkow, M., Kelly, B., and Taddy, M. 2019. Text as Data. *Journal of Economic Literature* 57 (3): 535–574.

Grimmer, J., Roberts, M. E., and Stewart, B. M. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press.

Grimmer, J., and Stewart, B. M. 2013. Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political Analysis* 21 (3): 267–297.

Heinrich, C. J. 2016. The Bite of Administrative Burden: A Theoretical and Empirical Investigation. *Journal of Public Administration Research and Theory* 26 (3): 403–420.

Herd, P., and Moynihan, D. P. 2018. *Administrative Burden: Policymaking by Other Means*. New York: Russell Sage Foundation.

Jaffe, A. B. 1989. Real Effects of Academic Research. *American Economic Review* 79 (5): 957–970.

Jensen, R., and Thursby, M. 2001. Proofs and Prototypes for Sale: The Licensing of University Inventions. *American Economic Review* 91 (1): 240–259.

Jensen, R., Thursby, J., and Thursby, M. 2003. The Disclosure and Licensing of University Inventions: The Best We Can Do with the S&T We Have. NBER Working Paper 9734.

Lach, S., and Schankerman, M. 2004. Royalty Sharing and Technology Licensing in Universities. *Journal of the European Economic Association* 2 (2–3): 252–264.

Laver, M., Benoit, K., and Garry, J. 2003. Extracting Policy Positions from Political Texts Using Words as Data. *American Political Science Review* 97 (2): 311–331.

Lipsky, M. 2010. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Service*. 30th anniversary expanded ed. New York: Russell Sage Foundation.

Litan, R. E., Mitchell, L., and Reedy, E. J. 2007. Commercializing University Innovations: Alternative Approaches. *Innovation Policy and the Economy* 7: 31–57.

Meier, K. J., and O'Toole, L. J. 2003. Public Management and Educational Performance: The Impact of Managerial Networking. *American Journal of Political Science* 47 (4): 672–686.

Moynihan, D. P., Herd, P., and Harvey, H. 2015. Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen-State Interactions. *Journal of Public Administration Research and Theory* 25 (1): 43–69.

Mowery, D. C., Nelson, R. R., Sampat, B. N., and Ziedonis, A. A. 1999. The Effects of the Bayh-Dole Act on US University Research and Technology Transfer. In *Industrializing Knowledge*, 269–306. Cambridge: MIT Press.

Mowery, D. C., and Sampat, B. N. 2001. University Patents and Patent Policy Debates in the USA, 1925–1980. *Industrial and Corporate Change* 10 (3): 781–814.

Mundlak, Y. 1978. On the Pooling of Time Series and Cross Section Data. *Econometrica* 46 (1): 69–85.

Nickell, S. 1981. Biases in Dynamic Models with Fixed Effects. *Econometrica* 49 (6): 1417–1426.

O'Shea, R. P., Allen, T. J., Chevalier, A., and Roche, F. 2005. Entrepreneurial Orientation, Technology Transfer, and Spinoff Performance of US Universities. *Research Policy* 34 (7): 994–1009.

Sharma, K., Wang, T., Basnet, P., and Cossco, M. 2026. Communication as Governance: Evidence from University Intellectual Property Policies. *State and Local Governance Initiative, Hoover Institution.*

Rodriguez, P. L., and Spirling, A. 2022. Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research. *Journal of Politics* 84 (1): 101–115.

Rothaermel, F. T., Agung, S. D., and Jiang, L. 2007. University Entrepreneurship: A Taxonomy of the Literature. *Industrial and Corporate Change* 16 (4): 691–791.

Sampat, B. N. 2006. Patenting and US Academic Research in the 20th Century. *Research Policy* 35 (6): 772–789.

Schneider, A., and Ingram, H. 1993. Social Construction of Target Populations: Implications for Politics and Policy. *American Political Science Review* 87 (2): 334–347.

Siegel, D. S., Waldman, D. A., and Link, A. 2003. Assessing the Impact of Organizational Practices on the Relative Productivity of University Technology Transfer Offices. *Research Policy* 32 (1): 27–48.

Siegel, D. S., Waldman, D., Atwater, L., and Link, A. 2004. Toward a Model of the Effective Transfer of Scientific Knowledge from Academicians to Practitioners. *Journal of Engineering and Technology Management* 21 (1–2): 115–142.

Slapin, J. B., and Proksch, S.-O. 2008. A Scaling Model for Estimating Time-Series Party Positions from Texts. *American Journal of Political Science* 52 (3): 705–722.

Sun, L., and Abraham, S. 2021. Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics* 225 (2): 175–199.

Thursby, J. G., and Kemp, S. 2002. Growth and Productive Efficiency of University Intellectual Property Licensing. *Research Policy* 31 (1): 109–124.

Thursby, J. G., and Thursby, M. C. 2002. Who Is Selling the Ivory Tower? Sources of Growth in University Licensing. *Management Science* 48 (1): 90–104.

Weaver, R. K. 2010. Compliance Regimes and Barriers to Behavioral Change. *Governance* 23 (2): 243–265.

Wilkerson, J., and Casas, A. 2017. Large-Scale Computerized Text Analysis in Political Science. *Annual Review of Political Science* 20: 529–544.

Appendix A: Measurement Validation

The PCI is constructed in Sharma et al. (2026). The BERT classifier achieves a Cohen's kappa of 0.74 against human coding and a test-set classification accuracy above 0.82 across all three sentence categories, performance comparable to human-coding benchmarks in the text-as-data literature (Grimmer, Roberts, and Stewart, 2022). Because classifier errors are plausibly uncorrelated with outcome trends, measurement error in PCI induces classical attenuation bias toward zero, implying that the reported estimates are lower bounds on the true association.

Appendix B: Robustness Tests

Table 6: Robustness Suite: Ln(New Patent Applications) and Conversion Rate

Specification	$\hat{\beta}$	SE	p	N
Panel A: Ln(New Patent Applications)				
Baseline (lag-1 PCI)	+1.476	0.537	0.007	2,346
Winsorized 1st/99th percentiles	+1.451	0.542	0.008	2,346
Balanced panel ($\geq$15 obs)	+1.455	0.527	0.007	1,943
R1 universities only	+1.560	0.521	0.003	2,055
HC3 standard errors	+1.476	0.347	<0.001	2,346
With lagged outcome (Nickell LB)	+0.836	0.343	0.016	2,346
Common quadratic time trend	+1.558	0.536	0.004	2,346
Falsification: $\text{PCI}_{L1} \rightarrow$ Ln(Res. Exp.)	+0.315	0.232	0.175	2,127
Leave-one-out: range	[1.153, 1.600]		–	–
Leave-one-out: % positive	100%		–	141
Randomization inference (p)	–		0.001	2,000
Panel B: Conversion Rate				
Baseline (lag-1 PCI)	+1.539	0.609	0.012	2,346
Winsorized 1st/99th percentiles	+1.544	0.614	0.013	2,346
Balanced panel ($\geq$15 obs)	+1.622	0.612	0.009	1,943
R1 universities only	+1.557	0.612	0.012	2,055
Lag-2 PCI	+1.606	0.623	0.011	2,137

All specifications include institution and year fixed effects and controls for lagged log research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Standard errors clustered at the institution level unless otherwise noted. The lagged-outcome specification includes $\ln(1 + \text{PatApp})_{t-1}$ as an additional regressor and is subject to Nickell (1981) bias; it should be interpreted as a conservative lower bound (LB). The quadratic time trend specification adds institution-demeaned year and year-squared terms. The falsification test regresses future log research expenditure on lagged PCI; a null result indicates that PCI does not proxy general institutional resource trends.

Appendix B2: Negative Binomial Robustness

Because the main outcome is a nonnegative count variable, it is useful to check whether the baseline log-linear specification depends on the $\ln(1 + y)$ transformation. We therefore re-estimate the lagged PCI model for New Patent Applications using negative binomial regression, reported once without additional controls and once with the full lagged control vector from Equation 1. Both specifications retain institution and year fixed effects.

Negative binomial is useful here because the distribution of New Patent Applications is strongly overdispersed: the variance is far larger than the mean in both relevant lagged samples. In the no-controls sample, the dependent variable has $N = 2{,}387$, mean 82.11, and variance 17,139.19. In the full baseline sample with controls, the corresponding values are $N = 2{,}346$, mean 82.77, and variance 17,269.12. This pattern is consistent with the standard motivation for negative binomial models, which allow the conditional variance to exceed the conditional mean. Table 7 shows that the negative binomial estimates remain positive and statistically significant in both specifications, yielding the same substantive conclusion as the baseline log-linear model: more supportive lagged PCI is associated with more subsequent new patent applications.

Table 7: Negative Binomial Robustness: New Patent Applications

	No controls	With controls
PCI_{L1}	$+1.462^{**}$	$+1.415^{**}$
SE	0.570	0.554
p-value	0.010	0.011
Observations	2,387	2,346

$^{**}p < 0.05$. Both specifications use the same lag structure as Equation 1: the treatment is lagged PCI ($PCI_{i,t-1}$). The "with controls" column additionally includes lagged log research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Institution and year fixed effects are included throughout, and standard errors are clustered at the institution level.

Appendix B3: Lag Sensitivity

Table 8: Lag Sensitivity: All Outcomes

Outcome	Contemp.	Lag 1	Lag 2	Lag 3
Ln(New Patent Apps)	$+1.527^{**}$	$+1.476^{***}$	$+1.655^{***}$	$+1.176^{**}$
Ln(Total Patent Apps)	$+1.495^{**}$	$+1.759^{***}$	$+1.867^{***}$	$+1.572^{**}$
Ln(Patents Issued)	$+1.234^{*}$	$+1.632^{**}$	$+1.763^{***}$	$+1.983^{***}$
Ln(Disclosures)	-0.271	-0.063	$+0.049$	-0.006
Conv. Rate	$+1.798^{***}$	$+1.539^{**}$	$+1.606^{**}$	$+1.182^{**}$
Ln(Licenses)	$+0.285$	$+0.515$	$+0.890$	$+0.548$
Ln(Startups)	$+0.104$	$+0.445$	$+0.200$	-0.321
Ln(License Income)	$+2.026$	$+2.250^{*}$	$+2.376^{*}$	$+2.416^{*}$

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include institution and year fixed effects and controls for lagged log research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Conversion rate equals Ln(New Patent Applications) − Ln(Disclosures).

Appendix C: Invention Quality Robustness Tests

C1. Grant Rate Falsification

A potential confound in the main analysis is that within-institution PCI improvements coincide with shifts in the quality composition of faculty inventions: if better inventors arrive or become more active at the same time as a policy revision, the conversion rate and patent application counts could rise independently of any change in TTO coordination efficiency. We test this interpretation using the USPTO grant rate, defined as the ratio of patents issued to new patent applications filed in the same year, $\ln(1 + \text{Issued}_{it}/(\text{Applied}_{it} + 1))$. Because USPTO examiners evaluate applications on technical and legal merit independently of the TTO, a quality composition story predicts that the grant rate should rise alongside applications when PCI improves. The coordination friction story predicts no grant rate response: the same quality pool of inventions is simply being processed more efficiently, so the share reaching grant status is unchanged.

Table 9 reports TWFE estimates of Equation 1 for the four outcomes simultaneously. PCI predicts a 9.7 percent increase in new patent applications ($b = 1.476$, $p = 0.007$) and a significant increase in patents issued ($b = 1.632$, $p = 0.013$), consistent with the main results. The grant rate estimate is $b = 0.017$ ($p = 0.923$), statistically indistinguishable from zero. The null grant rate result, observed alongside positive application and issuance effects, is inconsistent with a quality composition account and confirms that the PCI association reflects more efficient processing of the same underlying pool of faculty inventions rather than an improvement in average invention quality entering the pipeline.

Table 9: Invention Quality Falsification: Grant Rate Test

Outcome	$\hat{\beta}$	SE	p	N
Ln(New Patent Applications)	$+1.476^{***}$	0.537	0.007	2,346
Ln(Patents Issued)	$+1.632^{**}$	0.648	0.013	2,322
Ln(Grant Rate: Issued/Applied)	+0.055	0.174	0.750	2,322
Conversion Rate	$+1.539^{**}$	0.609	0.012	2,346

$^{*}p < 0.10$, $^{**}p < 0.05$. All specifications include institution and year fixed effects and control for lagged log total research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Standard errors clustered at the institution level. Grant Rate is defined as $\ln(1 + \text{Patents Issued}_{it}/(\text{New Applications}_{it} + 1))$. The null grant rate coefficient alongside significant application and issuance effects is inconsistent with a quality composition story and supports the coordination friction interpretation.

C2. Royalty Share Heterogeneity

Table 10 splits the sample at the median royalty share (40 percent) to examine whether the PCI effect is concentrated among institutions with stronger inventor financial incentives. The PCI coefficient for patent applications is $b = 1.942$ ($p = 0.033$) at high-royalty institutions and $b = 0.526$ ($p = 0.457$) at low-royalty institutions. The pattern is consistent with complementarity between financial incentives and communicative governance: institutions that invest in inventor engagement through royalty sharing appear to extract greater returns from improvements in policy document clarity. This finding does not discriminate between the quality and friction stories on its own, since high royalty shares may attract both more engaged and higher-quality inventors, but it is consistent with the broader heterogeneity pattern in Table 4, where effects are largest at private and research-intensive universities where inventor engagement is already strong.

Table 10: Royalty Share Split Sample

Subsample	$\hat{\beta}$	SE	p	N
Panel A: Ln(New Patent Applications)				
High royalty share ($\geq 40\%$)	$+1.942^{**}$	0.900	0.033	1,262
Low royalty share ($< 40\%$)	$+0.526$	0.705	0.457	1,084
Panel B: Conversion Rate				
High royalty share ($\geq 40\%$)	$+1.668^{*}$	0.969	0.088	1,262
Low royalty share ($< 40\%$)	$+0.869$	0.537	0.108	1,084

$^{*}p < 0.10$, $^{**}p < 0.05$. Median royalty share = 40.0%. All specifications include institution and year fixed effects and control for lagged log total research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. Standard errors clustered at the institution level. The concentration of the PCI effect at high-royalty institutions is consistent with complementarity between financial incentives and communicative governance quality; see Lach and Schankerman (2004) for evidence on royalty share effects on disclosure and licensing activity.

Appendix D: TLO Age Trajectory Control

A plausible time-varying confounder is organizational leadership change: a new TTO director could simultaneously revise IP governance documents toward more collaborative language and shift the office's patenting strategy, independently driving both PCI and patent outcomes. We address this directly by adding lagged TLO age and its interaction with a within-institution demeaned year trend to the baseline specification. This interaction allows institutions at different stages of the TTO life-cycle to follow distinct outcome trajectories, absorbing the most credible version of the leadership confound without requiring direct observation of director turnover.

Table 11 reports baseline and augmented estimates for the three outcomes central to the identification argument. The PCI coefficient strengthens rather than attenuates after controlling for TLO life-cycle trajectories: the patent application estimate rises from $b = 1.476$ to $b = 1.506$ ($p = 0.006$) and the conversion rate from $b = 1.539$ to $b = 1.548$ ($p = 0.011$). The disclosure coefficient remains statistically null in both specifications. This pattern is the opposite of what the leadership confound predicts: if organizational change were driving both policy rewrites and patenting improvements, absorbing TLO age trajectories should attenuate PCI, not sharpen it.

Table 11: TLO Age Trajectory Control

Outcome	Specification	$\hat{\beta}$	SE	p	N
Ln(New Patent Applications)	Baseline	$+1.476^{***}$	0.537	0.007	2,346
	+ TLO age × year	$+1.506^{***}$	0.544	0.006	2,346
Conversion Rate	Baseline	$+1.539^{**}$	0.609	0.012	2,346
	+ TLO age × year	$+1.548^{**}$	0.604	0.011	2,346
Ln(Disclosures)	Baseline	-0.063	0.561	0.910	2,346
	+ TLO age × year	-0.042	0.577	0.942	2,346

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include institution and year fixed effects and control for lagged log total research expenditure, lagged log licensing FTEs, lagged inventor royalty share, and lagged TLO age. The augmented specification adds the interaction between lagged TLO age and a within-institution demeaned year trend. Standard errors clustered at the institution level. PCI strengthening after this control is inconsistent with the leadership confound story, which predicts attenuation.

Appendix E: Staggered Adoption Robustness

Recent methodological work demonstrates that TWFE estimators with staggered treatment adoption can produce negatively weighted treatment effect estimates when treatment effects are heterogeneous across adoption cohorts (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). Although the primary identifying variation in Equation 1 is continuous rather than binary – PCI is a continuous annual measure and the TWFE estimate reflects a weighted average over the full distribution of year-to-year changes, not a binary adoption contrast – we implement the Callaway and Sant'Anna (2021) group-time average treatment effect (ATT) estimator as a direct robustness check. Revision cohorts are defined by the year of first large positive PCI revision ($|\Delta\text{PCI}| > 0.05$) among the 40 treated institutions identified in Section 5, with not-yet-treated institutions serving as the comparison group. The aggregated ATT across cohorts is consistent in sign and magnitude with the TWFE baseline estimate reported in Table 3, confirming that the main finding is not an artifact of differential timing weights across adoption cohorts. Full estimation output is available from the authors upon request.