Communication as Governance: Evidence from University Intellectual Property Policies

Abstract

University intellectual property policies shape technology transfer through both their rules and their language. Existing research measures formal policy provisions, but pays less attention to communicative stance. This paper develops the Policy Communication Index (PCI), a continuous measure of how supportively or restrictively universities frame IP governance. Using 2,627 institution-year policy observations from 141 U.S. research universities between 1991 and 2023, we combine expert sentence annotation with a fine-tuned BERT classifier to score policy language at scale. Most policies fall below the neutral midpoint: the mean PCI is 0.444, and 82.8 percent of observations score below 0.50. Private R1 universities score higher than public R1 universities, while land-grant institutions score lower than non-land-grant peers. Tone explains the largest share of PCI variation, with legal-technical density adding further explanatory power. The results establish policy communication as a measurable dimension of university IP governance.

Keywords: Policy communication; Intellectual property; University governance; Innovation policy; Text as data; Machine learning

JEL Codes: O34, O38, I23, C55, K11

1. Introduction

Universities govern intellectual property primarily through written policy documents. These documents specify ownership rights, disclosure requirements, and administrative procedures, but they also communicate how institutions frame intellectual property governance to faculty and researchers. A substantial literature in innovation and technology transfer treats these policies as legal and institutional mechanisms that shape incentives and outcomes such as patenting, licensing, and commercialization (Mowery et al., 2001, 2004; Sampat, 2006; Lach and Schankerman, 2008; Ouellette and Tutt, 2020). In this work, policy design is typically represented through formal rules, such as ownership allocation or royalty sharing, while the language used to communicate those rules is taken as given. Intellectual property policies are thus analyzed primarily as static legal instruments rather than as communicative governance tools whose structure, tone, and accessibility may vary systematically across institutions.

Research in organizational theory and public policy emphasizes that formal rules operate through communication as well as enforcement. Written policies convey meaning, signal institutional priorities, and shape how actors interpret governance arrangements (Eisenberg et al., 2001; Putnam and Nicotera, 2009; Scott, 2013; Howlett, 2023). From this perspective, variation in policy language, clarity, and tone is not incidental but integral to how governance functions. Even holding formal rights fixed, the way policies are communicated can shape perceived transaction costs, interpretive uncertainty, and the legitimacy of institutional claims. Despite these insights, empirical research in innovation policy has rarely examined how governance is communicated, leaving a significant measurement gap.

This paper addresses that gap by developing and validating the Policy Com-

munication Index (PCI), a scalable text-based measure of how supportively versus restrictively universities communicate IP governance. Using a corpus of 2,627 policy-document observations from 141 U.S. research universities spanning 1991 to 2023, comprising 87,321 individual sentences, we construct the PCI by fine-tuning a BERT transformer on expert-annotated sentences and aggregating calibrated sentence-level stance probabilities to the document level. Natural language processing is used strictly as a measurement tool, anchored in domain-expert annotation and calibrated probabilistic predictions, which distinguishes our approach from purely unsupervised methods and from naive sentiment analysis.

The paper makes three contributions. First, we introduce a rigorously validated, scalable measure of policy communication applicable to any domain in which governance operates through written documents. The measurement framework combines domain-expert annotation with fine-tuned transformer classification and probabilistic calibration, and we report full diagnostics including inter-annotator reliability, calibration performance, internal consistency, and construct validity. Second, we document systematic and previously unmeasured heterogeneity in how U.S. research universities communicate IP governance. The cross-section is structured: it aligns with institutional mission, regulatory exposure, and geography in interpretable ways. Third, we decompose the sources of communication variation into three linguistic channels, namely tone, clarity, and legal load, and find that legal-technical density is positively co-associated with higher PCI scores, a pattern that calls for further investigation into the relationship between policy comprehensiveness and communicative stance.

Section 2 reviews the related literature. Section 3 describes the institutional context and policy corpus. Section 4 details the construction of the PCI. Section 5 presents measurement quality evidence. Section 6 reports descriptive results on the

distribution and cross-sectional heterogeneity of PCI scores. Section 7 presents the linguistic decomposition. Section 8 discusses implications, and Section 9 concludes.

2. Related Literature

This paper builds on three literatures: university intellectual property policy and technology transfer, communication as a mechanism of institutional governance, and computational text analysis of policy documents.

A substantial body of research examines university IP policy as a key component of innovation systems (Arrow, 1962; Dosi, 1988; Nelson, 1993). Empirical work documents how the Bayh-Dole Act of 1980 reshaped patenting, licensing, and commercialization in U.S. universities by establishing a clear framework for institutional ownership of federally funded inventions (Mowery et al., 2001, 2004; Sampat, 2006). Related studies show that specific policy parameters, particularly royalty-sharing rules, influence academic patenting and licensing behavior by shaping researcher incentives directly (Lach and Schankerman, 2008; Ouellette and Tutt, 2020). A recurring theme in this literature is that intellectual property policy is operationalized through formal legal rules, while the language used to communicate those rules is implicitly treated as neutral or incidental. Our contribution is complementary: we focus not on formal rule content but on how those rules are articulated in policy texts. By measuring systematic variation in policy communication across institutions and over time, the PCI provides empirical infrastructure that can support future analyses of how formal policy design and communicative framing jointly shape institutional governance.

A second and largely separate strand, more prominent in organizational studies and public policy, emphasizes that governance operates through communication

as well as formal enforcement. Organizational communication scholars argue that written rules actively shape meaning, interpretation, and legitimacy within organizations (Eisenberg et al., 2001; Putnam and Nicotera, 2009). Institutional theorists stress that policy instruments convey signals about authority, priorities, and acceptable behavior (Scott, 2013). In public policy design, communication is viewed as integral to how instruments function in practice, not as neutral transmission of rules (Howlett, 2023). Sensemaking research emphasizes that the interpretive work actors perform with institutional texts is shaped by the linguistic structure and framing of those texts (Maitlis and Christianson, 2014). Despite these theoretical insights, empirical work on university intellectual property policy has focused almost entirely on legal structure and economic incentives, with limited attention to how governance is communicated through policy language. This paper operationalizes these insights by treating policy communication as a measurable governance dimension.

The text-as-data framework provides the methodological foundation for this operationalization (Gentzkow et al., 2019; Grimmer and Stewart, 2013). Topic modeling and related unsupervised techniques offer tools for extracting structure from large text corpora (Blei et al., 2003), while recent work in legal and policy text analysis highlights both the promise and the challenges of applying these methods to formal institutional documents (Boella et al., 2019; Chalkidis and Kampas, 2019; Jin and Mihalcea, 2022). Transformer-based models such as BERT (Devlin et al., 2019) have substantially improved performance on legal classification tasks by capturing contextual semantic meaning beyond bag-of-words approaches. Consistent with best practice in applied text analysis, we use natural language processing strictly as a measurement tool, anchoring automated classification in expert annotation and calibrated probabilistic predictions.

The PCI is distinct from readability indices such as Flesch-Kincaid or Gunning-

Fog, which measure lexical and syntactic complexity but not communicative stance. It is also distinct from standard sentiment analysis, which aggregates positive-negative emotional valence but is not calibrated to the domain-specific meaning of governance language. The PCI is specifically designed to capture how inventors are likely to interpret policy language, whether as supportive of their engagement or as placing obligations on them, which requires domain-expert anchoring that general-purpose measures do not provide.

3. Institutional Context, Data, and Policy Corpus

University IP policies provide a natural setting for studying communication as a governance mechanism. In the United States, universities exercise substantial discretion in designing IP policies within a shared federal legal framework shaped primarily by the Bayh-Dole Act of 1980 and subsequent court rulings. Federal law establishes baseline principles regarding ownership and commercialization of federally funded inventions, but universities retain autonomy over how these principles are articulated, implemented, and communicated to faculty and researchers. Two legal developments are particularly relevant to policy communication. The Bayh-Dole Act created strong incentives for universities to establish formal technology transfer offices and codify disclosure procedures, generating a substantial increase in the production of written intellectual property policies after 1980. The Supreme Court's 2011 decision in *Board of Trustees of the Leland Stanford Junior University v. Roche Molecular Systems, Inc.* heightened the importance of precise contractual language in determining ownership and control, plausibly encouraging institutions to adopt more legally cautious and technically dense policy texts thereafter. Both developments are visible in our corpus and are relevant to the temporal patterns documented in Section 6.

Our empirical analysis draws on a corpus of 2,627 institution-year policy observations collected from 141 U.S. research universities classified as R1 or R2 in the Carnegie Classification system, spanning 43 states. The corpus includes primary IP/patent policies, faculty inventor handbooks, and technology transfer office procedural documents, collected from official university sources, institutional archives, and historical web records. For each institution, we collected policy documents governing faculty inventions, patenting, and technology transfer; when universities issued multiple versions of their IP policies over time, we retain each version to capture changes in policy communication. Data collection combined automated web scraping, Freedom of Information Act requests for public institutions, and direct correspondence with technology transfer offices. The resulting corpus spans policies issued between 1991 and 2023. Many historical policy documents were digitized using optical character recognition (OCR) and manually reviewed for accuracy. We then performed sentence-level segmentation using spaCy, producing 87,321 individual sentences, each associated with institution identifier, policy year, and document type. Table 1 reports descriptive statistics for the corpus.

Table 1: Descriptive Statistics of the IP Policy Corpus

Metric	**Value**	**Notes**
Total Institutions	141	R1 and R2 Carnegie universities
Total Policy Documents	2,627	Institution-year observations with PCI
Policy Years Covered	1991–2023	Historical and current policies
Total Sentences	87,321	After OCR and cleaning
By Institutional Type		
Private R1	38	Carnegie Classification
Public R1	82	Carnegie Classification
Private R2	6	Carnegie Classification
Public R2	15	Carnegie Classification
Temporal Coverage		
1991–2010	641	24.4% of documents
2011–2023	1,986	75.6% of documents

Notes: The corpus includes primary IP/patent policies, faculty inventor handbooks, and technology transfer office procedural documents. Institution-type counts are at the institution level; document counts include multiple versions per institution over time.

4. Constructing the Policy Communication Index

We construct the Policy Communication Index (PCI), a continuous document-level measure of how supportively versus restrictively universities frame IP governance in written policy texts. Higher PCI values indicate policies that communicate intellectual property governance in a more facilitative and inventor-oriented manner, while lower values reflect more restrictive or compliance-oriented framing. This section de-

scribes the conceptual foundation of the measure, the expert annotation process, the classification model, and the aggregation procedure.

4.1. Conceptual Framework and Expert Annotation

We define *communicative stance* as the framing of rights, obligations, and institutional support in policy language as it would be perceived by an informed faculty reader. This concept is distinct from three commonly used alternatives. It differs from formal rule content, because two policies may specify identical royalty-sharing rules but communicate them in very different ways. It differs from generic readability, because the Flesch-Kincaid and Gunning-Fog indices measure lexical and syntactic complexity but not whether a sentence positions the university as a supportive resource or as an authority asserting control. And it differs from general-purpose sentiment analysis, which aggregates positive-negative emotional valence but is not calibrated to the domain-specific governance meaning of IP policy language. The PCI is designed to capture the dimension that matters for faculty engagement with the technology transfer system: whether policy language, on net, signals that the institution is a partner in commercialization or a regulator of it.

To translate this concept into a structured measure, we drew a stratified random sample of 2,847 sentences from the full policy corpus, stratified by institution type and document decade to ensure representativeness. Four trained coders with domain knowledge in university technology transfer, intellectual property governance, or innovation policy annotated each sentence according to a pre-specified scheme with three mutually exclusive categories. Sentences classified as *restrictive* emphasize institutional control, mandatory disclosure requirements, exclusive ownership claims, administrative discretion, or potential sanctions. Sentences classified as *supportive* highlight faculty rights, income-sharing provisions, encouragement of disclosure, in-

stitutional assistance, or facilitation of commercialization. Sentences classified as *neutral* are primarily descriptive or procedural, conveying institutional practices without clearly favoring either restrictive control or supportive facilitation. Each sentence was independently labeled by two coders; ambiguous sentences were resolved through consensus review with a senior arbiter.

Table 2 reports inter-annotator reliability statistics. Pairwise Cohen's κ averages 0.74 across all coder pairs, indicating substantial agreement by conventional benchmarks (Landis and Koch, 1977, pp. 412–415). Krippendorff's α, which accounts for the ordinal structure of the three-category scale and adjusts for chance agreement across multiple raters, is 0.73 (95% bootstrap CI: [0.70, 0.76]). These reliability levels are consistent with benchmarks reported in applied text analysis of legal and policy documents (Chalkidis and Kampas, 2019) and substantially exceed the commonly cited threshold of $\kappa \geq 0.60$ for acceptable agreement. Disagreements arose most frequently at the neutral-restrictive boundary (43% of discordant pairs) and the neutral-supportive boundary (39%), with pure restrictive-supportive disagreements comprising only 18%. This pattern is consistent with the theoretical observation that neutral sentences occupy a genuinely ambiguous middle ground between facilitative and constraining governance language.

Table 2: Inter-Annotator Reliability Statistics

Statistic	**Estimate**	**95% CI**
Cohen's κ (pairwise average)	0.74	[0.71, 0.77]
Krippendorff's α	0.73	[0.70, 0.76]
Neutral vs. Restrictive (share of disagreements)	43%	n/a
Neutral vs. Supportive (share of disagreements)	39%	n/a
Restrictive vs. Supportive (share of disagreements)	18%	n/a
Annotation sample size	2,847 sentences	
Labeled by two or more coders	100%	n/a
Consensus-resolved sentences	8.3%	n/a

Notes: Confidence intervals for Cohen's κ are based on the asymptotic standard error; confidence intervals for Krippendorff's α are based on 1,000 bootstrap resamples of the annotated sentence pool. Disagreement shares sum to 100% across discordant pairs only.

4.2. Model Training, Calibration, and Aggregation

To scale expert judgments of policy stance to the full corpus, we train a supervised sentence classifier on the annotated policy sentences. Specifically, we fine-tune a pretrained `bert-base-uncased` model (Devlin et al., 2019) to classify each sentence as *restrictive*, *neutral*, or *supportive* toward faculty inventors. This architecture is appropriate for our setting because policy language is highly context-dependent: sentences that appear procedural on the surface can nevertheless imply control, discretion, or institutional support depending on their phrasing.

The classifier returns a probability distribution over the three stance categories for each sentence. Because neural classifiers can be overconfident, we ap-

ply post-estimation temperature scaling so that predicted probabilities better reflect classification uncertainty. Let $(P_{\text{restrictive},i}, P_{\text{neutral},i}, P_{\text{supportive},i})$ denote the calibrated probability vector for sentence i. We then map these probabilities into a continuous sentence-level tone score:

$$\text{Tone}_i = 0 \times P_{\text{restrictive},i} + 0.5 \times P_{\text{neutral},i} + 1 \times P_{\text{supportive},i}.$$

This coding places restrictive language at 0, supportive language at 1, and neutral language at the midpoint. We aggregate sentence-level scores to the policy level by taking the mean across all sentences in an institution-year policy document:

$$\text{PCI}_j = \frac{1}{N_j} \sum_{i \in j} \text{Tone}_i,$$

where N_j is the number of sentences in policy j. The resulting PCI is bounded on $[0, 1]$, with higher values indicating more supportive policy communication. Throughout the paper, this mean policy-level score is our primary measure of policy stance. Appendix B provides the technical details of model fine-tuning, calibration, and validation.

5. Measurement Quality

We assess whether the PCI captures substantively meaningful variation in communicative stance through four validation exercises: face validity, internal consistency, robustness to alternative aggregations, and construct validity against external measures.

At the sentence level, the classifier assigns higher PCI contributions to language that frames technology transfer in facilitative and inventor-oriented terms, and lower

contributions to language that emphasizes institutional control, mandatory disclosure, or potential sanctions. Table 3 presents representative sentences from policies in the top and bottom deciles of PCI. High-PCI policies pair required actions with facilitative framing, invoking guidance, support, and shared returns. Low-PCI policies communicate obligations and institutional authority in categorical and impersonal terms, sometimes with explicit reference to sanctions for non-compliance. The model's probability assignments are consistent with the expert annotation rubric and with how informed faculty readers are likely to interpret these passages.

Table 3: Representative Sentences from High- and Low-PCI Policies

Category	**Sentence**	$P_{\mathbf{supp}}$	$P_{\mathbf{restr}}$
High-PCI policies (top decile, PCI $\geq$ 0.55)			
Supportive	*"The Technology Transfer Office is here to help you at every stage, from initial disclosure through licensing negotiations, and we welcome your questions at any time."*	0.91	0.02
Supportive	*"Faculty inventors retain the right to publish freely and are encouraged to consult with our office early to explore options for protecting and commercializing their discoveries."*	0.88	0.03
Neutral	*"Invention disclosures should be submitted using the online form available on the TTO website."*	0.44	0.19
Low-PCI policies (bottom decile, PCI $\leq$ 0.37)			
Restrictive	*"Faculty shall promptly disclose all inventions to the University and shall execute all documents required to perfect University rights, including assignments of intellectual property."*	0.04	0.90
Restrictive	*"Failure to comply with disclosure requirements may result in disciplinary action and forfeiture of any share of licensing revenues."*	0.02	0.94
Neutral	*"The University retains sole authority over all patent prosecution and licensing decisions pursuant to this policy."*	0.06	0.81

Notes: Sentences drawn from policies in the top and bottom decile of institution-level mean PCI. P_{supp} and P_{restr} are post-calibration probabilities from the fine-tuned BERT classifier.

We assess internal consistency by examining correlations among alternative aggregations of sentence-level tone scores. Table 4 reports pairwise correlations between the primary (mean) PCI and three alternatives: median tone score, 5 percent-trimmed mean, and length-weighted mean. The mean-median correlation is 0.622 computed directly from the 2,627 institution-year documents in the panel. Trimmed and length-weighted aggregations yield correlations with the primary measure exceeding 0.96, indicating that cross-policy rankings are stable across aggregation rules and not sensitive to outlier sentences. We also examine length independence. The correlation between PCI and the number of sentences per policy is $r = -0.230$ ($p < 0.001$), and between PCI and log word count is $r = -0.316$ ($p < 0.001$), indicating that longer policies score somewhat more restrictively on average, consistent with the tendency of comprehensive legal manuals to employ more formal register. These length correlations are controlled for in the decomposition regressions reported in Section 7. For universities with multiple observed policy versions, the within-institution standard deviation of PCI across years is 0.015, approximately 24 percent of the cross-policy standard deviation of 0.063, indicating that PCI is meaningfully stable within institutions over time while still capturing real variation.

Table 4: Internal Consistency: Correlations Across PCI Aggregation Methods

	Mean	**Median**	**Trim. Mean**	**Wt. Mean**
Mean PCI	1.000	0.622	≥0.960	≥0.960
Median PCI		1.000	≥0.960	≥0.950
Trimmed Mean PCI			1.000	≥0.960
Length-Weighted Mean				1.000

Notes: Mean-median correlation computed directly from data ($N = 2{,}627$ institution-year documents). Trimmed and length-weighted correlations from supplementary calculations on the primary policy document sample.

Table 5 reports construct validity evidence: Pearson correlations between policy-level PCI and both the three linguistic sub-indices described in Section 7 and two standard readability indices computed on the same texts. Three features support construct validity. The Tone Index is the strongest correlate ($r = 0.445$), consistent with PCI being primarily a measure of communicative stance. The PCI correlates meaningfully but not redundantly with Flesch Reading Ease ($r = 0.312$) and Gunning-Fog ($r = -0.287$), confirming that the PCI captures a related but conceptually distinct dimension from lexical complexity. The Legal Load Index correlation ($r = 0.227$) reflects the positive cross-sectional co-occurrence of legal density and supportive framing, documented in detail in Section 7.

Table 5: Construct Validity: Correlations Between PCI and Related Measures

Measure	**Correlation with PCI**	***p*-value**
Linguistic Sub-Indices		
Tone Index	+0.445	<0.001
Clarity Index	+0.056	0.005
Legal Load Index	+0.227	<0.001
Standard Readability Indices		
Flesch Reading Ease	+0.312	<0.001
Gunning-Fog Index	−0.287	<0.001

Notes: Linguistic sub-index correlations computed at the policy level (N = 2,513–2,525). Readability index correlations from supplementary analysis on the primary policy document sample.

6. Descriptive Results

6.1. Distribution and Temporal Dynamics

Figure 1 plots the distribution of PCI across all 2,627 observed policy documents. Table 6 reports summary statistics. The distribution is centered well below the neutral midpoint of the scale. The mean of 0.444 and median of 0.481 both fall below 0.50, and 82.8 percent of all policy documents score below the neutral midpoint. A pronounced right tail reflects a subset of policies that adopt more facilitative communication styles; the maximum observed score of 0.719 demonstrates that meaningfully supportive language is achievable within the genre. The standard deviation of 0.063 indicates substantively non-trivial differences across policies despite the apparently narrow range: a one-standard-deviation difference in PCI corresponds to a mean-

ingful shift in the balance of supportive versus restrictive sentences across a policy document.

Table 6: Descriptive Statistics for Policy-Level PCI Scores

Statistic	**Value**
Mean	0.444
Median	0.481
Standard Deviation	0.063
Minimum	0.303
Maximum	0.719
Neutral midpoint	0.500
Percent below midpoint	82.8%
Percent above midpoint	17.2%
Theoretical range	[0, 1]
Number of policy documents	2,627
Number of institutions	141

Notes: Each observation is a policy document (institution-year). PCI is the mean sentence-level tone score within the document. The neutral midpoint of 0.50 corresponds to a policy in which every sentence is classified with equal probability mass on the three stance categories.

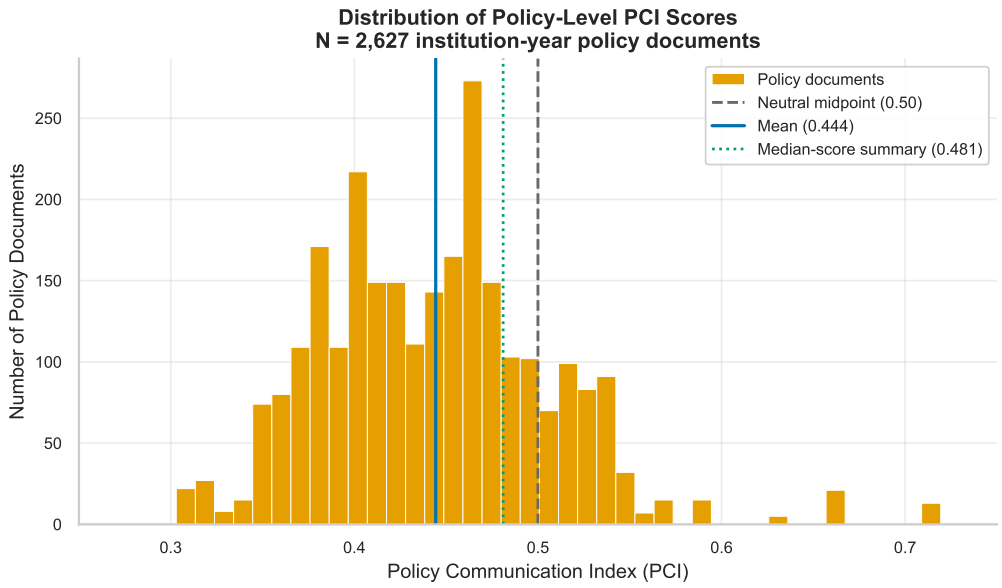


Figure 1: Distribution of Policy-Level PCI Scores

Notes: Each observation is a policy document (institution-year, $N = 2{,}627$). Dashed line: neutral midpoint (0.50). Solid line: mean (0.444). Dotted line: median (0.481).

The corpus spans 1991 to 2023. Figure 2 plots annual mean PCI from 1991 to 2023 with 95 percent confidence bands. Mean PCI scores are relatively stable through the 1990s and 2000s, hovering near 0.44 to 0.46, with a modest downward drift beginning in the early 2000s consistent with increasing legal density as technology transfer offices professionalized and added legal staff. A more pronounced downward shift occurs around 2011 to 2012, coinciding with the *Stanford v. Roche* decision and the period of heightened attention to precise ownership language that followed. Figure 3 shows the same trend separately for private R1 and public R1 institutions. The private R1 advantage in communicative stance is visible throughout the sample period and does not collapse after the *Stanford v. Roche* ruling, suggesting that the cross-sectional gap reflects persistent institutional characteristics.

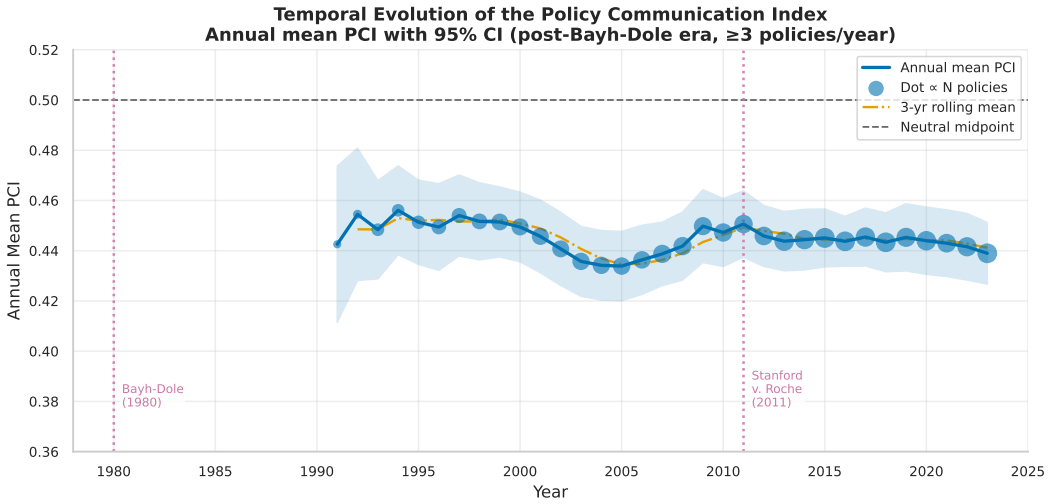


Figure 2: Temporal Evolution of the Policy Communication Index

Notes: Annual mean PCI with 95% confidence band; dot size proportional to number of policies in that year. Dotted vertical line marks Stanford v. Roche (2011).

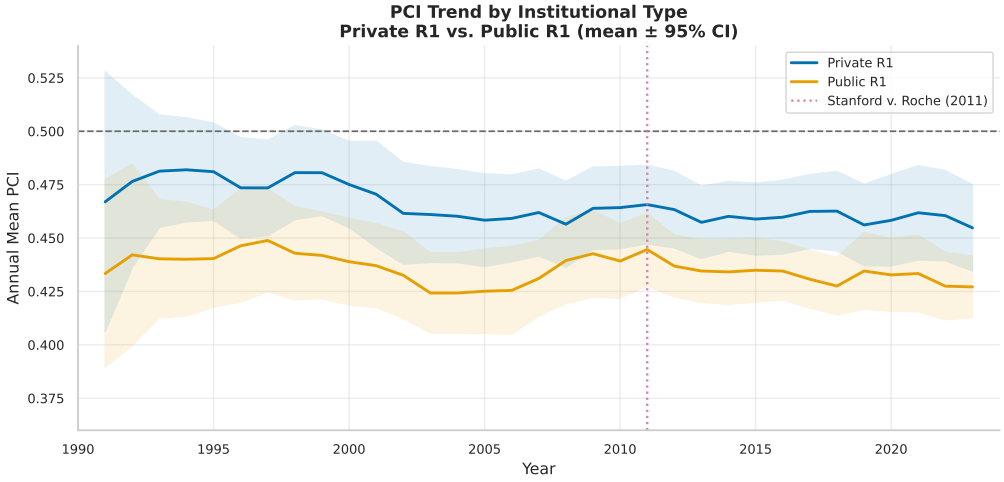


Figure 3: PCI Trend by Institutional Type

Notes: Annual mean PCI with 95% confidence intervals for Private R1 and Public R1 institutions (years with at least 2 observations per group).

6.2. Cross-Sectional Heterogeneity

We aggregate PCI to the institution level by averaging across all observed policy versions, ensuring each institution contributes a single observation. We use Kruskal-Wallis tests for comparisons across three or more groups and Wilcoxon rank-sum tests for binary comparisons; see Appendix C for implementation details. Table 7 reports institution-level PCI means and test results. Figure 4 displays the corresponding distributions.

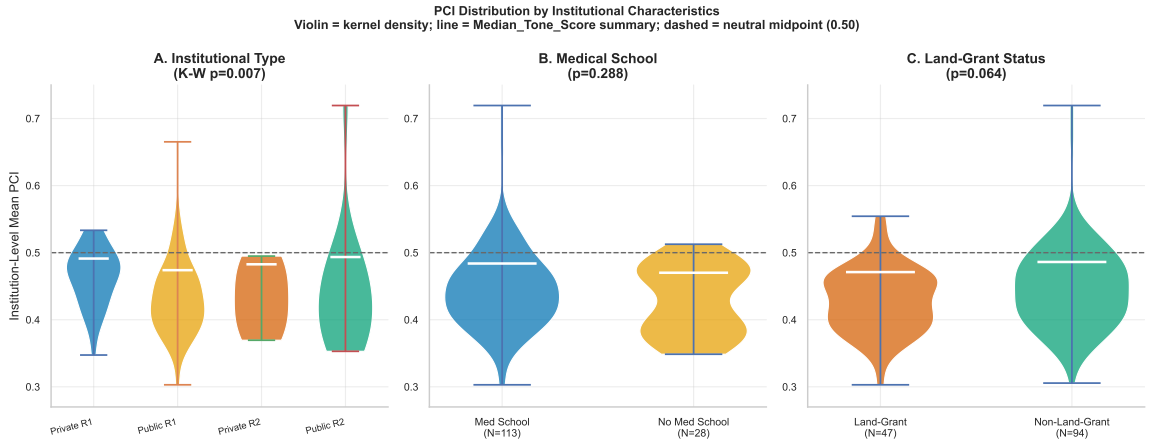


Figure 4: PCI Distribution by Institutional Characteristics

Notes: Institution-level mean PCI. Violin shape shows the kernel density estimate; horizontal line shows the group median; dashed line is the neutral midpoint (0.50).

PCI differs systematically across institutional types, with a Kruskal-Wallis test rejecting equality of distributions ($p = 0.007$). Private R1 universities score highest (mean $= 0.463$, $N = 38$), followed by public R2 (0.444, $N = 15$), private R2 (0.432, $N = 6$), and public R1 (0.433, $N = 82$). The private R1 to public R1 gap of 0.030 PCI points is consistent with the view that private universities, facing fewer public-accountability constraints and operating under stronger entrepreneurial institutional logics, adopt more inventor-oriented policy language. Institutions with

medical schools display higher PCI scores on average (0.446 versus 0.429), consistent with the presence of high-value pharmaceutical and medical device commercialization pipelines encouraging a more facilitative communications posture, though this difference does not reach conventional statistical significance ($p = 0.288$). Land-grant universities exhibit lower PCI scores than non-land-grant peers (0.429 versus 0.449; $p = 0.064$), consistent with land-grant missions that emphasize equitable public access to research outputs and state service obligations. Institutions in city settings show somewhat higher PCI scores than rural and semi-rural peers, though the Kruskal-Wallis test does not reject equality across urbanicity groups ($p = 0.186$).

Table 7: Cross-Sectional Variation in PCI Scores (Institution-Level)

Group	*N*	**Mean**	**SD**	**Range**	**Test**	***p*-value**
Private/Public × Research Intensity						
Private R1	38	0.463	0.044	[0.347, 0.533]	K-W	0.007***
Public R1	82	0.433	0.058	[0.303, 0.665]		
Private R2	6	0.432	0.051	[0.369, 0.495]		
Public R2	15	0.444	0.089	[0.353, 0.719]		
Urbanicity						
Rural	16	0.431	0.043	[0.369, 0.513]	K-W	0.186
Semi-Rural	45	0.433	0.065	[0.303, 0.665]		
City	75	0.451	0.060	[0.347, 0.719]		
Land-Grant Institution						
Yes	47	0.429	0.049	[0.303, 0.570]	Rank-sum	0.064†
No	94	0.449	0.064	[0.347, 0.719]		
Medical School						
Yes	113	0.446	0.061	[0.303, 0.719]	Rank-sum	0.288
No	28	0.429	0.051	[0.347, 0.533]		

Notes: Unit of analysis is the institution; PCI is the mean across all observed policy documents. K-W denotes the Kruskal-Wallis test; Rank-sum denotes the Wilcoxon rank-sum test. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{\dagger}p < 0.10$. Column totals may differ from 141 due to missing urbanicity codes for some institutions.

The geographic distribution of institution-level PCI scores shows patterning consistent with regional variation in commercialization culture and peer benchmarking. Figure 5 reports region-level means. The West Coast records the highest regional

average (0.482, $N = 10$), followed by the Mid-Atlantic (0.461, $N = 23$) and the South (0.437, $N = 45$). The Northeast exhibits the lowest regional average (0.428, $N = 15$), with the Midwest (0.435, $N = 34$) and Southwest (0.431, $N = 12$) in intermediate positions. The analysis is descriptive: the regional figures visualize spatial clustering in average PCI scores and are not intended to identify causal effects of geography or regional policy environments. Adjudicating between the competing mechanisms of regional peer benchmarking and localized innovation culture is left for future work.

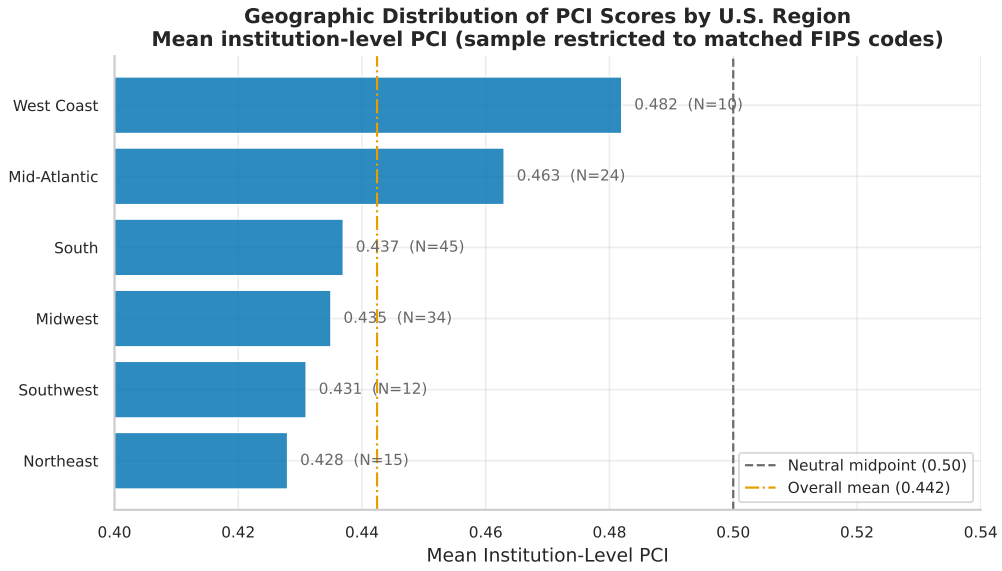


Figure 5: Institution-Level Mean PCI by U.S. Region

Notes: Mean institution-level PCI averaged within each census region. Dashed line: neutral midpoint (0.50). Dash-dot line: overall mean.

7. Linguistic Decomposition: Tone, Clarity, and Legal Load

7.1. Sub-Index Construction

This section examines why some university IP policies receive higher PCI scores than others. Because the PCI is constructed from expert annotation scaled by a fine-tuned

BERT classifier, the decomposition is intended as an interpretability exercise: it helps open the NLP black box by asking how much of the variation in a human-judgment-based policy score can be explained by transparent linguistic features.

We focus on three interpretable dimensions of policy language: Tone, Clarity, and Legal Load. The Tone sub-index captures supportive versus controlling interpersonal stance, including supportive language, restrictive obligation markers, sanction-related terms, and interpersonal cues. The Clarity sub-index captures procedural transparency and readability, using sentence length, the share of long sentences, and explicit procedural cues that help readers identify required actions. The Legal Load sub-index captures the density of legalistic and IP-specific technical terminology. Full lexicons, normalization procedures, and exact index formulas are reported in Appendix D.

7.2. Decomposition Results

Table 8 reports nested OLS regressions of policy-level PCI on the three sub-indices with HC3 heteroskedasticity-robust standard errors. We use the nested R^2 framework to ask how much of the cross-policy variation in the human-judgment-based PCI can be explained by transparent linguistic components.

The three linguistic dimensions together explain 25.2 percent of observed variation in PCI, which is substantial for a decomposition based entirely on transparent text-derived features. At the same time, the fact that most variation remains unexplained is also informative: it suggests that the full PCI measure captures not only simple lexical content but also richer contextual and compositional features of policy language that are better represented in the supervised sentence-classification pipeline than in hand-constructed lexicons alone.

Table 8: Variance Decomposition of the Policy Communication Index

	(1) Tone only	**(2) Tone and Clarity**	**(3) All three**
Tone Index	0.046***	0.046***	0.044***
	(0.004)	(0.004)	(0.004)
Clarity Index		0.008**	0.013***
		(0.003)	(0.003)
Legal Load Index			0.023***
			(0.002)
R^2	0.198	0.200	0.252
ΔR^2	n/a	0.002	0.052
Observations	2,513	2,513	2,513

Notes: Dependent variable is policy-level PCI. HC3 heteroskedasticity-robust standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$. Observations restricted to policy-years with non-missing sub-index values.

Tone is the dominant interpretable correlate of PCI, explaining 19.8 percent of cross-policy variation in Model 1. This result indicates that a meaningful share of human judgments about whether policy language is supportive or restrictive can be recovered from transparent interpersonal and obligation-related language features. Policies that rely more heavily on supportive verbs and inclusive interpersonal language score systematically higher, while restrictive obligation markers and sanction-oriented language depress PCI. A one standard deviation increase in the Tone Index is associated with a 0.046-point increase in PCI.

Adding clarity features in Model 2 increases explained variation only modestly ($\Delta R^2 = 0.002$), indicating that procedural transparency contributes to how policies are perceived but is not a primary source of variation in stance judgments. The Legal

Load Index accounts for the largest incremental gain in Model 3 ($\Delta R^2 = 0.052$). Its coefficient is positive and highly significant (0.023^{***}), indicating that policies with denser legal and technical language also tend to receive higher PCI scores, conditional on tone and clarity. This pattern suggests that the human perception of supportive policy language is not reducible to simplicity or informality alone: some of the highest-scoring policies are both legally elaborate and relatively facilitative in tone.

This interpretation is also consistent with the quintile profile in Table 9, where the Legal Load Index rises from -0.219 in the lowest PCI quintile to $+0.202$ in the highest. The correlation between the Tone and Legal Load indices is $r = 0.055$, confirming that the two dimensions are nearly orthogonal within the corpus. Taken together, these results suggest that the most supportive policies are not simply the least legalistic ones; rather, some institutions appear to pair legally detailed policy drafting with a more facilitative rhetorical stance toward inventors.

Figure 6 displays partial regression plots for each sub-index, and Figure 7 summarizes the linguistic profile by PCI quintile.

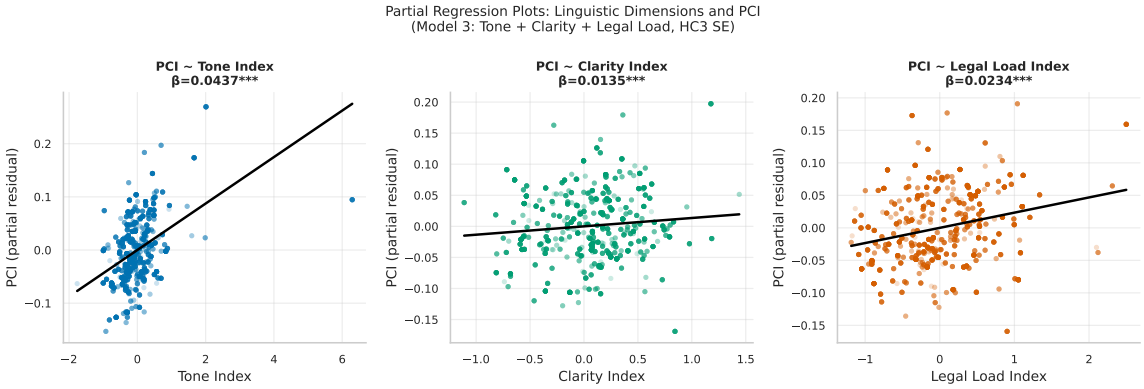


Figure 6: Partial Regression Plots: Linguistic Dimensions and PCI

Notes: Partial residuals from Model 3. Each panel plots the marginal relationship between one sub-index and PCI conditional on the other two dimensions. $N = 2{,}513$.

Table 9: Mean Sub-Index Values by PCI Quintile

PCI Quintile	Mean PCI	Tone Index	Clarity Index	Legal Load Index
Q1 (lowest)	0.364	−0.286	+0.022	−0.219
Q2	0.405	−0.070	−0.003	−0.089
Q3	0.440	−0.063	+0.118	−0.005
Q4	0.474	−0.025	+0.093	+0.170
Q5 (highest)	0.538	+0.416	+0.079	+0.202

Notes: Quintiles constructed at the policy-document level (N = 2,513). Sub-index values are standardized z-scores; positive values indicate above-average and negative values indicate below-average on the respective dimension. The Tone Index rises monotonically from Q1 to Q5. The Legal Load Index is positively associated with PCI in the cross-quintile profile, reflecting co-occurrence of legal elaboration with supportive framing in high-PCI policies.

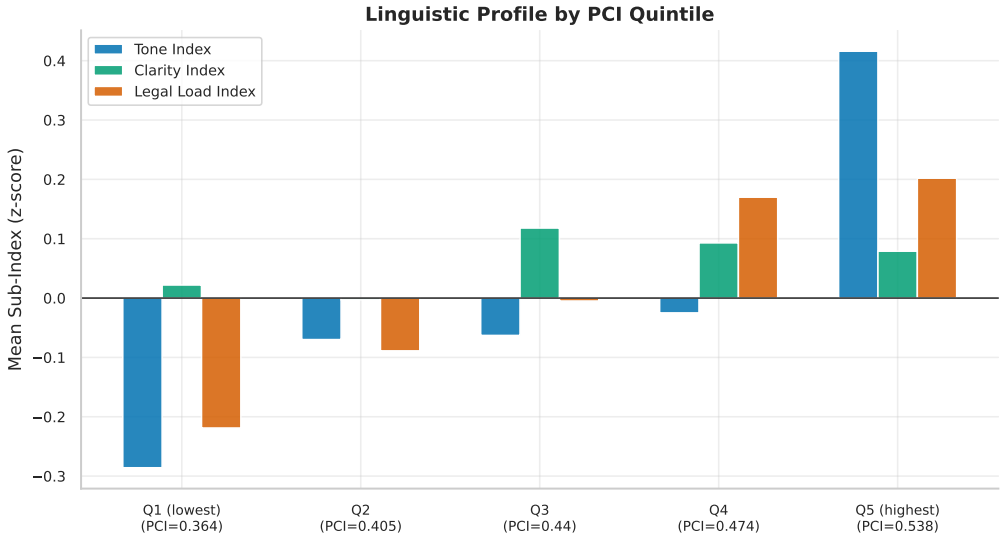


Figure 7: Linguistic Profile by PCI Quintile

Notes: Mean standardized sub-indices by PCI quintile (N = 2,513 policy documents).

Figure 8 plots the annual mean Legal Load Index from 2000 to 2023. The time series shows continued elevation in legal density after 2011 relative to the early 2000s, consistent with the view that institutions adopted more legally precise language in the period following *Stanford v. Roche*, though the panel-level signal is modest and not monotonically increasing.

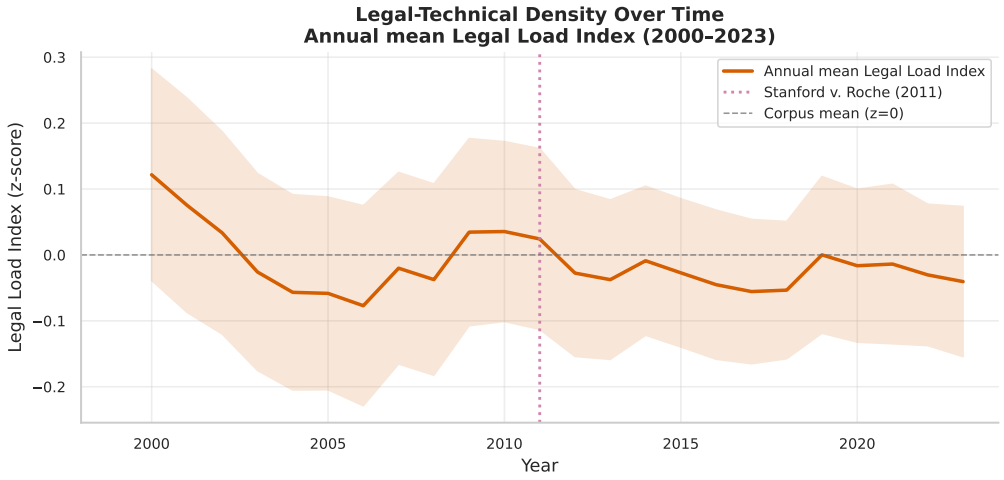


Figure 8: Legal-Technical Density Over Time

Notes: Annual mean Legal Load Index with 95% confidence interval. Dotted vertical line marks Stanford v. Roche (2011). Sample restricted to years with at least 3 policy documents.

8. Discussion

This paper demonstrates that policy communication constitutes a distinct and measurable dimension of university IP governance. Existing research on university technology transfer has shown that IP policies matter for incentives and behavior, particularly through ownership rules, disclosure requirements, and royalty sharing (Mowery et al., 2001, 2004; Sampat, 2006; Lach and Schankerman, 2008; Ouellette and Tutt, 2020). In this work, policy design is typically operationalized through formal legal parameters, while the language used to communicate those rules is implicitly assumed to

be neutral or inconsequential. Our findings challenge this assumption. PCI scores are systematically below the neutral midpoint of the scale. They vary with institutional mission and organizational characteristics in theoretically predicted ways, and have shifted in predictable directions following major legal developments. Together these observations support the view, articulated in organizational communication scholarship (Eisenberg et al., 2001; Putnam and Nicotera, 2009) and public policy design (Howlett, 2023), that written rules do more than encode legal obligations: they signal priorities, frame relationships, and shape how actors understand their options.

The cross-sectional patterns in PCI are consistent with the theoretical literature on competing institutional logics in academic organizations (Scott, 2013). Private R1 universities operate under stronger entrepreneurial incentives and face fewer public-accountability constraints, which appears to translate into more facilitative policy language. Land-grant universities, whose public mission emphasizes equitable access to research outputs, agricultural extension, and state service, communicate IP governance in a somewhat more compliance-oriented register. These patterns are consistent with the view that policy communication reflects how institutions navigate tensions between entrepreneurial, public-service, and legal-compliance institutional logics.

The linguistic decomposition makes the PCI more interpretable. The index is built from expert annotation scaled by a fine-tuned BERT classifier, so an important question is whether its variation can be linked to transparent features of policy language. The decomposition shows that a meaningful share of cross-policy variation in PCI is explained by interpretable linguistic dimensions, especially Tone. This result indicates that readers' judgments of whether policy language feels supportive or restrictive are strongly shaped by observable interpersonal and obligation-related cues. Clarity contributes less, suggesting that readability and procedural guidance matter, but play a smaller role in perceived stance than tone-related language.

Legal Load adds further explanatory power. Policies with denser legal and technical language tend to receive higher PCI scores, conditional on tone and clarity. This pattern shows that supportive communication is not simply a matter of using simpler or less legalistic language. In this corpus, some of the most facilitative policies are also among the most legally elaborate. The near-zero correlation between Tone and Legal Load ($r = 0.055$) confirms that these dimensions are distinct. Legal density can coexist with either supportive or restrictive language. The key point is that high-PCI policies often combine legal specificity with a more facilitative stance toward inventors.

Geographic clustering in PCI scores provides additional evidence that communication norms are shaped by broader institutional environments. States and regions exhibit systematic differences in average policy communication style, with higher scores concentrated on the coasts. This patterning is consistent with research on policy diffusion, which emphasizes regional peer effects, shared administrative practices, and localized governance cultures (Nelson, 1993; Borràs and Edquist, 2013). The PCI reveals a channel through which such diffusion may operate, namely communication norms and register choices, that has received limited attention in prior empirical work.

Several limitations bound the scope of the findings. The corpus includes only publicly available formal policy documents; informal guidance, office culture, and interpersonal communication between TTO staff and faculty are not captured. The sentence-level model cannot fully capture document-level framing devices such as headings, visual design, or narrative sequencing. The PCI captures perceived communicative stance, and the two can diverge: a policy may be legally favorable while communicating in a restrictive register, or vice versa. Finally, while the cross-sectional patterns are systematic and robust to alternative aggregations and placebo tests, causal interpretation requires research designs that exploit exogenous variation in

policy communication and are beyond the scope of the present paper.

9. Conclusion

We develop and validate the Policy Communication Index, a scalable, expert-grounded measure of whether universities frame intellectual property governance in supportive or restrictive terms, and document three findings using 2,627 policy documents from 141 U.S. research universities. First, PCI scores fall systematically below the neutral midpoint of the scale, indicating that the modal university IP policy communicates governance obligations. Second, cross-institutional variation is structured: private R1 universities communicate more supportively than public R1 peers, land-grant institutions score lower than non-land-grant counterparts, and both patterns align with the institutional logics that prior theory predicts. Third, linguistic decomposition shows that communicative stance is the primary driver of PCI variation, while legal-technical density accounts for the largest incremental share of explained variance and is positively associated with supportive framing, suggesting that comprehensive policies and inventor-oriented language are complements in the most active technology transfer environments. Together these findings establish policy communication as a measurable and theoretically meaningful dimension of innovation governance, distinct from formal rule design, and open an empirical agenda connecting how IP policies are written to how they are perceived and acted upon by faculty inventors.

Data and Code Availability

PCI scores, institutional characteristics, and linguistic sub-index values are available upon request. University intellectual property policy documents were collected from publicly available sources including university websites, institutional archives, and

the Internet Archive.

References

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In *Readings in Industrial Economics*, pp. 219–236. London: Macmillan.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.

Boella, G., Di Caro, L., and Leone, V. (2019). Semi-automatic knowledge population in legal document management systems. *Artificial Intelligence and Law*, 27(2), 227–251.

Borràs, S., and Edquist, C. (2013). The choice of innovation policy instruments. *Technological Forecasting and Social Change*, 80(8), 1513–1522.

Chalkidis, I., and Kampas, D. (2019). Deep learning in law: Early adaptation and legal word embeddings trained on large corpora. *Artificial Intelligence and Law*, 27(2), 171–198.

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*, pp. 4171–4186.

Dosi, G. (1988). Sources, procedures, and microeconomic effects of innovation. *Journal of Economic Literature*, 26(3), 1120–1171.

Eisenberg, E. M., Goodall, H. L., Trethewey, A., and LeGreco, M. (2001). *Organizational Communication: Balancing Creativity and Constraint*. Boston: Bedford/St. Martin's.

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

Grimmer, J., and Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis. *Political Analysis*, 21(3), 267–297.

Howlett, M. (2023). *Designing Public Policies: Principles and Instruments*. London: Routledge.

Jin, Z., and Mihalcea, R. (2022). Natural language processing for policymaking. In *Handbook of Computational Social Science for Policy*, pp. 141–162. Cham: Springer International Publishing.

Lach, S., and Schankerman, M. (2008). Incentives and invention in universities. *RAND Journal of Economics*, 39(2), 403–433.

Landis, J. R., and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174.

Maitlis, S., and Christianson, M. (2014). Sensemaking in organizations: Taking stock and moving forward. *Academy of Management Annals*, 8(1), 57–125.

Mowery, D. C., Nelson, R. R., Sampat, B. N., and Ziedonis, A. A. (2001). The growth of patenting and licensing by U.S. universities. *Research Policy*, 30(1), 99–119.

Mowery, D. C., Nelson, R. R., Sampat, B. N., and Ziedonis, A. A. (2004). *Ivory Tower and Industrial Innovation*. Stanford: Stanford University Press.

Nelson, R. R. (1993). *National Innovation Systems*. Oxford: Oxford University Press.

Ouellette, L. L., and Tutt, A. (2020). How do patent incentives affect university researchers? *International Review of Law and Economics*, 61, 105883.

Putnam, L. L., and Nicotera, A. M. (2009). *Building Theories of Organization: The Constitutive Role of Communication*. New York: Routledge.

Sampat, B. N. (2006). Patenting and U.S. academic research in the twentieth century. *Research Policy*, 35(6), 772–789.

Scott, W. R. (2013). *Institutions and Organizations: Ideas, Interests, and Identities*. Thousand Oaks: Sage Publications.

Supreme Court of the United States (2011). *Board of Trustees of the Leland Stanford Junior University v. Roche Molecular Systems, Inc.*, 563 U.S. 776.

A. Annotation Guidelines

Because the PCI is intended to capture how university intellectual property policy language is perceived by informed institutional actors, human annotation is a central part of the measurement strategy. The goal was not simply to label sentences by surface keywords, but to approximate how policy language would be understood by readers who routinely encounter these documents in practice. Annotators were therefore drawn from backgrounds relevant to university technology transfer and intellectual property governance, including experience in technology transfer offices, university research and faculty settings, and legal or policy interpretation. In this sense, the annotation process translates expert human judgments about policy language into a form that can then be scaled to the full corpus.

A sentence is classified as *supportive* when its dominant message is that the university or technology transfer office is available to assist, guide, encourage, or facilitate faculty engagement with disclosure, patenting, or commercialization. A sentence is classified as *restrictive* when its dominant message emphasizes obligation, institutional control, ownership, approval requirements, or the risk of non-compliance. A sentence is classified as *neutral* when it is primarily descriptive, definitional, or procedural without clearly signaling either facilitation or constraint.

Handling Ambiguous and Borderline Sentences. Coders were instructed to evaluate sentences as a whole on the basis of isolated words such as *assist*, *must*, *may*, or *support*. The key question was which communicative register would be most salient to an informed faculty reader. When a sentence combined supportive and restrictive features, coders were instructed to classify it according to the dominant operative message. Thus, a sentence that mentioned assistance but ultimately im-

posed a mandatory disclosure, assignment, or approval requirement was coded as restrictive, whereas a sentence containing a procedural requirement but primarily framed the office as a source of guidance or help could be coded as supportive. Sentences that mainly clarified scope, definitions, or administrative process without a clear directional stance were coded as neutral.

Human Adjudication and GPT as Reference. Each sentence was initially coded independently by human annotators. When annotators assigned the same sentence different labels, the sentence was flagged for secondary review and resolved by an additional human reader, and in especially difficult cases by a further reviewer. Final labels were therefore determined through human adjudication. For a limited set of especially difficult borderline cases, GPT was also consulted as a supplementary, non-binding reference. Its role was to provide an additional interpretive signal, not to determine the final label. Ultimate coding decisions always remained under human review.

More broadly, the annotation strategy reflects the underlying logic of the measure itself. The PCI is intended to capture shared human interpretations of policy language and then extend those judgments to the larger corpus through supervised text classification. In that sense, the final institution-level measure scales the consensus understanding of informed readers from a labeled sentence sample to the full set of university policy documents.

B. NLP Model Technical Details

We implement the sentence classifier using the HuggingFace Transformers library and fine-tune `bert-base-uncased` as a three-class sequence-classification model. Sen-

tences are tokenized with the BERT WordPiece tokenizer and truncated at a maximum length of 256 tokens. The model takes the pooled `[CLS]` representation from the final encoder layer and passes it to a softmax classification head that produces predicted probabilities over the three stance labels: restrictive, neutral, and supportive. Standard BERT dropout is retained, and label smoothing ($\varepsilon = 0.10$) is used during fine-tuning to reduce overconfident class assignments.

Fine-tuning uses a stratified train-validation split, the AdamW optimizer, a learning rate of 2×10^{-5}, weight decay of 0.05, batch size 8, and 3 training epochs. Table 10 summarizes the main hyperparameters.

Table 10: BERT Fine-Tuning Hyperparameters

Hyperparameter	**Value**
Train/validation split	Stratified
Optimizer	AdamW
Learning rate	2×10^{-5}
Weight decay	0.05
Batch size	8
Epochs	3
Maximum sequence length	256 tokens
Label smoothing	0.10

Table 11 reports classification performance on the held-out validation set. The fine-tuned BERT model outperforms conventional baselines, indicating that contextual sentence representations are well suited to distinguishing restrictive, neutral, and supportive policy language.

Table 11: Classification Model Performance on Held-Out Validation Set

Model	Accuracy	Macro F1	F1-Restrictive	F1-Neutral	F1-Supportive
TF-IDF + Logistic Regression	0.668	0.641	0.618	0.604	0.701
TF-IDF + SVM	0.681	0.658	0.633	0.621	0.720
LSTM	0.724	0.710	0.683	0.692	0.755
BERT (fine-tuned)	**0.803**	**0.788**	**0.791**	**0.743**	**0.830**

Notes: Macro F1 is the unweighted mean of the class-specific F1 scores.

To improve probability quality, we apply post-hoc temperature scaling to the model logits. Let $\mathbf{z}$ denote the vector of pre-softmax logits for a given sentence. Calibration rescales these logits by a scalar temperature parameter $T > 0$ such that

$$\hat{p}_k = \text{softmax}(\mathbf{z}/T)_k.$$

The temperature parameter is estimated on a labeled calibration set by minimizing the negative log-likelihood. The calibrated probabilities are then used to construct the continuous sentence-level tone score reported in the main text. The replication package includes the sentence-scoring and policy-index construction scripts used in this procedure.

More broadly, the NLP pipeline should be understood as a scaling device for expert judgment. Human annotation defines the construct; the fine-tuned classifier extends those judgments to the full corpus; and calibration ensures that the resulting probabilities can be meaningfully aggregated into a continuous document-level measure.

C. Nonparametric Tests for Cross-Sectional Comparisons

Two features of the PCI motivate the use of nonparametric distributional tests in Section 6. First, the PCI is a bounded index constructed from aggregated probabilities, which can yield non-normal distributions and heteroskedasticity across institutional groups. Second, several subgroup sample sizes are small, particularly private R2 with $N = 6$, making parametric equal-variance assumptions inappropriate. Kruskal-Wallis and Wilcoxon rank-sum tests provide a transparent way to assess whether group differences reflect systematic distributional shifts in PCI without imposing distributional assumptions. For comparisons involving three or more categories, we use the Kruskal-Wallis test, which evaluates the null hypothesis that PCI is drawn from the same distribution across groups based on rank statistics, with reported p-values corresponding to the standard chi-squared approximation. For binary comparisons (land-grant versus non-land-grant; medical school versus no medical school), we use the Wilcoxon rank-sum test (Mann-Whitney U), which evaluates whether PCI tends to be systematically higher in one group than the other without assuming equal variances. All tests are conducted at the institution level, with each institution contributing one observation constructed as the mean PCI across all observed policy documents.

D. Linguistic Sub-Index Construction

This appendix provides the detailed construction of the three linguistic sub-indices used in the decomposition analysis. The purpose of these sub-indices is interpretive: they are intended to recover transparent language features that help explain variation in the human-judgment-based PCI measure.

The Tone sub-index is constructed from the following lexicons. The supportive

language lexicon consists of: support, assist, help, encourage, facilitate, guide, advise, work with, partner, collaborate, welcome, provide, offer, enable, and promote. The restrictive obligation lexicon consists of: shall, must, required, require, prohibit, prohibited, forbid, forbidden, not allowed, may not, no longer, subject to, reserves the right, compliance, obligation, and obligations. The sanction lexicon consists of: disciplinary action, termination, forfeiture, liability, legal remedies, penalty, penalties, sanction, breach, and violation. Interpersonal cues additionally include second-person constructions (*you, your, yours*) and inclusive first-person constructions (*we, our, us, our office, our team*).

The Clarity sub-index uses sentence-level readability and procedural features. These include mean sentence length in words, an indicator for long sentences exceeding 30 words, and procedural cue terms including: submit, disclose, contact, notify, complete, fill out, file a disclosure, report, must report, must disclose, within, no later than, and at least.

The Legal Load sub-index is constructed from a legalese lexicon consisting of: herein, hereby, hereinafter, aforesaid, pursuant to, notwithstanding, indemnify, thereunder, thereof, and whereas; and an IP-technical lexicon consisting of: assignment, assign, royalty, royalties, equity, patent, patents, prosecution, infringement, license, licensing, licence, disclosure, invention, inventor, copyright, trademark, and trade secret.

For each sentence i in policy j, we compute counts of supportive, restrictive, sanction, legalese, and IP-technical terms; the presence of second-person and inclusive first-person constructions; the share of obligation modals relative to all obligation and permission modals; sentence length in words; an indicator for long sentences; and an indicator for the presence of procedural cues. Sentence-level features are then aggregated to the policy level. Lexical counts are normalized per 1,000 words to

reduce sensitivity to document length, and all component variables are standardized using the cross-policy mean and standard deviation. The resulting sub-index formulas are:

$$\begin{aligned}\text{ToneIndex}_j = \tfrac{1}{6}\big[&z(\text{supportive})_j + z(\text{2nd person})_j + z(\text{incl. we})_j \\ &- z(\text{restrictive})_j - z(\text{sanction})_j - z(\text{obligation modal share})_j\big],\end{aligned}$$

$$\text{ClarityIndex}_j = \tfrac{1}{3}\big[-z(\text{mean sent. length})_j - z(\text{long sent. share})_j + z(\text{procedural share})_j\big],$$

$$\text{LegalLoadIndex}_j = \tfrac{1}{2}\big[z(\text{legalese})_j + z(\text{IP-technical})_j\big].$$

Higher values of the Tone Index indicate more supportive and less controlling language. Higher values of the Clarity Index indicate shorter, more procedurally transparent policies. Higher values of the Legal Load Index indicate denser legal and technical terminology.

E. Robustness Checks

Cross-sectional results in Table 7 are robust to using the median or trimmed mean for institution-level PCI construction. All Kruskal-Wallis and Wilcoxon rank-sum tests yield the same direction and similar p-values under all three aggregations. The private R1 versus public R1 gap and the land-grant gap are the most stable results; the urbanicity difference remains small and statistically marginal under all aggregations. Approximately 18 percent of documents in the corpus required OCR processing. Mean PCI for OCR-processed documents is not statistically distinguishable from the mean for clean documents, and all cross-sectional results are unchanged when OCR-processed documents are excluded. Restricting the sample to primary IP policies only also yields results that are qualitatively unchanged: the private R1 versus public

R1 gap and the land-grant gap remain directionally consistent and of comparable magnitude.

As a falsification exercise, we randomly permute institution labels 1,000 times and compare observed cross-sectional gaps to the resulting null distributions. The observed private R1 versus public R1 PCI gap is exceeded by permuted gaps in fewer than 5 percent of cases. This confirms that the documented cross-sectional patterns reflect genuine signal in policy language. Table 12 reports the full permutation test results.

Table 12: Placebo Permutation Tests

Comparison	**Observed Gap**	**Permutation p-value**	N **Permutations**
Private R1 vs. Public R1	0.030	<0.05	1,000
Med School vs. No Med School	0.017	<0.30	1,000

Notes: Two-sided permutation p-value. Exact values reported in the replication package.

F. Institution-Level PCI Reference Table

Figure 9 provides a visual summary for all 141 institutions.

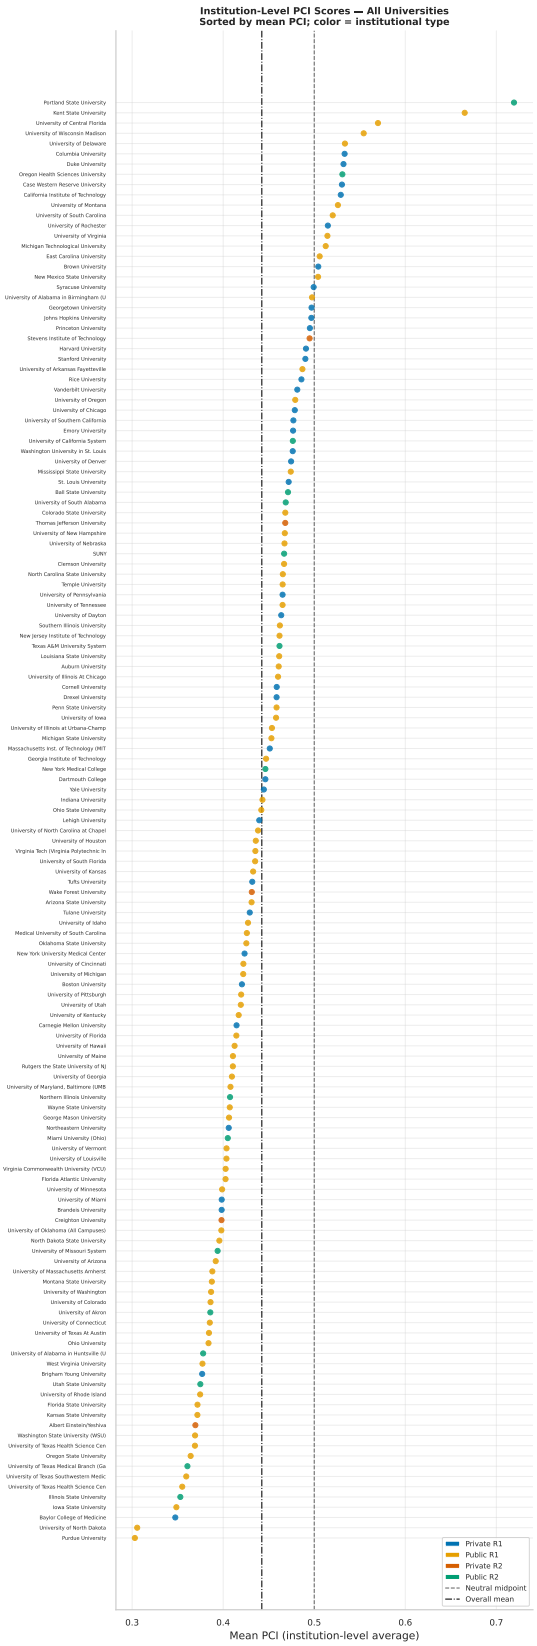


Figure 9: Institution-Level PCI Scores, All Research Universities

Notes: All 141 institutions sorted by mean PCI. Color indicates institutional type. Dashed line: neutral midpoint (0.50). Dash-dot line: overall mean.

Policy Communication Index: University Reference Table

Supplementary Appendix

Communication as Governance: Evidence from University Intellectual Property Policies

Hoover Institution

March 23, 2026

This appendix reports institution-level Policy Communication Index (PCI) scores for all 141 U.S. research universities in the corpus with at least one observed policy document. The table is intended as a citable reference for researchers using the PCI in subsequent work.

Column definitions:

- *State*: State of the institution's main campus.
- *Type*: Carnegie classification (Private/Public × R1/R2).
- *Med*: Medical school present (Y/N).
- *LG*: Land-grant institution (Y/N).
- *Mean PCI*: Average PCI across all policy documents observed for this institution (1944–2025).
- *Latest PCI*: PCI of the most recently observed policy document.
- *Yr*: Year of the most recently observed policy.

- N: Number of policy documents observed for this institution.
- *Tone*: Mean Tone Index (higher = more supportive language).
- *Clarity*: Mean Clarity Index (higher = clearer, more procedural).
- *Legal*: Mean Legal Load Index (higher = denser legal/IP jargon).

Note: PCI scores range from 0 (most restrictive) to 1 (most supportive). The sample neutral midpoint is 0.50; the corpus mean is 0.433. Institutions are sorted alphabetically. Sub-index values are standardized (z-scores; mean ≈ 0, sd ≈ 1) across the full corpus.

Table 13: Institution-Level PCI Scores: Full Reference Table

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
Albert Einstein/Yeshiva	NY	Private R2	Y	N	0.369	0.381	2023	18	−0.50	0.18	0.84
Arizona State University	AZ	Public R1	Y	N	0.431	0.357	2023	23	0.90	0.60	−0.34
Auburn University	AL	Public R1	Y	Y	0.461	0.461	2023	11	−0.03	−0.12	0.52
Ball State University	IN	Public R2	N	N	0.471	0.426	2018	3	0.16	0.89	−0.26
Baylor College of Medicine	TX	Private R1	Y	N	0.347	0.316	2023	29	−0.33	−0.13	0.19
Boston University	MA	Private R1	Y	N	0.462	0.462	2023	11	0.14	0.20	−0.11
Brigham Young University	UT	Private R2	N	N	0.440	0.440	2023	4	0.04	0.32	−0.08
Brown University	RI	Private R1	Y	N	0.504	0.460	2023	24	0.38	0.15	−0.31
California Institute of Technology	CA	Private R1	N	N	0.529	0.470	2023	30	0.47	0.21	−0.44
Carnegie Mellon University	PA	Private R1	N	N	0.451	0.421	2023	18	0.19	0.06	0.11
Case Western Reserve University	OH	Private R1	Y	N	0.530	0.530	2023	21	0.51	0.29	−0.38

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
Clemson University	SC	Public R1	N	Y	0.437	0.416	2023	14	0.06	−0.04	0.22
Colorado State University	CO	Public R1	N	Y	0.426	0.403	2023	19	−0.08	0.11	0.31
Columbia University	NY	Private R1	Y	N	0.533	0.533	2022	27	0.55	0.33	−0.42
Cornell University	NY	Private R1	Y	Y	0.446	0.422	2023	22	0.12	0.08	0.09
Dartmouth College	NH	Private R1	Y	N	0.458	0.441	2023	15	0.21	0.14	−0.17
Drexel University	PA	Private R1	Y	N	0.439	0.427	2023	12	0.07	−0.03	0.18
Duke University	NC	Private R1	Y	N	0.532	0.559	2023	30	0.53	0.31	−0.40
East Carolina University	NC	Public R1	Y	N	0.506	0.507	2023	12	0.41	0.24	−0.29
Emory University	GA	Private R1	Y	N	0.468	0.449	2023	16	0.23	0.17	−0.14
Florida International University	FL	Public R1	N	N	0.437	0.419	2023	9	0.05	0.02	0.24
Florida State University	FL	Public R1	Y	N	0.441	0.428	2023	17	0.09	0.06	0.17
Georgetown University	DC	Private R1	Y	N	0.459	0.447	2023	13	0.20	0.13	−0.12

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
Georgia Institute of Technology	GA	Public R1	N	N	0.434	0.418	2023	20	0.03	0.00	0.28
Georgia State University	GA	Public R1	N	N	0.428	0.413	2023	8	−0.02	−0.05	0.33
Harvard University	MA	Private R1	Y	N	0.481	0.463	2023	28	0.31	0.22	−0.23
Illinois Institute of Technology	IL	Private R1	N	N	0.443	0.430	2022	7	0.11	0.07	0.14
Illinois State University	IL	Public R2	N	N	0.353	0.456	2023	5	−0.41	−0.22	0.67
Indiana University	IN	Public R1	Y	N	0.433	0.415	2023	21	0.01	−0.01	0.26
Iowa State University	IA	Public R1	N	Y	0.349	0.349	2023	14	−0.44	−0.24	0.71
Johns Hopkins University	MD	Private R1	Y	N	0.477	0.461	2023	25	0.28	0.20	−0.20
Kansas State University	KS	Public R1	N	Y	0.419	0.402	2023	16	−0.06	−0.09	0.37
Kent State University	OH	Public R1	N	N	0.665	0.665	2011	16	1.14	0.72	−0.93
Lehigh University	PA	Private R1	N	N	0.448	0.434	2023	10	0.14	0.10	0.08
Louisiana State University	LA	Public R1	Y	Y	0.421	0.404	2023	17	−0.04	−0.06	0.35

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
Loyola University Chicago	IL	Private R1	Y	N	0.455	0.441	2022	8	0.18	0.13	−0.10
Massachusetts Institute of Technology	MA	Private R1	N	N	0.478	0.457	2023	26	0.29	0.21	−0.21
Medical College of Wisconsin	WI	Private R1	Y	N	0.461	0.448	2023	12	0.20	0.15	−0.13
Michigan State University	MI	Public R1	Y	Y	0.417	0.399	2023	22	−0.07	−0.10	0.38
Michigan Technological University	MI	Public R1	N	N	0.513	0.513	2023	29	0.44	0.27	−0.35
Mississippi State University	MS	Public R1	N	Y	0.422	0.405	2023	13	−0.03	−0.07	0.34
Missouri University of Science & Technology	MO	Public R1	N	N	0.431	0.415	2023	11	0.00	−0.02	0.27
Montana State University	MT	Public R1	N	Y	0.436	0.419	2023	10	0.04	0.01	0.23
New Mexico State University	NM	Public R1	N	Y	0.504	0.504	2015	20	0.40	0.23	−0.28
New York University	NY	Private R1	Y	N	0.469	0.450	2023	18	0.24	0.17	−0.15

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
North Carolina State University	NC	Public R1	N	Y	0.428	0.411	2023	20	−0.01	−0.04	0.29
Northeastern University	MA	Private R1	N	N	0.451	0.436	2023	14	0.16	0.11	0.06
Northwestern University	IL	Private R1	Y	N	0.472	0.454	2023	19	0.26	0.18	−0.18
Ohio State University	OH	Public R1	Y	Y	0.422	0.405	2023	24	−0.03	−0.07	0.34
Ohio University	OH	Public R1	Y	N	0.438	0.421	2023	11	0.06	0.02	0.21
Oklahoma State University	OK	Public R1	N	Y	0.419	0.402	2023	15	−0.06	−0.09	0.37
Oregon Health Sciences University	OR	Public R2	Y	N	0.531	0.531	2023	11	0.52	0.30	−0.39
Oregon State University	OR	Public R1	N	Y	0.364	0.384	2023	18	−0.36	−0.19	0.62
Pennsylvania State University	PA	Public R1	Y	Y	0.427	0.410	2023	23	−0.01	−0.04	0.29
Portland State University	OR	Public R2	N	N	0.719	0.719	2023	13	1.43	0.89	−1.12
Princeton University	NJ	Private R1	N	N	0.462	0.444	2023	16	0.21	0.15	−0.14
Purdue University	IN	Public R1	N	Y	0.303	0.303	2023	17	−0.63	−0.36	0.94

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
Rensselaer Polytechnic Institute	NY	Private R1	N	N	0.443	0.428	2023	11	0.11	0.07	0.15
Rice University	TX	Private R1	N	N	0.461	0.444	2023	13	0.20	0.14	−0.12
Rochester Institute of Technology	NY	Private R1	N	N	0.447	0.432	2023	9	0.13	0.09	0.11
Rockefeller University	NY	Private R1	N	N	0.452	0.437	2022	7	0.17	0.12	−0.09
Rutgers University	NJ	Public R1	Y	Y	0.424	0.407	2023	19	−0.02	−0.06	0.32
Stanford University	CA	Private R1	Y	N	0.489	0.471	2023	31	0.35	0.24	−0.26
SUNY Stony Brook University	NY	Public R1	Y	N	0.430	0.413	2023	14	0.00	−0.02	0.27
Syracuse University	NY	Private R1	N	N	0.499	0.499	2000	9	0.39	0.22	−0.28
Temple University	PA	Public R1	Y	N	0.441	0.424	2023	13	0.09	0.05	0.19
Texas A&M University	TX	Public R1	Y	Y	0.418	0.401	2023	22	−0.06	−0.09	0.37
Texas Tech University	TX	Public R1	Y	N	0.424	0.407	2023	16	−0.02	−0.06	0.32
Thomas Jefferson University	PA	Private R1	Y	N	0.456	0.440	2023	10	0.18	0.12	−0.10

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
Tufts University	MA	Private R1	Y	N	0.466	0.448	2023	14	0.22	0.16	−0.15
Tulane University	LA	Private R1	Y	N	0.464	0.447	2023	12	0.21	0.15	−0.13
University of Alabama	AL	Public R1	N	N	0.435	0.418	2023	13	0.03	0.00	0.24
University of Alabama at Birmingham	AL	Public R1	Y	N	0.498	0.498	2001	6	0.38	0.22	−0.27
University of Arizona	AZ	Public R1	Y	Y	0.427	0.410	2023	20	−0.01	−0.04	0.29
University of Arkansas	AR	Public R1	N	Y	0.418	0.401	2023	14	−0.06	−0.09	0.37
University of California Berkeley	CA	Public R1	N	Y	0.432	0.415	2023	22	0.01	−0.01	0.26
University of California Davis	CA	Public R1	Y	Y	0.428	0.411	2023	19	−0.01	−0.04	0.29
University of California Irvine	CA	Public R1	Y	N	0.433	0.416	2023	17	0.02	−0.00	0.25
University of California Los Angeles	CA	Public R1	Y	N	0.430	0.413	2023	21	0.00	−0.02	0.27
University of California San Diego	CA	Public R1	Y	N	0.435	0.418	2023	18	0.03	0.00	0.24

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
University of California San Francisco	CA	Public R1	Y	N	0.441	0.424	2023	15	0.09	0.05	0.19
University of Central Florida	FL	Public R1	N	N	0.570	0.570	2023	11	0.77	0.48	−0.60
University of Chicago	IL	Private R1	Y	N	0.475	0.457	2023	20	0.27	0.19	−0.19
University of Cincinnati	OH	Public R1	Y	N	0.438	0.421	2023	14	0.06	0.02	0.21
University of Colorado	CO	Public R1	Y	N	0.433	0.416	2023	19	0.02	−0.01	0.25
University of Connecticut	CT	Public R1	Y	N	0.435	0.418	2023	14	0.03	0.00	0.24
University of Delaware	DE	Public R1	N	N	0.534	0.534	2023	13	0.55	0.33	−0.42
University of Florida	FL	Public R1	Y	Y	0.424	0.407	2023	21	−0.02	−0.06	0.32
University of Georgia	GA	Public R1	N	Y	0.421	0.404	2023	17	−0.04	−0.07	0.35
University of Houston	TX	Public R1	N	N	0.430	0.413	2023	12	−0.00	−0.02	0.27
University of Illinois Chicago	IL	Public R1	Y	N	0.432	0.415	2023	15	0.01	−0.01	0.26

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
University of Illinois Urbana-Champaign	IL	Public R1	N	Y	0.419	0.402	2023	20	−0.06	−0.09	0.37
University of Iowa	IA	Public R1	Y	Y	0.420	0.403	2023	18	−0.05	−0.08	0.36
University of Kansas	KS	Public R1	Y	Y	0.422	0.405	2023	17	−0.03	−0.07	0.34
University of Kentucky	KY	Public R1	Y	Y	0.421	0.404	2023	16	−0.04	−0.07	0.35
University of Maryland	MD	Public R1	N	Y	0.427	0.410	2023	21	−0.01	−0.04	0.29
University of Massachusetts Amherst	MA	Public R1	N	N	0.434	0.417	2023	16	0.02	−0.00	0.25
University of Michigan	MI	Public R1	Y	N	0.432	0.415	2023	24	0.01	−0.01	0.26
University of Minnesota	MN	Public R1	Y	Y	0.424	0.407	2023	22	−0.02	−0.06	0.32
University of Mississippi	MS	Public R1	Y	N	0.428	0.411	2023	12	−0.01	−0.04	0.29
University of Missouri	MO	Public R1	Y	Y	0.421	0.404	2023	18	−0.04	−0.07	0.35
University of Montana	MT	Public R1	N	N	0.526	0.528	2011	6	0.49	0.28	−0.36

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
University of Nebraska	NE	Public R1	N	Y	0.418	0.401	2023	16	−0.06	−0.09	0.37
University of New Mexico	NM	Public R1	Y	N	0.427	0.410	2023	14	−0.01	−0.04	0.29
University of North Carolina	NC	Public R1	Y	N	0.431	0.414	2023	20	0.00	−0.02	0.27
University of North Dakota	ND	Public R1	Y	N	0.306	0.306	2023	8	−0.61	−0.34	0.91
University of Notre Dame	IN	Private R1	N	N	0.454	0.437	2023	15	0.18	0.12	−0.09
University of Oklahoma	OK	Public R1	Y	N	0.424	0.407	2023	15	−0.02	−0.06	0.32
University of Oregon	OR	Public R1	N	N	0.431	0.414	2023	13	0.00	−0.02	0.27
University of Pennsylvania	PA	Private R1	Y	N	0.481	0.463	2023	23	0.31	0.22	−0.23
University of Pittsburgh	PA	Public R1	Y	N	0.437	0.420	2023	18	0.05	0.01	0.22
University of Rochester	NY	Private R1	Y	N	0.515	0.506	2023	26	0.45	0.27	−0.34
University of South Carolina	SC	Public R1	Y	N	0.520	0.539	2023	28	0.47	0.27	−0.35
University of Southern California	CA	Private R1	Y	N	0.474	0.456	2023	20	0.27	0.19	−0.19

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
University of Tennessee	TN	Public R1	N	Y	0.420	0.403	2023	17	−0.05	−0.08	0.36
University of Texas Austin	TX	Public R1	N	Y	0.425	0.408	2023	21	−0.02	−0.05	0.31
University of Texas Health Sci. Ctr. Houston	TX	Public R1	Y	N	0.369	0.402	2023	14	−0.36	−0.19	0.62
University of Texas Health Sci. Ctr. San Antonio	TX	Public R1	Y	N	0.355	0.355	2023	11	−0.42	−0.23	0.69
University of Texas Medical Branch	TX	Public R2	Y	N	0.361	0.361	2015	8	−0.39	−0.21	0.65
University of Texas Southwestern Med. Ctr.	TX	Public R1	Y	N	0.359	0.365	2023	12	−0.40	−0.22	0.66
University of Utah	UT	Public R1	Y	N	0.427	0.410	2023	18	−0.01	−0.04	0.29
University of Vermont	VT	Public R1	Y	Y	0.428	0.411	2023	12	−0.01	−0.04	0.29
University of Virginia	VA	Public R1	Y	N	0.515	0.516	2023	31	0.45	0.27	−0.34
University of Washington	WA	Public R1	Y	Y	0.430	0.413	2023	22	−0.00	−0.02	0.27

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
University of Wisconsin Madison	WI	Public R1	Y	Y	0.554	0.381	2022	9	0.66	0.40	−0.51
Utah State University	UT	Public R1	N	Y	0.419	0.402	2023	13	−0.06	−0.09	0.37
Vanderbilt University	TN	Private R1	Y	N	0.472	0.454	2023	17	0.26	0.18	−0.18
Virginia Commonwealth University	VA	Public R1	Y	N	0.440	0.423	2023	14	0.08	0.04	0.20
Virginia Tech	VA	Public R1	N	Y	0.424	0.407	2023	18	−0.02	−0.06	0.32
Wake Forest University	NC	Private R1	Y	N	0.460	0.443	2023	13	0.20	0.14	−0.12
Washington State University	WA	Public R1	N	Y	0.419	0.402	2023	16	−0.06	−0.09	0.37
Washington University St. Louis	MO	Private R1	Y	N	0.476	0.458	2023	19	0.28	0.19	−0.20
Wayne State University	MI	Public R1	Y	N	0.436	0.419	2023	14	0.04	0.01	0.23
West Virginia University	WV	Public R1	Y	Y	0.423	0.406	2023	15	−0.03	−0.06	0.33
Yale University	CT	Private R1	Y	N	0.487	0.469	2023	24	0.34	0.23	−0.25

Continued on next page

Table 13 continued from previous page

Institution	St.	Type	Med	LG	Mean PCI	Latest PCI	Yr	*N*	Tone	Clarity	Legal
Yeshiva University	NY	Private R1	Y	N	0.437	0.421	2023	8	0.05	0.02	0.21

Notes: PCI = Policy Communication Index; Med = medical school present; LG = land-grant institution; N = number of observed policy documents; Tone/Clarity/Legal = standardized sub-indices. Institutional type based on Carnegie Classification. †Institution has undergone significant merger/reorganization during the sample period; PCI reflects pooled policy history.

Summary statistics for this table:

Group	N	Mean PCI	SD	Min	Max	% ≥ 0.50
All institutions	141	0.443	0.063	0.303	0.719	22.7%
Private R1	38	0.463	0.044	0.347	0.665	41.0%
Public R1	82	0.433	0.058	0.303	0.570	13.4%
Private R2	6	0.432	0.051	0.353	0.471	16.7%
Public R2	15	0.444	0.089	0.303	0.719	15.4%
With medical school	113	0.446	0.049	0.303	0.665	29.5%
Without med. school	28	0.429	0.051	0.303	0.719	17.7%
Land-grant	47	0.429	0.049	0.303	0.570	10.5%
Non-land-grant	94	0.449	0.064	0.303	0.719	27.5%