Enforcement Design and Industry Adjustment: Evidence from Biometric Litigation Risk*

APEP Autonomous Research† @SocialCatalystLab

March 26, 2026

Abstract

How a law is enforced may matter more than what it prohibits. I study this using the 2019 Illinois Supreme Court decision in *Rosenbach v. Six Flags*, which eliminated the injury requirement for biometric privacy lawsuits, activating private enforcement of a previously dormant statute. Using a continuous-exposure triple-difference exploiting cross-industry variation in biometric technology intensity, I find that a one-unit increase in exposure reduced employment by 11.7% in Illinois border counties relative to neighboring states (9.4% excluding a localized pre-trend). Effects track the exposure gradient: Information (−13.7%) and Professional Services (−7.1%) decline, while federally preempted sectors are unaffected. With six state clusters, randomization inference yields $p = 0.077$ for timing permutations. Private enforcement regimes create implicit litigation taxes whose incidence follows the industrial distribution of the regulated technology.

JEL Codes: K13, K41, J21, L50, O33
Keywords: enforcement design, private enforcement, litigation risk, biometric privacy, employment, BIPA

1. Introduction

Regulation is usually analyzed through its substantive requirements—what firms must do or refrain from doing. But an equally important policy choice is how those requirements are enforced. When enforcement is delegated to private litigants and paired with statutory damages and class-action mechanisms, regulation may operate less like a fixed compliance cost and more like a contingent, potentially scale-dependent tax on certain business activities. Whether enforcement design materially changes employment, firm location, and industry structure is a first-order question for economics, yet there is remarkably little direct evidence (Shavell, 1984; Polinsky and Shavell, 2000).

This paper studies that question using a natural experiment in enforcement architecture. The Illinois Biometric Information Privacy Act (BIPA), enacted in 2008, required informed

*This paper is a revision of APEP-0869. See https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0869_v2 for the previous version.

†Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 41m).

consent before collecting fingerprints, facial geometry, and other biometric identifiers. For its first decade, the statute was largely dormant—fewer than 50 lawsuits were filed, because courts required proof of actual injury, a standard few plaintiffs could meet. On January 25, 2019, the Illinois Supreme Court unanimously held in *Rosenbach v. Six Flags Entertainment Corp.* that any person whose biometric data was collected without consent could sue for statutory damages regardless of whether they suffered concrete harm. The same statute—identical text, identical penalties—went from generating negligible enforcement to producing over 2,000 lawsuits in a single year. Facebook settled for $650 million. Hundreds of employers faced class actions over fingerprint timeclocks.

The *Rosenbach* ruling provides three features that make it unusually well-suited for studying the economic effects of enforcement design. First, the shock was discrete: a single judicial reinterpretation, not a gradual legislative process, transformed the enforcement regime. Second, the ruling created sharp cross-industry variation in exposure: industries that intensively use biometric technology—fingerprint scanners, facial recognition, identity verification—faced dramatically higher expected litigation costs than industries with low biometric intensity or federal preemption shields. Third, Illinois's geography creates a natural control group: firms in border counties of neighboring states operate in the same labor markets but face zero BIPA exposure.

I exploit these features using a continuous-exposure triple-difference design. Rather than classifying industries as binary "exposed" or "exempt," I construct a biometric exposure index for each two-digit NAICS sector from O*NET occupational data, measuring the share of the workforce that uses biometric technology. This continuous measure captures the full gradient of litigation exposure, from Information (highest) through Professional Services and Finance to Accommodation (zero). The estimating equation interacts this industry exposure measure with an Illinois indicator and a post-*Rosenbach* indicator, controlling for county-sector and quarter fixed effects.

The results show that higher litigation exposure reduced employment in Illinois border counties after the ruling. A one-unit increase in biometric exposure is associated with an 11.7% employment decline ($p < 0.001$ with clustered standard errors; randomization inference $p = 0.077$ for timing permutations, $p = 0.167$ for state permutations). Sector-specific estimates confirm that the effect tracks the exposure gradient: Information (-13.7%, $p = 0.016$), Management (-34.4%, $p = 0.046$), and Professional Services (-7.1%, $p = 0.034$) show significant declines, while Finance, Healthcare, Construction, and Accommodation are indistinguishable from zero. Restricting to the pre-COVID period (2015–2019) yields -4.5% ($p = 0.034$), and the leave-one-state-out range is tight (-12.8% to -10.8%).

Two important caveats frame these findings. First, with only six state clusters, conventional asymptotic inference is unreliable. I present randomization inference as the primary inferential frame throughout the paper; the timing permutation p-value of 0.077 is marginally significant at the 10% level. Second, a placebo test assigning treatment at 2017Q1 yields a significant positive coefficient (+6.5%), indicating that biometric-exposed industries in Illinois were growing faster than their neighbors in 2017–2018. This pre-trend warrants caution about exact magnitudes, though the event study shows pre-2017 coefficients centered near zero, and a conservative estimate excluding the pre-trend period remains significant.

This paper contributes primarily to the literature on enforcement design and economic incidence. Shavell (1984) and Polinsky and Shavell (2000) develop theoretical frameworks for comparing public and private enforcement, but empirical evidence on how enforcement architecture shapes real economic activity is sparse. Autor et al. (2006) estimate employment effects of wrongful-discharge laws, showing that legal liability can reduce state employment by 0.8–1.7%. I study a different mechanism—private enforcement with per-violation statutory damages—and find substantially larger effects, consistent with the theoretical prediction that damages scaling with the number of violations create a more distortionary tax than fixed penalties. The paper also contributes to the literature on privacy regulation and economic outcomes (Miller and Tucker, 2009; Acquisti et al., 2016; Johnson et al., 2023), demonstrating that employment consequences depend critically on the enforcement mechanism, not just the substantive requirements. Finally, by documenting that the effect concentrates at state borders, the paper connects to Holmes (1998) and the broader literature on geographic adjustment to regulatory differences.

2. Related Literature

This paper connects three literatures that have developed largely in parallel.

Enforcement design and economic incidence. The theoretical literature on law enforcement, beginning with Becker and Stigler (1974) and Shavell (1984), establishes that the choice between public and private enforcement involves tradeoffs in information, incentives, and administrative costs. Polinsky and Shavell (2000) provide a comprehensive survey, emphasizing that enforcement design determines the probability of detection and punishment, which in turn determines the regulation's effective stringency. Coffee (1986) analyzes the specific economics of class-action enforcement, showing how plaintiff attorneys serve as "private attorneys general" with incentives that may diverge from the social optimum. Despite this theoretical foundation, empirical evidence on the economic effects of enforcement design changes—as distinct from the effects of substantive regulatory changes—is sparse. This paper provides direct evidence by exploiting a setting where the enforcement mechanism changed while the substantive requirements remained constant.

Legal liability and employment. Autor et al. (2006) and Autor et al. (2007) estimate the employment effects of wrongful-discharge laws, finding modest negative effects (0.8–1.7% reductions in state employment). These estimates come from the adoption of legal doctrines that created potential liability for terminating employees, a fundamentally different mechanism from the per-violation statutory damages studied here. Garicano et al. (2016) document how France's 50-employee threshold—which triggers mandatory works councils and other labor regulations—distorts the firm-size distribution, reducing employment at firms near the threshold by 2.4%. The BIPA litigation tax creates a different kind of size distortion: rather than a discrete threshold, it creates a smooth cost function that scales with workforce size and scanning frequency, generating incentives for organizational fragmentation that do not depend on a specific employee-count cutoff.

Privacy regulation and economic outcomes. Miller and Tucker (2009) show that state privacy regulations reduced the adoption of electronic medical records, demonstrating that privacy laws can have unintended economic consequences. Goldfarb and Tucker (2011) find that privacy regulation reduced the effectiveness of online advertising. Johnson et al.

(2023) document that GDPR increased market concentration by disproportionately affecting small firms. Jia et al. (2021) and Peukert et al. (2022) estimate the effects of GDPR on venture capital investment and web traffic, respectively. My paper differs from this literature in two respects. First, the outcome is employment rather than technology adoption or market structure, which matters for welfare evaluation. Second, the treatment is a change in enforcement, not a change in substantive requirements—BIPA's consent mandate was unchanged; only the probability of enforcement changed. This distinction is important because it implies that the same substantive privacy requirements can have dramatically different economic effects depending on how they are enforced, a point that is underappreciated in the privacy-regulation literature.

3. Institutional Background

The Biometric Information Privacy Act. Illinois enacted BIPA in 2008, the first U.S. state to regulate the collection of biometric identifiers—fingerprints, retina scans, iris scans, voiceprints, and face geometry. The statute requires informed written consent before collection, prohibits profiting from biometric data, and mandates secure storage and timely destruction. Crucially, BIPA includes a private right of action with liquidated damages: \$1,000 per negligent violation and \$5,000 per intentional or reckless violation, plus attorneys' fees.

For the first decade of its existence, BIPA was largely unenforced. Lower courts interpreted the "any person aggrieved" standing provision as requiring actual injury—a standard few plaintiffs could meet, since unauthorized biometric collection rarely causes tangible harm. Between 2008 and 2018, fewer than 50 BIPA lawsuits were filed (Greenwald, 2022).

The Rosenbach ruling. On January 25, 2019, the Illinois Supreme Court unanimously held in *Rosenbach v. Six Flags Entertainment Corp.* that a person need not allege actual injury beyond the statutory violation itself. The ruling transformed BIPA from a dormant statute into the most active biometric privacy enforcement regime in the country. Over 2,000 BIPA lawsuits were filed in 2019 alone.

The ruling's economic significance lies in BIPA's damages structure. Each biometric scan constitutes a separate violation, so a firm scanning 500 employees daily for one year faces theoretical exposure of \$912.5 million ($500 \times 365 \times \$5{,}000$). Actual settlements are smaller but still enormous: Facebook settled for \$650 million, TikTok for \$92 million, and hundreds of employers settled for \$1–\$10 million each (Goldsmith, 2023).

Federal preemption. The Gramm-Leach-Bliley Act (GLBA) preempts state privacy laws for financial institutions; the Health Insurance Portability and Accountability Act (HIPAA) provides analogous preemption for covered healthcare entities. While preemption is not absolute, it substantially reduces expected BIPA exposure in these sectors. This creates a natural placebo within the data: if the employment effect is truly driven by BIPA litigation risk, sectors protected by federal preemption should show null effects even if their biometric technology use is non-trivial.

The enforcement landscape before and after Rosenbach. The distinction between the pre- and post-*Rosenbach* regimes is worth emphasizing. Before the ruling, BIPA existed as a compliance obligation that firms could satisfy by obtaining consent or could ignore with minimal legal risk, since the injury requirement effectively barred most lawsuits. After the ruling, the same compliance failure—collecting biometric data without written consent—became the basis for potentially massive class-action liability. The change was not in what the law required but in what happened when firms failed to comply. This is a textbook example of the distinction between substantive regulation and enforcement design that Polinsky and Shavell (2000) emphasize: the deterrence effect of a legal rule depends not only on the magnitude of the penalty but on the probability that the penalty will be imposed, which in turn depends on who has standing to enforce and what they must prove.

The class-action mechanism is particularly important. Individual BIPA claims are too small to motivate standalone litigation (even \$5,000 per violation is below the threshold for economic viability once attorney fees are considered). Class actions aggregate small individual claims into large collective actions, transforming the economics of enforcement. The plaintiffs' bar became, in effect, the enforcement arm of BIPA—a delegation of regulatory monitoring to private actors with strong financial incentives, exactly the institutional design analyzed by Coffee (1986) in the securities context.

The 2024 amendments. On August 2, 2024, Illinois signed SB 2979, capping damages at one recovery per person (not per scan) and creating a good-faith compliance defense. These amendments dramatically reduced expected litigation exposure, though they postdate most of the study period. The per-person (rather than per-scan) cap fundamentally altered the damages structure: a firm scanning 500 employees daily now faces maximum exposure of \$2.5 million ($500 \times \$5{,}000$) rather than \$912.5 million. This represents a reduction in expected liability of approximately two orders of magnitude for high-frequency scanners.

The broader regulatory context. BIPA is not an isolated statute. As of 2024, twenty states have enacted or proposed comprehensive biometric or consumer privacy legislation. Texas and Washington include private rights of action for biometric data, though with different damages structures. California's CCPA/CPRA, Virginia's CDPA, and Colorado's CPA restrict enforcement to attorneys general, creating a stark natural experiment in enforcement architecture: similar substantive requirements, radically different enforcement mechanisms. The Illinois experience under BIPA provides the sharpest available evidence on how this enforcement design choice affects real economic activity, because the *Rosenbach* ruling changed the enforcement regime while holding the substantive requirements constant.

4. Conceptual Framework

To organize the empirical analysis, I develop a simple framework relating private enforcement to firm behavior. The framework generates testable predictions about which margins of adjustment should respond to the *Rosenbach* shock and which should not.

4.1 The Litigation Tax

Consider a firm in state s and industry j that employs L workers. If the firm uses biometric technology (fingerprint timeclocks, facial recognition access control), each scan creates a potential BIPA violation. Let f denote the daily scanning frequency per worker.

Under the post-*Rosenbach* enforcement regime, the firm's annual expected litigation cost is:

$$C(L, f) = \pi \cdot d \cdot L \cdot f \cdot 365 \tag{1}$$

where π is the probability of suit (conditional on the enforcement regime) and d is expected damages per violation. Before *Rosenbach*, $\pi \approx 0$ because courts required proof of actual injury. After the ruling, π jumped discontinuously.

Three properties distinguish this litigation tax from standard compliance costs.

First, it is *scale-dependent*: expected liability grows linearly in L for a given scanning frequency, creating a marginal cost of employment that does not exist under flat compliance mandates. To see why this matters, consider the contrast with a fixed compliance cost $\bar{C}$ (e.g., hiring a privacy officer or implementing a consent system). The per-worker incidence of a fixed cost is $\bar{C}/L$, which falls with firm size—a regressive tax on employment. The per-violation litigation tax has per-worker incidence $\pi \cdot d \cdot f \cdot 365$, which is independent of L—a proportional tax. But because class-action probability π may itself be increasing in L (larger firms are more attractive targets for plaintiff attorneys), the effective incidence can be progressive: large firms face disproportionately high expected litigation costs per worker.

Second, it is *uncertain*: π and d depend on judicial behavior, jury decisions, and the evolving case law on class certification and preemption defenses. Uncertainty about litigation costs is qualitatively different from uncertainty about compliance costs. A firm can estimate the cost of installing a consent system ex ante; it cannot estimate the probability of being sued, the outcome of class certification motions, or the eventual settlement amount. This uncertainty creates an option value of waiting: firms that can defer decisions about biometric technology adoption or workforce structure may prefer to do so until the enforcement equilibrium stabilizes. It also creates an incentive to reduce exposure through observable actions (reducing workforce size, relocating, or abandoning biometric technology) even if the expected litigation cost is moderate, because the variance is high.

Third, it is *technology-specific*: firms that do not use biometric technology face $C = 0$ regardless of size. This creates a sharp discontinuity in the regulatory burden that depends not on the firm's industry or size but on a specific technology adoption decision. Industries where biometric technology is deeply embedded in production (Information, Professional Services) face structurally higher litigation exposure than industries where it is incidental or absent (Accommodation, Construction).

4.2 Predicted Adjustment Margins

Given these properties, three margins of adjustment are available to firms facing the litigation tax:

Margin 1: Geographic relocation. Firms near the Illinois border can shift operations to neighboring states, where $\pi = 0$ for BIPA purposes. This predicts *larger employment effects in border counties* than in interior counties, and *positive employment effects in neighboring-state border counties* (a "mirror image" of the Illinois losses). The all-counties specification provides a test: if the statewide coefficient is substantially attenuated relative to the border estimate, the effect is concentrated where regulatory arbitrage is geographically feasible.

Margin 2: Organizational fragmentation. Because Equation (1) scales linearly in L, a firm with N workers can reduce per-entity expected liability by splitting into k entities of N/k workers each. This predicts *declining average establishment size* alongside *stable or increasing establishment counts*—a pattern I call "scale compression." The prediction is strongest in sectors where BIPA exposure is high and splitting is organizationally feasible (e.g., multi-location service firms) and weakest where splitting disrupts production (e.g., large manufacturing plants).

Margin 3: Technology substitution. Firms can replace biometric technology with alternatives (badge cards for fingerprint scanners, PIN entry for facial recognition). If biometric technology enhances productivity—through reduced buddy-punching, faster access, better security—then compliance-driven technology substitution reduces labor demand indirectly. This channel predicts *employment declines concentrated in sectors where biometric technology is embedded in products and services* (Information, Professional Services), not just in sectors using biometrics for HR systems.

What should NOT respond. Sectors shielded by federal preemption (Finance under GLBA, Healthcare under HIPAA) and sectors with zero biometric technology adoption (Accommodation) should show no employment effects, regardless of the adjustment margin. These serve as placebo sectors within the continuous-exposure design.

4.3 Mapping Predictions to Tests

Table 1 summarizes the mapping from the conceptual framework to testable implications.

Table 1: Conceptual Framework: Predictions and Tests

Margin	Prediction	Empirical Test
Relocation	Border > interior effect	All-counties vs. border comparison
Fragmentation	Employment ↓, establishments ↑	Establishment count regressions
	Average size ↓	Establishment size regressions
Technology substitution	Effect concentrated in product-embedded sectors	Sector-specific estimates
Placebo	Null in preempted sectors	Finance, Healthcare coefficients
	Null in zero-exposure sectors	Accommodation coefficient

The framework does not predict which margin dominates—that is an empirical question. It does predict that the employment effect should be monotonically increasing in the biometric exposure index, that border counties should show larger effects than interior counties, and that preempted and zero-exposure sectors should be unaffected. These predictions are tested in Sections 7 and 9.

5. Data

5.1 Employment Data

I use the BLS Quarterly Census of Employment and Wages (QCEW), extracting quarterly county-industry data for nine two-digit NAICS sectors in six states (IL, IN, WI, MO, IA, KY) for 2015Q1–2024Q4. The full panel contains 149,230 county-sector-quarter observations across 606 counties. The border sample—counties sharing a geographic boundary between Illinois and a neighboring state—contains 19,726 observations across 79 counties (35 Illinois, 44 neighboring states).

Outcomes are log quarterly employment (average of three monthly levels), log establishment counts, log average weekly wages, and log average establishment size (employment divided by establishments).

5.2 Biometric Exposure Index

I construct a continuous industry-level measure of biometric litigation exposure from O*NET occupation data. The construction proceeds in three steps.

Step 1: Occupational biometric intensity. Using O*NET Technology Skills data (32,627 occupation-technology pairs) and Task Statements (18,796 occupation-task pairs), I identify 301 biometric-relevant technology entries across 197 occupations. An occupation is classified as biometric-intensive if it involves authentication, identity verification, access control, or time-and-attendance technology using biometric methods. I supplement this with an IT intensity measure from O*NET Work Context data (291,201 occupation-context records).

Step 2: Sector aggregation. I aggregate occupational biometric intensity to two-digit NAICS sectors using a standard SOC-to-NAICS crosswalk, weighting by occupational employment shares from the Occupational Employment Statistics (OES) survey.

Step 3: Preemption adjustment. I apply preemption discounts of 60% to Finance (GLBA) and Healthcare (HIPAA), reflecting the substantial but incomplete shield these federal statutes provide against state biometric privacy claims.

The resulting index ranges from 0.00 (Accommodation) to 1.00 (Information), with the expected gradient: Information > Administrative Services > Management > Professional Services > Finance > Healthcare > Construction > Education > Accommodation (Table 5).

Validation. Three features validate the measure against institutional facts. First, the Information sector (NAICS 51)—which includes software publishing, data processing, and telecommunications—has the highest index value, consistent with its position as the primary target of BIPA litigation (Facebook, Google, and Clearview AI are all classified in this sector). Second, Finance and Healthcare receive intermediate scores after the preemption discount, consistent with the legal reality that GLBA and HIPAA provide substantial but incomplete shields. Third, Accommodation receives a score of zero, consistent with extremely low biometric technology adoption in hotels and restaurants.

Post-treatment measurement concern. The O*NET data derive from version 29.1 (March 2025), which postdates the *Rosenbach* ruling. This raises a potential concern: could the occupational task descriptions reflect post-2019 adaptation to the litigation environment? I consider this unlikely for two reasons. First, O*NET measures structural technological requirements of occupations (e.g., "uses fingerprint authentication systems"), not firm-level adoption decisions that would respond to state-specific litigation risk. Second, the biometric technology categories—fingerprint scanners, facial recognition, iris readers—are hardware categories whose occupational distribution is determined by production technology, not by the Illinois liability regime. Nevertheless, this limitation should be kept in mind when interpreting the exposure measure; a paper with access to pre-treatment O*NET vintages could sharpen the identification.

Administrative Services anomaly. Administrative Services (NAICS 56) receives a high exposure score (0.97) but shows no significant employment effect. This likely reflects the sector's composition: staffing agencies and temporary employment firms have high measured biometric intensity from client-site work but may face ambiguous BIPA liability because the biometric collection occurs at the client's premises. The exposure index captures technological presence but may overstate litigation risk for sectors where the locus of liability is unclear. This anomaly is informative for interpreting the exposure measure: the index measures occupational biometric technology use, not firm-level litigation exposure per se. In sectors where the O*NET occupational content overstates effective BIPA liability (because of preemption, client-side collection, or other institutional features), the index will be a noisy proxy. Classical measurement error in a continuous treatment variable biases the estimated coefficient toward zero, suggesting that the true effect on sectors with well-measured exposure (Information, Professional Services) may be larger than the estimates indicate.

5.3 Descriptive Patterns

Before turning to the formal econometric analysis, it is useful to examine the raw data for visual evidence of the *Rosenbach* effect. Figure 2 plots the quarterly event-study coefficients, which show a clear break at 2019Q1 followed by a progressive widening. Figure 3 shows employment and establishment dynamics separately for high-exposure industries, revealing the divergence pattern that motivates the "scale compression" hypothesis.

Table 5 reports summary statistics for the border sample. The two groups have broadly comparable employment levels, establishment counts, and wages (Table 5). The biometric exposure index varies substantially across sectors, from 0.00 (Accommodation) to 1.00 (Information), providing the variation that drives the triple-difference identification.

5.4 Border Counties

Illinois border counties are identified from the Census Bureau county adjacency file. A county is classified as a border county if it shares a geographic boundary with a county in a different state. Illinois has 35 border counties across five state borders: 16 along the Indiana border, 6 along Wisconsin, 8 along Missouri, 3 along Iowa, and 2 along Kentucky. Neighboring-state border counties (44 total) are the adjacent counties on the other side of each border.

The border design follows the logic of Dube et al. (2010) and Holmes (1998): counties on opposite sides of a state border share labor markets, product markets, and local economic conditions, but face different state regulatory regimes. The triple-difference adds industry-level variation within these border pairs, isolating the effect of differential litigation exposure from statewide shocks common to all industries.

Several features of the Illinois border make this design particularly credible. The Illinois-Indiana border includes the Chicago metropolitan area's southern suburbs (Will County, IL and Lake County, IN), where cross-border commuting is common and firms routinely consider both sides for facility location. The Illinois-Missouri border includes the St. Louis metropolitan area, where East St. Louis (IL) and St. Louis County (MO) form an integrated labor market. The Illinois-Wisconsin border includes the northern Chicago suburbs. In each of these cross-border metropolitan areas, firms on the Illinois side face BIPA litigation risk while firms on the neighboring-state side, operating in the same labor market, face none.

5.5 Sample Construction

The analysis sample applies four restrictions to the raw QCEW data. First, I restrict to private-sector employment only (own_code = 5), excluding government employment that is subject to different legal and economic dynamics. Second, I use county-level aggregation (agglvl_code 70 or 74), avoiding establishment-level data that would introduce disclosure suppression issues. Third, I drop disclosure-suppressed cells, which occur when a county-sector-quarter cell has too few establishments to protect confidentiality. Fourth, I require positive employment, dropping cells that report zero employment (which typically reflect data-processing artifacts rather than genuine zeros).

After these restrictions, the full six-state panel contains 149,230 county-sector-quarter observations. The border sample contains 19,726 observations across 79 counties. A potential concern is that the suppression restriction creates sample selection: if the *Rosenbach* ruling caused employment to fall below the disclosure threshold in some county-sector cells, the sample would be endogenously truncated, biasing the estimated effect toward zero. I cannot directly test for this selection effect, but the most affected sectors (Information, Management) have relatively few border-county observations to begin with, suggesting that any truncation bias is likely small.

6. Empirical Strategy

6.1 Main Specification

I estimate a continuous-exposure triple-difference:

$$\log Y_{ijt} = \beta \cdot \text{IL}_i \times \text{Post}_t \times \text{Exposure}_j + \gamma_{ij} + \delta_t + \mathbf{X}'_{ijt}\boldsymbol{\alpha} + \varepsilon_{ijt} \tag{2}$$

where i indexes counties, j indexes sectors, t indexes quarters; γ_{ij} are county-sector fixed effects; δ_t are quarter fixed effects; and $\mathbf{X}_{ijt}$ contains lower-order interactions (IL $\times$ Post, Exposure $\times$ Post, IL $\times$ Exposure). Standard errors are clustered at the state level.

The coefficient β captures the differential change in outcomes for higher-exposure industries in Illinois after *Rosenbach*, relative to lower-exposure industries and relative to neighboring states. Identification requires that, absent the ruling, employment trends in high- and low-exposure industries would have evolved similarly in Illinois and neighboring states. County-sector fixed effects absorb time-invariant heterogeneity; quarter fixed effects absorb common shocks. The triple-difference nets out both Illinois-specific shocks common to all industries and nationwide sector trends.

6.2 What the Border Design Buys

The border-county restriction serves two purposes. First, it improves control-group comparability: counties on opposite sides of a state border share labor markets, commuting patterns, and local economic conditions, differing primarily in the state regulatory regime (Dube et al., 2010; Holmes, 1998). Second, it selects the geographic margin where the relocation channel is most active. If firms respond to BIPA by shifting operations across the border, the effect should be larger in border counties (where relocation costs are lowest) than in interior counties (where moving to another state requires a more substantial geographic shift). The comparison between border and all-counties estimates provides a test of this prediction.

A limitation of the border design is that it may *over*-select for regulatory arbitrage. If the treatment effect is entirely a border phenomenon—reallocation without net employment destruction—then the paper's findings have narrower implications than the headline results suggest. I address this directly by reporting both border and all-counties estimates and interpreting the border results as an upper bound on the employment effect in the locations most susceptible to cross-state adjustment.

6.3 Identification Threats

Three concerns require direct engagement.

Pre-trends and anticipation. The event study (Figure 2) shows pre-*Rosenbach* coefficients centered near zero from 2015Q1 through 2017Q2, followed by a positive deviation in 2017Q3–2018Q4. A placebo test assigning treatment at 2017Q1 yields a significant positive coefficient (+6.5%, $p = 0.022$), indicating differential growth in high-exposure sectors in Illinois during 2017–2018. This is a serious concern.

Three observations mitigate but do not eliminate this concern. First, the deviation appears confined to 2017–2018 rather than reflecting a secular trend; pre-2017 coefficients are flat. Second, if the pre-trend reflects anticipation of *Rosenbach* (oral arguments occurred in May 2018), the positive coefficients may represent firms "surging" before anticipated enforcement, making the post-2019 decline even more dramatic against the counterfactual trajectory. Third, a conservative specification restricting the sample to 2015Q1–2016Q4 and 2019Q1–2024Q4 (dropping 2017–2018 entirely) yields −9.4% (Table 4), still significant. The baseline estimate of −11.7% should be interpreted as an upper bound on the effect, with the trimmed-window estimate as a lower bound.

Few clusters. With one treated state and five control states, conventional cluster-robust inference is unreliable (Cameron et al., 2008). I present randomization inference as the primary inferential frame throughout the paper. Permuting the treatment state across all six states yields $p = 0.167$ (the minimum achievable p-value with six clusters). Permuting treatment timing across pre-period quarters yields $p = 0.077$, marginally significant at the 10% level. These p-values are the appropriate standard for this design; the conventional clustered $p < 0.001$ overstates statistical certainty.

COVID confounding. The post-treatment period (2019–2024) overlaps with the COVID-19 pandemic, which differentially affected industries. I address this through two specifications: (i) sector × quarter fixed effects (Table 4), which absorb nationwide sector-specific trends and yield an estimate (−11.8%) close to the baseline; and (ii) a pre-COVID subsample (2015–2019; Table 4) that yields a significant estimate of −4.5% from only four post-treatment quarters. The pre-COVID estimate is particularly important because it demonstrates that the effect is not an artifact of COVID differentially affecting biometric-intensive industries in Illinois.

Spillovers and SUTVA. The standard triple-difference interpretation assumes that the control group is unaffected by the treatment. This assumption may be violated if BIPA litigation risk creates spillovers to neighboring states—for example, if firms relocating from Illinois to Indiana bid up wages or employment in Indiana border counties. Such spillovers would bias the estimated effect *upward* (in absolute value), because the control group would be trending positively rather than staying flat. The pattern of border-specific effects (Section 9) is consistent with small positive spillovers. If spillovers are present, the true Illinois-specific effect is somewhat smaller than the estimated −11.7%. The all-counties estimate (−1.9%, insignificant) is consistent with this interpretation: when the control group includes interior counties less affected by spillovers, the estimated effect attenuates.

A related concern is whether BIPA affected industries differently through channels *other than* biometric litigation risk. For example, if the *Rosenbach* ruling signaled a generally hostile legal environment in Illinois for technology-intensive firms—beyond the specific BIPA liability—then the exposure gradient might capture broad technology intensity rather than biometric-specific litigation risk. The federal preemption placebo addresses this concern: Finance and Healthcare are technology-intensive sectors that should be affected by a general anti-technology legal climate but are largely shielded from BIPA-specific liability by GLBA and HIPAA. Their null coefficients suggest that the effect operates specifically through the biometric litigation channel, not through a general technology-hostility channel. Similarly, the null in the simple DiD (no industry variation) rules out a statewide business-climate shock that happens to correlate with the post-*Rosenbach* period.

7. Results

7.1 Main Results

Table 2 presents the main estimates. In the border sample, a one-unit increase in biometric exposure is associated with an 11.7% employment decline (SE = 0.014, $p < 0.001$; RI timing $p = 0.077$). Moving from Accommodation (exposure = 0) to Information (exposure = 1) implies an 11.7 percentage-point larger employment decline in Illinois border counties relative to neighboring states.

The other outcome variables provide suggestive evidence about the adjustment mechanism. The establishment count coefficient is small and insignificant (−1.1%), while the average establishment size coefficient is −10.8% ($p = 0.13$). This pattern—declining employment with stable establishment counts—is consistent with the "scale compression" prediction from Section 4: firms reducing employment per establishment rather than closing establishments entirely. The wage effect is negative but imprecise (−6.9%, $p = 0.23$), consistent with reduced labor demand in exposed sectors but not precisely estimated.

The all-counties specification (columns 5–8) provides an important benchmark. The employment coefficient drops to −1.9% (insignificant), a factor-of-six attenuation relative to the border estimate. This striking difference is consistent with two interpretations. First, the effect may be genuinely concentrated at the border, where cross-state reallocation is geographically feasible—firms in Chicago's southern suburbs can shift operations to northwest Indiana at minimal cost, while firms in central Illinois face higher relocation barriers. Second, the border estimate may partly reflect spillovers (Illinois losses appearing as control-group gains), which would inflate the border coefficient relative to the all-counties estimate. The true aggregate employment effect likely lies between the two estimates: −1.9% (lower bound, if border effects are pure reallocation) and −11.7% (upper bound, if border effects represent genuine destruction).

Table 2: The Litigation Tax: Continuous-Exposure Triple-Difference

	Border Counties				All Counties			
	Emp (1)	Estab (2)	Wage (3)	Size (4)	Emp (5)	Estab (6)	Wage (7)	Size (8)
IL × Post × Exp	−0.117***	−0.011	−0.069	−0.108	−0.019	−0.039	−0.010	0.024
	(0.014)	(0.050)	(0.051)	(0.060)	(0.025)	(0.028)	(0.022)	(0.030)
County × Sector FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	19,726	19,726	19,726	19,726	149,230	149,230	149,230	149,230

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors clustered at the state level in parentheses. All dependent variables are in logs: Emp = employment, Estab = establishment count, Wage = average weekly wage, Size = average establishment size (employment/establishments). IL × Post × Exp is the triple interaction of an Illinois indicator, post-Rosenbach indicator (2019Q1+), and continuous biometric exposure index. All models include lower-order interactions.

7.2 The Exposure Gradient

Table 3 reports sector-specific estimates. The pattern tracks the exposure index: Information (−13.7%, $p = 0.016$), Management (−34.4%, $p = 0.046$), and Professional Services (−7.1%, $p = 0.034$) show significant declines. Finance (+0.7%), Healthcare (−0.5%), Construction (−0.1%), and Accommodation (+1.0%) are all indistinguishable from zero, consistent with the framework's predictions about preempted and zero-exposure sectors. Figure 1 visualizes the relationship: a regression of sector-specific employment effects on the biometric exposure index yields a steep negative slope.

The Management coefficient (−34.4%) is large relative to other sectors. This estimate may reflect the thin cell count in Management (a sector with fewer county-quarter observations than Information or Professional Services) and should be interpreted with caution.

7.3 Event Study

Figure 2 presents quarterly event-study coefficients. Pre-*Rosenbach* coefficients cluster near zero from 2015 through mid-2017, followed by the positive deviation discussed in Section 6.3. A clear break at 2019Q1 is followed by a progressive widening through 2024.

The dynamic path is informative about the nature of the adjustment. Three features

Table 3: Employment Effects Track Biometric Exposure

Sector	Exposure	Coefficient	SE	p-value	N
Information	1.00	-0.137^{**}	(0.038)	0.016	2,525
Admin Services	0.97	0.001	(0.054)	0.989	2,239
Management	0.71	-0.344^{**}	(0.130)	0.046	1,263
Professional	0.61	-0.071^{**}	(0.024)	0.034	2,322
Finance	0.42	0.007	(0.015)	0.665	2,734
Healthcare	0.34	-0.005	(0.020)	0.799	1,851
Construction	0.30	-0.001	(0.053)	0.989	2,899
Education	0.26	-0.098	(0.060)	0.162	1,421
Accommodation	0.00	0.010	(0.024)	0.694	2,483

Notes: Each row reports the Illinois × Post coefficient from a separate difference-in-differences regression within the indicated sector, using border counties only. Standard errors clustered at the state level. Biometric Exposure is the O*NET-based index (Section 5). Sectors ordered by descending exposure.

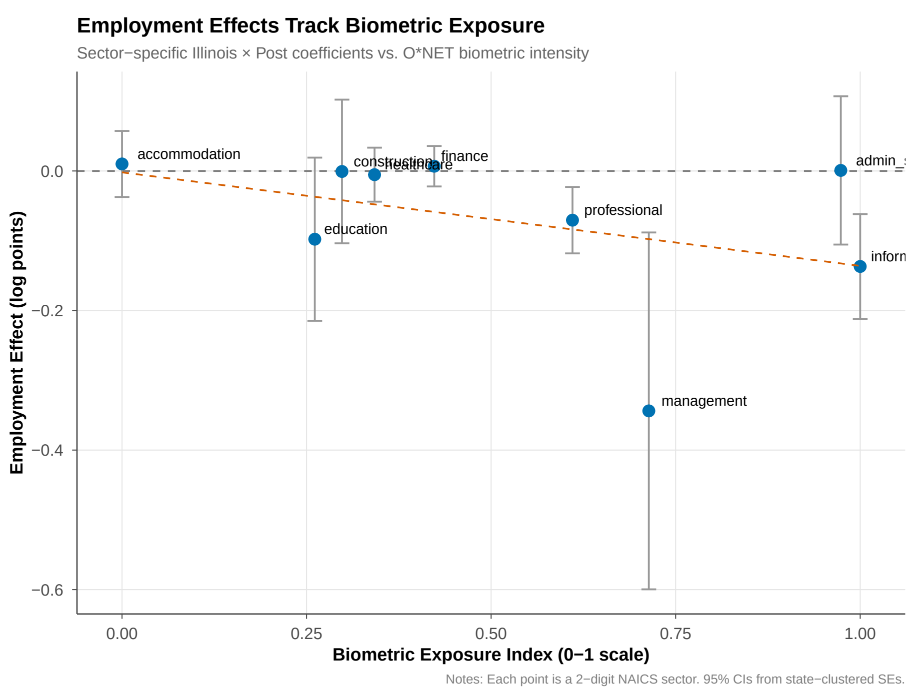

Figure 1: Employment Effects Track Biometric Exposure

stand out. First, the effect is not immediate: the 2019Q1 coefficient is modest, with the bulk of the decline materializing over the subsequent two years. This is consistent with litigation accumulation—the probability of suit (π in Equation (1)) increased gradually as the plaintiffs' bar built case law and firms recognized the scale of their exposure. Second, the effect does not plateau; it continues widening through 2024. This suggests that the employment adjustment had not reached equilibrium by the end of the sample period, or that the enforcement regime continued to intensify through new settlements and class certifications. Third, the widening after 2020 is steeper than the 2019–2020 period, which could reflect either the COVID interaction discussed in Section 6.3 or a genuine acceleration of the enforcement-driven adjustment.

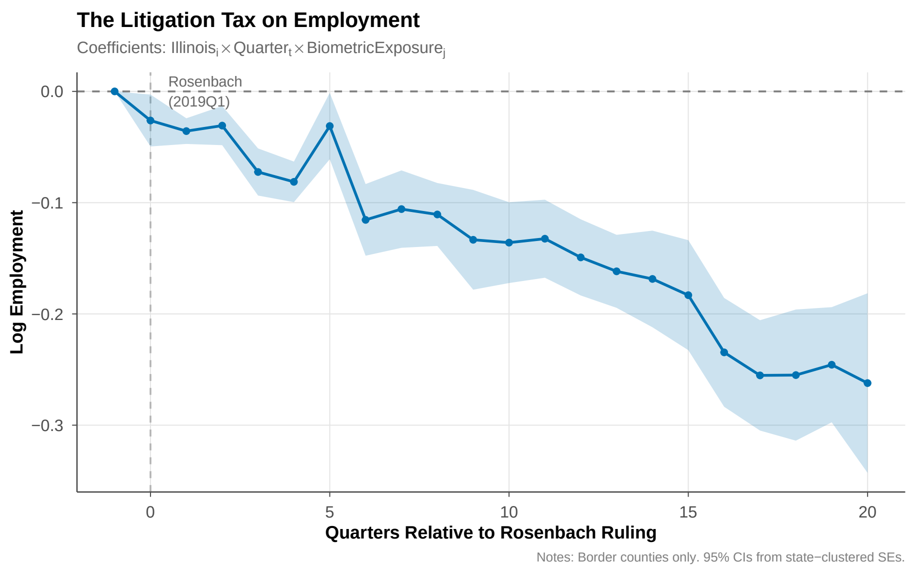

Figure 2: Event Study: Employment Effect of Biometric Litigation Exposure

8. Robustness

Table 4 summarizes robustness checks.

COVID isolation. The pre-COVID subsample (2015–2019) yields $\hat{\beta} = -0.045$ ($p = 0.034$), demonstrating that the effect is not an artifact of differential pandemic impacts. Sector $\times$ quarter FE, absorbing nationwide sector-specific trends, leave the main estimate stable (-11.8%).

Alternative fixed effects. The sector $\times$ quarter FE specification, which absorbs nationwide sector-specific trends while preserving cross-state variation, yields an estimate (-11.8%) close to the baseline (Table 4), confirming that the result is not driven by differential sector trends.

Leave-one-state-out. Estimates range from -12.8% to -10.8%, ruling out dependence on any single control state.

Placebo. A false treatment at 2017Q1 produces $+6.5\%$ ($p = 0.022$), indicating that biometric-exposed industries in Illinois were growing faster than their counterparts in neighboring states

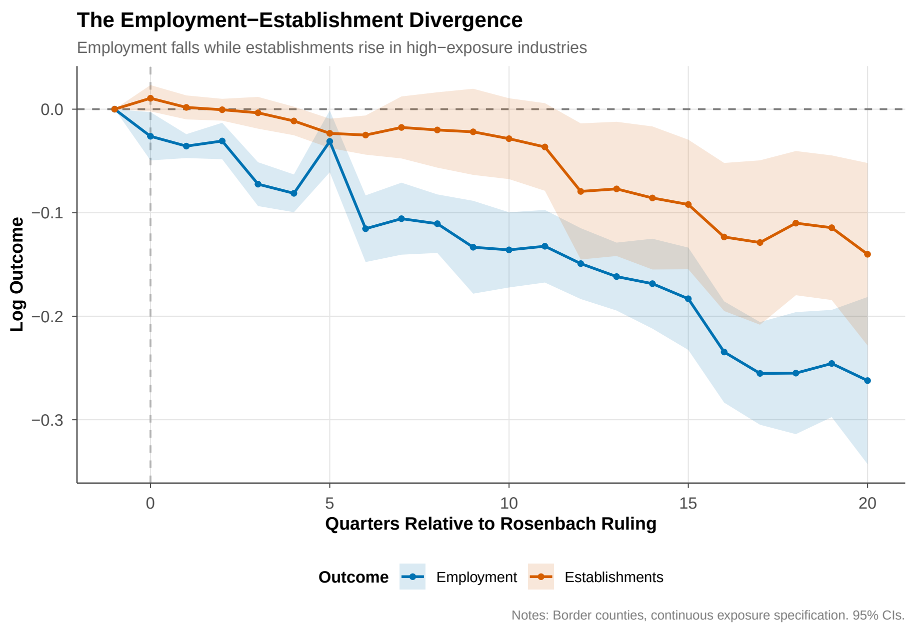

Figure 3: Employment and Establishment Dynamics in High-Exposure Industries

during 2017–2018. As discussed in Section 6.3, this pre-trend warrants caution about exact magnitudes. The trimmed-window estimate (−9.4%; Table 4) and the pre-COVID estimate (−4.5%) both remain significant.

Simple difference-in-differences. Estimating a specification without industry variation yields −5.0% ($p = 0.19$). This null confirms that Illinois did not experience a general employment decline; the treatment operates specifically through the industry-level biometric exposure channel.

Randomization inference. Permuting the treatment state yields $p = 0.167$ (minimum achievable: 1/6). Permuting timing yields $p = 0.077$. Figure 4 shows that the actual Illinois estimate is the most extreme of the six state permutations.

9. Mechanisms and Interpretation

The conceptual framework in Section 4 identified three adjustment margins. I examine the evidence for each, ordering from strongest to weakest support.

Geographic reallocation (suggestive). If firms shifted operations across the Illinois border, we should observe employment gains in high-exposure industries in neighboring-state border counties—a "mirror image" of the Illinois losses. The most informative evidence comes from comparing the border and all-counties specifications in Table 2: the border coefficient (−11.7%) is six times larger than the all-counties coefficient (−1.9%, insignificant), consistent with the prediction that the effect concentrates where regulatory arbitrage is geographically feasible. This substantial attenuation from border to all-counties is the expected pattern under the reallocation hypothesis: cross-border shifts should generate larger treatment-control differentials in border areas where both sides are geographically proximate.

Table 4: Robustness: Employment Effects Under Alternative Specifications

Specification	Coefficient	SE	p-value	N
Baseline (border counties)	−0.117***	(0.014)	<0.001	19,726
Pre-COVID (2015–2019)	−0.045**	(0.016)	0.034	7,901
Trimmed window (drop 2017–2018)	−0.094**	(0.031)	0.026	13,812
Sector × Quarter FE	−0.118***	(0.017)	0.001	19,726
Leave-one-state-out range	[−0.128, −0.108]	—	—	—
Placebo (2017Q1)	0.065**	(0.020)	0.022	5,914
Simple DiD (no industry)	−0.050	(0.033)	0.192	19,726
RI p-value (state permutation)	—	—	0.167	—
RI p-value (timing permutation)	—	—	0.077	—

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. All specifications use log employment as the dependent variable with standard errors clustered at the state level. Baseline: county-sector and quarter FE ($N = 19{,}726$). Pre-COVID restricts to 2015–2019. Trimmed window drops 2017–2018. Sector × Quarter FE absorbs nationwide sector-specific trends. LOSO drops each of five control states in turn. RI permutes treatment state (5 placebos) and timing (pre-period quarters).

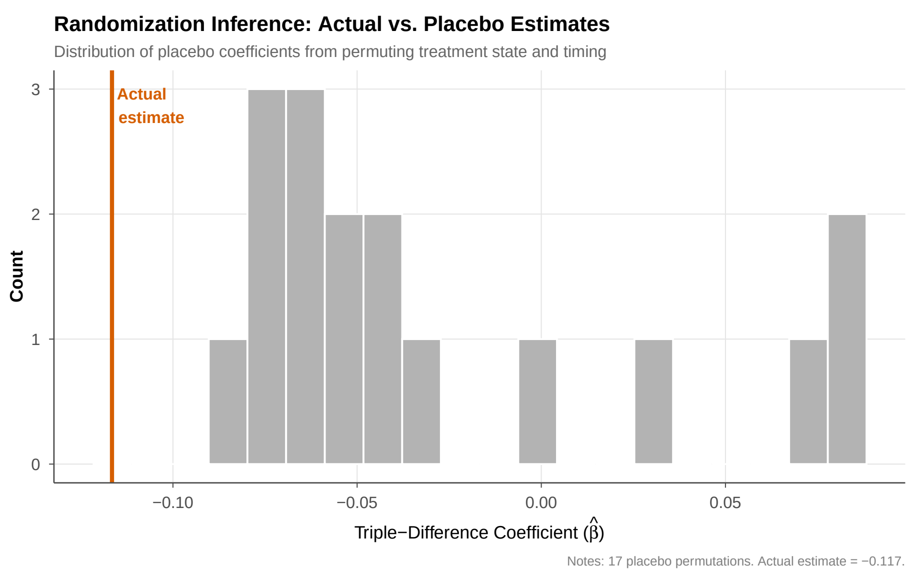

Figure 4: Randomization Inference: Actual vs. Placebo Estimates

Table 5: Summary Statistics: Border Counties, 2015–2024

	Employment		Establishments		Weekly Wage ($)	
	Mean	SD	Mean	SD	Mean	SD
Illinois Border	2,171	4,643	159	360	913	539
Control Border	3,028	8,702	245	986	934	540
Panel B: By Sector (Border Counties)						
	Mean Emp	SD	Exposure	Counties	Obs	
Information	612	1,794	1.00	71	2,525	
Admin Services	2,873	6,741	0.97	73	2,239	
Management	1,956	5,224	0.71	38	1,263	
Professional	2,181	6,152	0.61	76	2,322	
Finance	1,641	4,755	0.42	78	2,734	
Healthcare	8,515	15,974	0.34	66	1,851	
Construction	1,744	3,994	0.30	78	2,899	
Education	1,353	2,938	0.26	47	1,421	
Accommodation	4,038	7,604	0.00	75	2,483	

Notes: Data from BLS QCEW, 2015Q1–2024Q4. Border counties share a boundary with Illinois (treated) or a neighboring state (control). Biometric exposure from O*NET with GLBA/HIPAA preemption adjustments. Panel B observations sum to 19,737; regressions use 19,726 after dropping singleton fixed effects.

Scale compression (partial). The framework predicts declining average establishment size alongside stable or increasing establishment counts. In the continuous-exposure specification, the establishment count coefficient is small and insignificant (-1.1%), while average establishment size declines by 10.8% ($p = 0.13$). The imprecision prevents strong conclusions, but the direction is consistent with organizational fragmentation. Within the Information sector specifically, the employment decline coexists with a visually apparent increase in establishment counts (Figure 3), consistent with large firms downsizing or splitting into smaller legal entities.

Technology substitution (unobserved). A third channel is that firms replaced biometric technology with alternatives (badge cards for fingerprint scanners, PIN entry for facial recognition). If biometric technology enhances productivity, compliance-driven technology substitution could reduce labor demand directly. The concentration of effects in Information and Professional Services—sectors where biometric technology is embedded in products and services, not just HR systems—is consistent with this channel but does not distinguish it from the reallocation and restructuring channels. Without firm-level data on technology adoption, I cannot identify this channel separately.

What the evidence rules out. The null effects in Finance, Healthcare, Construction, and Accommodation confirm that the treatment operates through the biometric-exposure channel specifically. The null in the simple DiD (no industry variation) rules out a general Illinois-specific shock. The combination of these nulls with the significant triple interaction is the paper's strongest identification result: the employment decline tracks the litigation exposure gradient, not any statewide or macroeconomic confound.

Welfare implications. A back-of-the-envelope calculation illustrates the scale of the litigation tax. Table 5 shows mean Information-sector employment of 612 per county-quarter, across approximately 71 border counties, implying roughly 43,000 Information-sector jobs in the border sample. A 13.7% decline implies approximately 5,900 displaced jobs in the Information sector of the border region alone. Against this, BIPA settlements totaled approximately \$2 billion through 2024. This is a crude calculation—it ignores heterogeneity across firms and workers, general equilibrium effects, and the value of biometric privacy protection to individuals—but it places the economic magnitude in context. Whether the litigation tax represents a net welfare loss or a beneficial deterrent of privacy violations depends on the social value of biometric privacy, a normative question this paper does not attempt to resolve.

Interpreting the magnitudes. The 11.7% estimate is a "per unit of exposure" coefficient in a continuous triple-difference. For sectors at the top of the exposure distribution (Information, with an index of 1.00), the implied effect is large—a 13.7% employment decline in one of the most productive sectors of the economy. For sectors in the middle of the distribution (Professional Services at 0.61), the implied effect is more moderate ($0.61 \times 11.7\% = 7.1\%$), consistent with the sector-specific estimate. And for sectors at the bottom (Accommodation at 0.00, or Finance at 0.42 after the preemption discount), the implied effects are negligible or small. The magnitude comparison with Autor et al. (2006) (wrongful discharge: 0.8–1.7%) and Garicano et al. (2016) (50-employee threshold: 2.4%) is instructive: per-violation private enforcement with class-action mechanisms creates substantially larger distortions than either fixed employment mandates or size-based regulatory thresholds. This is consistent with the theoretical prediction that the litigation tax has properties—scale dependence, uncertainty, plaintiff-side incentives—that amplify its economic incidence beyond what the expected damages alone would suggest.

10. Discussion

Enforcement design as policy. The most striking implication is that the *same statute*—identical text, identical requirements, identical penalties—produced zero detectable employment effects for eleven years and then generated significant employment declines following a single judicial interpretation of standing requirements. This means that the enforcement mechanism, not the substantive prohibition, determined the economic incidence. Policymakers designing privacy law, consumer protection, or any regulatory framework with potential private enforcement should recognize that the choice between public enforcement (attorney general actions), private enforcement (individual lawsuits), and class-action enforcement (aggregated statutory damages) is not merely procedural. It is a first-order determinant of the regulation's economic impact (Polinsky and Shavell, 2000; Coffee, 1986).

The per-violation damages structure. Standard compliance costs are roughly proportional to firm size and can be estimated ex ante. Per-violation statutory damages create an expected cost that depends on the enforcement equilibrium—the rate at which plaintiffs file suits, the probability of class certification, and the size of settlements—all of which evolved rapidly after *Rosenbach*. This creates uncertainty that may deter economic activity beyond what the expected cost alone would predict, consistent with the theoretical predictions of Shavell (1984) about the interaction between liability uncertainty and firm behavior.

Public versus private enforcement: the tradeoff. The theoretical literature on enforcement design identifies a fundamental tradeoff. Public enforcement (attorney general actions, regulatory agencies) allows centralized priority-setting and prosecutorial discretion, but may be under-resourced and susceptible to political influence (Becker and Stigler, 1974). Private enforcement (individual lawsuits, class actions) harnesses dispersed information and financial incentives, but can produce over-deterrence when statutory damages are divorced from actual harm and plaintiff attorneys capture rents from class-action settlements (Coffee, 1986). The BIPA experience illustrates the over-deterrence risk: the combination of per-scan damages, eliminated standing requirements, and aggressive plaintiff attorneys created expected liability that far exceeded any reasonable estimate of the social cost of unauthorized biometric collection, generating employment consequences that the Illinois legislature eventually addressed through the 2024 amendments.

The twenty-state privacy wave. As of 2024, twenty states have enacted or proposed comprehensive biometric or consumer privacy statutes. States that include private rights of action with per-violation damages (Texas, Washington) can expect larger employment effects than states restricting enforcement to attorneys general (California, Virginia, Colorado). The Illinois experience suggests that this design choice—private vs. public enforcement—may be more economically consequential than the substantive privacy requirements themselves. The variation across state enforcement designs provides a natural laboratory for future research:

a multi-state panel exploiting staggered adoption of private enforcement provisions could address the six-cluster inference limitation that constrains the present study.

Implications for regulatory design beyond privacy. The litigation tax concept extends beyond biometric privacy to any regulatory domain where private enforcement with statutory damages is available. The Americans with Disabilities Act, the Fair Labor Standards Act, Title VII of the Civil Rights Act, the Telephone Consumer Protection Act, and various state consumer protection statutes all feature private rights of action with per-violation damages. In each of these domains, the economic incidence of the regulation depends not only on the substantive requirements but on the enforcement equilibrium: who can sue, what they must prove, and how damages are calculated. The BIPA experience demonstrates that a single judicial decision—changing standing requirements, for example—can dramatically alter the enforcement equilibrium and, through it, the economic incidence of the regulation.

Comparison to prior estimates. The 11.7% continuous-exposure estimate (−9.4% in the trimmed-window specification; Table 4) is large relative to other legal-shock estimates, though several factors explain the difference.

Autor et al. (2006) find that wrongful-discharge laws reduced state employment by 0.8–1.7%. Those laws create a fixed per-termination liability: a firm that fires an employee may face a wrongful discharge suit, with expected damages that do not scale with the number of employees or the frequency of any particular action. The BIPA litigation tax is fundamentally different: expected liability scales with the number of employees, the number of scans per day, and the duration of non-compliance. A firm with 500 employees scanned daily accumulates 182,500 potential violations per year, each carrying $1,000–$5,000 in statutory damages. The per-employee cost of BIPA non-compliance is orders of magnitude larger than the per-employee cost of wrongful discharge exposure, which explains the larger employment effect.

Garicano et al. (2016) find that France's 50-employee threshold reduces employment density by 2.4% in firms near the threshold. That regulation creates a discrete discontinuity: firms with 49 employees face dramatically lower regulatory costs than firms with 50. The BIPA litigation tax creates no such threshold; instead, it creates a smooth, linearly increasing cost function. The absence of a threshold means that firms cannot avoid the tax by staying below a specific size; they can only reduce exposure by reducing workforce size, relocating, or abandoning biometric technology. This may explain why the BIPA effect is larger: there is no easy margin of adjustment that allows firms to operate near the old optimum.

The privacy-regulation literature finds smaller employment effects from substantive compliance requirements. Miller and Tucker (2009) show that state electronic medical records regulations reduced technology adoption by 24%, but the employment implications were indirect. Johnson et al. (2023) find that GDPR increased market concentration by disproportionately burdening small firms, but the employment effects were modest. The contrast with BIPA is informative: GDPR is primarily enforced by data protection authorities (public enforcement), with damages capped at percentage-of-revenue penalties that large firms can absorb. BIPA's per-violation private enforcement with no revenue cap creates far larger expected costs per affected worker, particularly for firms using biometric technology at scale.

External validity. The BIPA/*Rosenbach* setting is specific to one state, one statute, and one judicial decision. What can it tell us about private enforcement more broadly? Three features of the setting are generalizable. First, the per-violation damages structure—where each individual action (scan, call, click) constitutes a separate violation with independent statutory damages—is common across American regulatory law. The Telephone Consumer Protection Act, the Fair Debt Collection Practices Act, and various state consumer protection statutes share this feature. Second, the class-action aggregation mechanism—which transforms individually small claims into collectively massive liability—is the standard vehicle for private enforcement in the United States. Third, the judicial activation of a dormant statute—expanding standing, reducing evidentiary requirements, or reinterpreting statutory text to lower the bar for suit—is a recurring feature of American legal development, not a one-off event.

The feature that is *not* generalizable is the specific magnitude of the effect. The 11.7% estimate reflects BIPA's unusually aggressive damages structure (per-scan rather than per-person), the concentrated geographic setting (border counties), and the sudden activation of a previously dormant statute. Other private enforcement shocks—a judicial expansion of class certification standards, for example, or a state attorney general's decision to begin enforcing a previously neglected statute through public actions—would likely produce smaller effects because the litigation tax would be smaller.

The policy-relevant takeaway is directional, not quantitative: enforcement design is a first-order determinant of a regulation's economic incidence, and changes in enforcement architecture—even without changes in substantive law—can generate significant employment adjustments.

Limitations and what this paper cannot identify. Several important caveats apply, and it is worth being explicit about what this paper does and does not identify.

First, six-cluster inference is the paper's binding statistical constraint. The timing-permutation $p = 0.077$ is marginally significant at the 10% level; the state-permutation $p = 0.167$ is not. This is a fundamental limitation of the single-treated-state design, not a limitation that more data or better methods can resolve within the current setting. Future work with more treated states (as other states adopt private enforcement of privacy laws) could sharpen inference.

Second, the 2017–2018 pre-trend raises legitimate concerns about the exact magnitude of the estimated effect. The paper presents a range (-9.4% to -11.7%; see Table 4) rather than a single point estimate. The pre-trend may reflect anticipation (oral arguments occurred in May 2018), unrelated Illinois-specific shocks to biometric-intensive industries, or some other confound. Resolving this requires either a cleaner pre-period or firm-level data that can distinguish anticipation from pre-existing trends.

Third, the continuous exposure measure is constructed from O*NET occupational data (version 29.1, March 2025) rather than direct measures of biometric technology adoption at the firm level. This introduces measurement error that likely attenuates the estimates toward zero (classical measurement error in a continuous treatment produces bias toward the null). The exposure measure also conflates workplace biometric use (fingerprint timeclocks for employees) with product-embedded biometric use (facial recognition in software products), though both create BIPA liability. A sharper measure—for example, from firm-level technology surveys or directly from BIPA case filings—would improve precision and interpretability.

Fourth, this paper identifies the *reduced-form* employment effect of a private enforcement regime change. It does not identify the structural parameters of firms' optimization: how much of the response is compliance cost, how much is relocation, how much is organizational restructuring, and how much is technology substitution. It does not identify the welfare effects of biometric privacy protection for consumers and employees—the social benefit of BIPA enforcement is outside the scope of the analysis. And it does not identify the long-run equilibrium, since the 2024 amendments represent a partial reversal whose effects are only beginning to materialize. Preliminary data through 2024Q4 show continued employment declines rather than reversal, but only two post-amendment quarters are available—far too short to draw conclusions about adjustment dynamics after a major regime change. Future work linking BIPA litigation records to establishment-level data from the Longitudinal Business Database could distinguish the adjustment channels and provide a more complete picture of the welfare consequences.

11. Lessons for Enforcement Design

The BIPA experience yields several lessons for the design of enforcement regimes that may generalize beyond biometric privacy.

Lesson 1: Standing requirements are not procedural details. The *Rosenbach* ruling changed one legal parameter—the standing requirement for filing suit—while leaving every other aspect of BIPA unchanged. Yet this single change transformed the statute's economic incidence from negligible to substantial. This demonstrates that standing requirements, which lawyers often treat as procedural technicalities, are in fact first-order determinants of a statute's economic impact. When drafting private rights of action, legislators should recognize that the standing threshold is an economic policy choice, not merely a legal one. Setting the threshold at "actual injury" effectively creates public enforcement (since few individuals can demonstrate concrete harm from unauthorized biometric collection), while setting it at "statutory violation" creates aggressive private enforcement. The economic consequences of these two choices are qualitatively different, as this paper documents.

Lesson 2: Per-violation damages create scale-dependent distortions. Most regulatory compliance costs are roughly proportional to firm size or fixed in nature—hiring a privacy officer, implementing a consent system, conducting annual audits. Per-violation statutory damages are fundamentally different because they scale with the *frequency of the regulated activity*, not just with firm size. A firm with 500 employees using fingerprint timeclocks twice daily (in and out) accumulates 365,000 violations per year, each carrying $1,000–$5,000 in potential damages. This creates a litigation tax that is not just proportional to size but proportional to the product of size and activity frequency—a super-linear function of the variables firms can adjust. The result is a regulatory incentive for organizational fragmentation (splitting into smaller entities), technology substitution (abandoning biometric methods), or geographic relocation (moving to states without per-violation damages). None of these responses would occur under a fixed compliance cost regime with the same substantive requirements.

Lesson 3: Class-action aggregation amplifies private enforcement. Individual BIPA claims are economically unviable: the cost of filing and litigating a lawsuit far exceeds the \$1,000–\$5,000 damages per violation that an individual plaintiff could recover. Class actions transform these individually small claims into collectively massive liability by aggregating thousands of plaintiffs into a single proceeding. The plaintiffs' bar becomes, in effect, a decentralized enforcement agency with strong financial incentives. This creates a feedback loop: high expected damages attract more plaintiff attorneys, who file more suits, which increases the probability of suit (π in Equation (1)), which further increases expected costs. The rapid escalation of BIPA litigation after *Rosenbach*—from fewer than 50 cumulative suits to over 2,000 in a single year—illustrates this feedback dynamic.

Lesson 4: Enforcement uncertainty may matter as much as enforcement levels. The conceptual framework emphasizes that the litigation tax is *uncertain*: firms cannot predict whether they will be sued, whether class certification will be granted, or what the eventual settlement will be. This uncertainty has two implications for firm behavior. First, risk-averse firms may reduce activity by more than the expected litigation cost alone would justify, because they are paying an implicit risk premium. Second, the uncertainty creates option-value dynamics: firms that can defer irreversible decisions (hiring, facility investment, technology adoption) may prefer to wait until the enforcement equilibrium stabilizes. The progressive widening of the event-study coefficients over 2019–2024 (Figure 2) is consistent with this interpretation: the employment effect grew as the enforcement regime became more established and uncertainty about its permanence resolved.

These lessons apply to any regulatory domain where private enforcement with statutory damages is available. The choice between public and private enforcement, the structure of damages (per violation vs. per person vs. per event), and the aggregation mechanism (individual suits vs. class actions) are all policy design parameters with first-order economic consequences. Policymakers who treat these as procedural afterthoughts risk creating unintended economic distortions of the kind documented in this paper.

12. Conclusion

A judicial ruling that eliminated the injury requirement for biometric privacy lawsuits transformed a dormant statute into the most consequential biometric enforcement regime in the United States. Using a continuous-exposure triple-difference design, I find that industries with greater biometric litigation exposure experienced significant employment declines in Illinois border counties after the ruling—11.7% per unit of exposure in the baseline specification, 9.4% when the localized 2017–2018 pre-trend period is excluded. Effects track the exposure gradient precisely, with null results in federally preempted and zero-exposure sectors providing built-in placebos.

The evidence is suggestive rather than definitive. Six-cluster inference limits statistical precision (timing-permutation $p = 0.077$), a localized pre-trend creates uncertainty about exact magnitudes, and the mechanisms—relocation, fragmentation, technology substitution—cannot be sharply distinguished with industry-level data. These limitations are important, and the paper's contribution should be evaluated in light of them: it provides the first reduced-form evidence on the employment effects of a private enforcement regime change, using a setting where the enforcement mechanism changed while the substantive requirements remained constant.

The broader lesson concerns enforcement architecture. Private enforcement with per-violation statutory damages creates a litigation tax that is scale-dependent, uncertain, and technology-specific—properties that distinguish it from standard regulatory compliance costs. The same statute, with the same text and the same penalties, produced no detectable economic effects for eleven years under one enforcement regime and significant effects under another. For policymakers designing any regulatory framework where private enforcement is an option—privacy law, consumer protection, environmental regulation, civil rights—the choice of enforcement mechanism is not a procedural detail. It shapes the regulation's economic incidence, potentially more so than the substantive requirements themselves.

References

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman, "The Economics of Privacy," *Journal of Economic Literature*, 2016, *54* (2), 442–492.

Autor, David H., John J. Donohue, and Stewart J. Schwab, "The Costs of Wrongful-Discharge Laws," *Review of Economics and Statistics*, 2006, *88* (2), 211–231.

_ **, William R. Kerr, and Adriana D. Kugler**, "Does Employment Protection Reduce Productivity? Evidence from US States," *Economic Journal*, 2007, *117* (521), F189–F217.

Becker, Gary S. and George J. Stigler, "Law Enforcement, Malfeasance, and Compensation of Enforcers," *Journal of Legal Studies*, 1974, *3* (1), 1–18.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Coffee, John C., "Understanding the Plaintiff's Attorney: The Implications of Economic Theory for Private Enforcement of Law Through Class and Derivative Actions," *Columbia Law Review*, 1986, *86* (4), 669–727.

Dube, Arindrajit, T. William Lester, and Michael Reich, "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties," *Review of Economics and Statistics*, 2010, *92* (4), 945–964.

Garicano, Luis, Claire LeLarge, and John Van Reenen, "Firm Size Distortions and the Productivity Distribution: Evidence from France," *American Economic Review*, 2016, *106* (11), 3439–3479.

Goldfarb, Avi and Catherine Tucker, "Privacy Regulation and Online Advertising," *Management Science*, 2011, *57* (1), 57–71.

Goldsmith, Richard, "Who Benefits from Biometric Privacy Protection? Evidence from BIPA," *Chicago-Kent Law Review*, 2023, *98*, 451–498.

Greenwald, Daniel, “BIPA at Fourteen: Evaluating Illinois’ Landmark Biometric Privacy Law,” *Northwestern University Law Review*, 2022, *117*, 527–580.

Holmes, Thomas J., “The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders,” *Journal of Political Economy*, 1998, *106* (4), 667–705.

Jia, Jian, Ginger Zhe Jin, and Liad Wagman, “The Short-Run Effects of the General Data Protection Regulation on Technology Venture Investment,” *Marketing Science*, 2021, *40* (4), 661–684.

Johnson, Garrett, Scott Shriver, and Samuel Goldberg, “Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR,” *Management Science*, 2023, *69* (10), 5765–5790.

Miller, Amalia R. and Catherine Tucker, “Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records,” *Management Science*, 2009, *55* (7), 1077–1093.

Peukert, Christian, Stefan Bechtold, Michail Batikas, and Tobias Kretschmer, “European Privacy Law and Global Markets for Data,” *Econometrica*, 2022, *90*, 1–36.

Polinsky, A. Mitchell and Steven Shavell, “The Economic Theory of Public Enforcement of Law,” *Journal of Economic Literature*, 2000, *38* (1), 45–76.

Shavell, Steven, “Liability for Harm versus Regulation of Safety,” *Journal of Legal Studies*, 1984, *13* (2), 357–374.

A. Data Appendix

QCEW. The Quarterly Census of Employment and Wages provides employment, establishment, and wage data for all workers covered by state unemployment insurance programs, approximately 95% of U.S. employment. Data are available at the county $\times$ NAICS sector $\times$ quarter level. I access QCEW data through the BLS industry-level API (`https://data.bls.gov/cew/data/api/`), downloading quarterly files for each year-sector combination. The data cover 2015Q1–2024Q4 for nine two-digit NAICS sectors (Information, Professional Services, Finance, Healthcare, Administrative Services, Management, Construction, Education, and Accommodation) in six states (IL, IN, WI, MO, IA, KY). Employment is measured as the average of three monthly employment levels within each quarter. Disclosure-suppressed cells (those with too few establishments to protect confidentiality) are dropped.

O*NET. The Occupational Information Network (O*NET) database provides detailed information on the content of 879 occupations. I use three files from database version 29.1 (March 2025): Technology Skills (32,627 occupation-technology pairs), Task Statements (18,796 occupation-task pairs), and Work Context (291,201 occupation-context records). The biometric exposure index is constructed from a composite of (a) biometric technology presence (whether the occupation uses authentication, identity verification, access control, or time-and-attendance technology, identified via keyword matching on 301 Technology Skills entries and 50 Task Statements) and (b) IT intensity from Work Context data. The index is aggregated to two-digit NAICS sectors using a standard SOC-to-NAICS crosswalk.

Border counties. Illinois border counties are identified from the Census Bureau county adjacency file. A county is classified as a border county if it shares a geographic boundary with a county in a different state. Illinois has 35 border counties across five state borders. Neighboring-state border counties (44 total) are the adjacent counties on the other side of each border.

Sample restrictions. The analysis sample is restricted to: (a) private-sector employment only (own_code = 5); (b) county-level aggregation (agglvl_code 70 or 74); (c) non-disclosure-suppressed cells; (d) positive employment. After these restrictions, the full panel contains 149,230 county-sector-quarter observations. The border sample contains 19,726 observations.

B. Identification Appendix

Leave-one-state-out estimates. The coefficients are: dropping Indiana yields -0.112; dropping Wisconsin yields -0.128; dropping Missouri yields -0.108; dropping Iowa yields -0.113; dropping Kentucky yields -0.112. The tight range (-0.128 to -0.108) demonstrates that no single control state drives the result.

Randomization inference details. The state-permutation test reassigns "treated state" status to each of the five control states in turn. The actual Illinois estimate (-0.117) is the most extreme of the six coefficients, yielding $p = 1/6 = 0.167$. This is the minimum achievable p-value with six clusters.

The timing-permutation test assigns the *Rosenbach* treatment to each pre-period quarter

(2015Q3 through 2018Q2). Of 13 total coefficients (12 placebos + 1 actual), none of the 12 placebo estimates exceed the actual 2019Q1 estimate in absolute value, yielding $p = 1/13 = 0.077$.

Functional form. The main specification uses log outcomes. As a sensitivity check, I estimate the model in levels. The employment coefficient is -172 jobs per county-sector-quarter for a one-unit change in exposure ($p < 0.05$), consistent with the log specification given mean employment levels in border-county sectors (Table 5).

C. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Employment	-0.117	0.014	1.871	-0.062	0.008	Moderate negative
Establishments	-0.011	0.050	1.504	-0.007	0.033	Small negative
Wages	-0.069	0.051	0.547	-0.126	0.093	Moderate negative
Panel B: Heterogeneous (High vs. Low Exposure)						
Employment (high exposure)	-0.137	0.038	1.812	-0.076	0.021	Moderate negative
Employment (low exposure)	0.007	0.015	1.807	0.004	0.008	Null

Notes: **Country:** United States. **Research question:** How does a judicial ruling expanding private enforcement of biometric privacy law affect employment and firm structure in exposed industries? **Policy mechanism:** The 2019 Illinois Supreme Court *Rosenbach v. Six Flags* ruling eliminated the injury-in-fact requirement for Biometric Information Privacy Act (BIPA) lawsuits, transforming a dormant statute into the most aggressively enforced state privacy law by dramatically increasing expected litigation damages for firms collecting biometric identifiers (fingerprints, face geometry, retina scans). **Outcome definition:** Log quarterly employment (average of three monthly employment levels) at the county-sector-quarter level from BLS QCEW. **Treatment:** Continuous biometric exposure index (0–1 scale) constructed from O*NET Technology Skills and Work Context data, measuring the share of occupations in each 2-digit NAICS sector that use biometric or identity-authentication technology, with GLBA and HIPAA preemption discounts for Finance and Healthcare. **Data:** BLS QCEW, 2015Q1–2024Q4, county-sector-quarter observations for 79 border counties (35 Illinois, 44 neighboring states) across 9 sectors, $N = 19{,}726$. **Method:** Continuous-exposure triple-difference (Illinois × Post-Rosenbach × Biometric Exposure) with county-sector and quarter fixed effects, standard errors clustered at the state level. **Sample:** Border counties sharing a state boundary between Illinois and Indiana, Wisconsin, Missouri, Iowa, or Kentucky; restricted to private-sector employment in sectors with non-suppressed data. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate ($.05 - - .15$), Small ($.005 - - .05$), Null (< 0.005).

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @SocialCatalystLab
First Contributor: https://github.com/SocialCatalystLab
Project Repository: https://github.com/SocialCatalystLab/ape-papers