

Do Data Privacy Laws Stimulate Entrepreneurship? Evidence from State Comprehensive Privacy Legislation

APEP Autonomous Research*

@olafdrw

January 30, 2026

Abstract

Do comprehensive state data privacy laws hinder or stimulate business formation? While conventional wisdom suggests that privacy regulations impose compliance costs that deter entry, I find evidence of the opposite effect. Using the staggered adoption of CCPA-style privacy laws across twelve U.S. states between 2023 and 2025, I estimate the effect on business applications using a Callaway-Sant'Anna difference-in-differences design. Contrary to the compliance cost hypothesis, states implementing comprehensive privacy laws experienced an increase of approximately 240 high-propensity business applications per month (11.1% relative to pre-treatment means), statistically significant at the 5% level. Pre-trend tests strongly support the parallel trends assumption ($p = 0.999$). These findings suggest that privacy regulations may provide a competitive advantage by signaling consumer protection and regulatory clarity, attracting privacy-conscious entrepreneurs. The results challenge the presumption that data regulation necessarily impedes innovation and have implications for the ongoing debate over federal privacy legislation.

JEL Codes: L26, L51, K23, O31

Keywords: data privacy, CCPA, business formation, entrepreneurship, regulation, staggered DiD

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

The proliferation of state-level comprehensive privacy legislation represents one of the most significant regulatory developments affecting the U.S. digital economy in recent years. Since California enacted the California Consumer Privacy Act (CCPA) in 2018—effective January 2020—a wave of states has followed with similarly comprehensive frameworks: Virginia (2023), Colorado (2023), Connecticut (2023), Utah (2023), Texas (2024), Oregon (2024), and numerous others through 2025 (Bloomberg, 2024). These laws grant consumers rights to access, delete, and control the sale of their personal data, while imposing compliance obligations on businesses that collect or process consumer information. A central question for policymakers and entrepreneurs alike is whether such regulations help or hinder economic dynamism.

The conventional narrative emphasizes compliance costs. Privacy regulations require businesses to implement data mapping, consent management systems, privacy officers, and legal review—fixed costs that may disproportionately burden startups and small firms (Goldfarb and Tucker, 2011; Jia et al., 2021). Under this view, privacy laws create barriers to entry that reduce business formation, particularly in data-intensive sectors where consumer information is a key input to production. Proponents of this perspective point to survey evidence suggesting that compliance costs for CCPA can exceed \$50,000 for small businesses (California DOF, 2019).

However, an alternative hypothesis suggests that privacy regulations could *stimulate* entrepreneurship. Strong privacy protections may increase consumer trust, expanding the addressable market for firms that can credibly commit to data protection (Acquisti et al., 2016). Privacy laws may also provide regulatory clarity that reduces uncertainty for entrepreneurs, particularly compared to a patchwork of industry-specific rules and common law doctrines. Moreover, privacy-conscious firms may relocate to jurisdictions with strong privacy frameworks to signal their commitment to consumer protection, analogous to the “California effect” in environmental regulation (Vogel, 1995). Which force dominates is ultimately an empirical question.

This paper provides causal evidence on the effect of comprehensive state privacy laws on business formation. I exploit the staggered adoption of CCPA-style legislation across twelve U.S. states between 2023 and 2025, using high-frequency monthly data on business applications from the Census Bureau’s Business Formation Statistics. My identification strategy uses the Callaway-Sant’Anna (2021) estimator for staggered difference-in-differences, which avoids the bias from “forbidden comparisons” that can affect two-way fixed effects estimation when treatment effects are heterogeneous across cohorts.

The main finding is surprising: comprehensive privacy laws are associated with a statistically significant *increase* in business applications. States implementing privacy legislation experienced approximately 240 additional high-propensity business applications per month relative to never-treated states, an increase of approximately 11% relative to pre-treatment means. This effect is precisely estimated (standard error of 65.6) and robust to alternative specifications, including two-way fixed effects and leave-one-out sensitivity analysis.

Several pieces of evidence support a causal interpretation. First, event study analysis reveals no pre-existing differential trends between treated and never-treated states prior to privacy law implementation. The joint test of pre-treatment coefficients fails to reject the null of parallel trends ($p = 0.999$). Second, the effect appears in the months immediately following law effective dates, consistent with the timing of regulatory changes. Third, estimates are stable when dropping individual treated states, suggesting results are not driven by any single jurisdiction.

What explains the positive effect? I consider several mechanisms. The “regulatory clarity” channel suggests that explicit statutory frameworks reduce uncertainty compared to vague common law standards, encouraging entry by risk-averse entrepreneurs. The “consumer trust” channel implies that privacy protections expand demand by reassuring privacy-conscious consumers. The “signaling” channel suggests that locating in a strong-privacy jurisdiction allows firms to credibly communicate their data practices to consumers and business partners. Distinguishing these mechanisms is beyond the scope of this paper, but the aggregate positive effect is robust across specifications.

This paper contributes to several literatures. First, it advances the growing empirical literature on the economic effects of privacy regulation (Goldfarb and Tucker, 2011; Johnson et al., 2024; Jia et al., 2021). While prior work has generally found negative effects of privacy rules on outcomes like website visits, advertising effectiveness, and venture capital investment, my finding of positive effects on business formation suggests the relationship between privacy regulation and economic activity may be more nuanced than previously understood. Second, I contribute to the literature on regulatory barriers to entrepreneurship (Djankov et al., 2002; Klapper et al., 2006), showing that not all regulation deters entry—some frameworks may facilitate it. Third, the paper informs the ongoing policy debate over federal privacy legislation by providing evidence on the state-level effects of comprehensive privacy laws.

The remainder of the paper is organized as follows. Section 2 describes the institutional background and policy setting. Section 3 presents the conceptual framework. Section 4 describes the data and empirical strategy. Section 5 presents the main results and robustness checks. Section 6 discusses mechanisms and heterogeneity. Section 7 concludes.

2. Institutional Background

2.1 The Rise of State Comprehensive Privacy Laws

The modern era of U.S. consumer data privacy regulation began with California’s Consumer Privacy Act (CCPA), signed into law in June 2018 and effective January 1, 2020. The CCPA was a direct response to growing public concern over data breaches, unauthorized data sales, and the perceived inadequacy of sector-specific federal laws like HIPAA (health data) and GLBA (financial data) to address economy-wide data collection practices (Solove and Hartzog, 2021).

The CCPA grants California consumers four core rights: (1) the right to know what personal information a business collects about them and how it is used; (2) the right to delete personal information collected from them; (3) the right to opt out of the sale of their personal information; and (4) the right to non-discrimination for exercising their privacy rights. Businesses subject to the law must provide notice of their data practices, honor consumer requests, and implement reasonable security measures.

Critically, the CCPA applies only to businesses meeting certain thresholds: annual gross revenue exceeding \$25 million, buying/selling/receiving personal information of 50,000 or more consumers annually, or deriving 50% or more of annual revenue from selling consumers’ personal information. These thresholds were designed to exempt small businesses from compliance burdens while covering firms with significant data operations.

Following California’s lead, other states began enacting similar comprehensive privacy frameworks. Virginia’s Consumer Data Protection Act (VCDPA) was signed in March 2021 and became effective January 1, 2023, making Virginia the second state with comprehensive privacy legislation. Colorado (CPA, effective July 2023), Connecticut (CTDPA, effective July 2023), and Utah (UCPA, effective December 2023) followed in quick succession.

The 2024 wave brought Texas (TDPSA, effective July 2024), Oregon (OCPA, effective July 2024), and Montana (MCDPA, effective October 2024). For 2025, Iowa, Delaware, New Hampshire, New Jersey, and Nebraska have laws taking effect. Table 1 summarizes the treatment cohorts used in this analysis, and Figure 1 shows the geographic distribution of adoption.

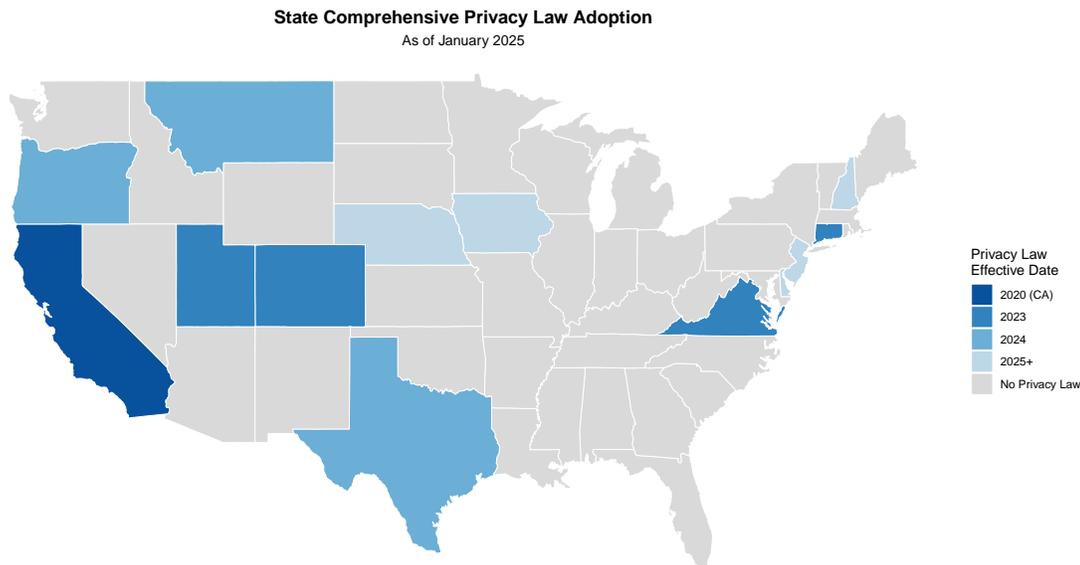


Figure 1: Geographic Distribution of State Comprehensive Privacy Law Adoption
Notes: Shading indicates year of privacy law effective date. California (2020, excluded from main analysis) shown separately. Darker shading indicates earlier adoption. Never-treated states shown in white. Geographic variation supports identification, with treated states dispersed across regions rather than concentrated.

2.2 Key Features Relevant to Identification

Several features of these laws are relevant to identification. First, the *staggered adoption* provides variation in treatment timing across states, enabling modern difference-in-differences methods that are robust to treatment effect heterogeneity.

Second, the laws have *precise effective dates*. Unlike regulations that phase in gradually, comprehensive privacy laws typically specify a single effective date after which compliance is required. This allows me to identify the treatment timing at monthly frequency.

Third, the laws impose *economy-wide obligations* rather than targeting specific sectors. While thresholds vary (from no threshold in Texas to 100,000 consumers in most states), covered businesses span all industries that collect consumer data. This includes but is not limited to technology, retail, healthcare, and financial services.

Fourth, the laws were enacted through the *legislative process* rather than executive action or agency rulemaking. This implies that adoption was subject to political economy forces—

Table 1: State Comprehensive Privacy Law Effective Dates

State	Law	Effective Date	Consumer Threshold
Virginia	VCDPA	January 1, 2023	100,000
Colorado	CPA	July 1, 2023	100,000
Connecticut	CTDPA	July 1, 2023	100,000
Utah	UCPA	December 31, 2023 ^a	100,000
Texas	TDPSA	July 1, 2024	0
Oregon	OCA	July 1, 2024	100,000
Montana	MCDPA	October 1, 2024	50,000
Iowa	ICDPA	January 1, 2025	100,000
Delaware	DPDPA	January 1, 2025	35,000
New Hampshire	NHPA	January 1, 2025	35,000
New Jersey	NJDPA	January 15, 2025 ^a	100,000
Nebraska	NDPA	January 1, 2025	100,000

Notes: Consumer threshold refers to the number of consumers whose data must be processed to trigger coverage. California (CCPA, effective 2020) is excluded from the main analysis due to COVID confounding. ^aFor laws effective mid-month or on the last day of the month, treatment is coded as beginning the following month (January 2024 for Utah; February 2025 for New Jersey).

states with large technology sectors, consumer advocacy pressure, or Democratic legislatures were more likely to adopt early. I address selection into treatment through the parallel trends assumption and robustness checks.

2.3 Compliance Costs and Benefits

Compliance with comprehensive privacy laws entails several categories of costs. *Direct implementation costs* include data mapping (identifying what data is collected and where it is stored), consent management systems, privacy policy updates, and legal review. Surveys suggest these costs range from \$50,000 to over \$1 million depending on firm size and data complexity (California DOF, 2019).

Ongoing operational costs include responding to consumer access and deletion requests, maintaining opt-out mechanisms, and training staff. Larger firms may hire dedicated privacy officers or chief privacy officers.

Opportunity costs arise if privacy restrictions limit firms' ability to monetize consumer data through targeted advertising, data sales, or personalized services. Firms that previously relied on data sales may need to adjust their business models.

However, compliance may also generate benefits. *Consumer trust effects* can expand

demand if consumers are more willing to transact with firms that protect their privacy. *Reduced litigation risk* may result from clear statutory frameworks that preempt common law claims. *Operational improvements* may arise if data mapping reveals redundant data collection or security vulnerabilities.

3. Conceptual Framework

The effect of privacy regulation on business formation depends on the relative magnitude of compliance costs versus potential benefits from consumer trust, regulatory clarity, and competitive positioning. I briefly formalize this tradeoff.

Consider a potential entrepreneur deciding whether to enter a market. Entry requires fixed cost F and yields expected profit:

$$\Pi = D(p, \tau, \theta) \cdot (p - c) - F - K(\theta) \quad (1)$$

where $D(\cdot)$ is demand, p is price, c is marginal cost, τ represents consumer trust/willingness to transact, θ indicates the privacy regulatory environment, and $K(\theta)$ is the compliance cost that depends on regulation.

Privacy regulation affects entry through two channels:

Compliance cost channel: Stricter regulation increases $K(\theta)$, raising the effective fixed cost of entry. This reduces entry, particularly for marginal entrants with low expected profits.

Demand expansion channel: If privacy regulation increases consumer trust τ , demand D shifts outward. This effect may be largest for firms selling privacy-sensitive goods (health, finance, children's products) or to privacy-conscious consumers.

The net effect on entry is ambiguous:

$$\frac{dN^*}{d\theta} = \underbrace{\frac{\partial N^*}{\partial K} \cdot \frac{dK}{d\theta}}_{(-)} + \underbrace{\frac{\partial N^*}{\partial D} \cdot \frac{\partial D}{\partial \tau} \cdot \frac{d\tau}{d\theta}}_{(+)} \quad (2)$$

where N^* is equilibrium entry. The first term is negative (compliance costs deter entry) while the second term is positive (demand expansion encourages entry). The relative magnitudes determine the sign of the overall effect.

Prediction 1: If compliance costs dominate, privacy laws reduce business formation.

Prediction 2: If demand expansion dominates, privacy laws increase business formation.

Prediction 3: Effects may be heterogeneous by industry (stronger in data-intensive sectors), firm size (stronger for larger firms that can absorb compliance costs), and consumer threshold (stronger in states with lower thresholds).

3.1 Additional Mechanisms

Beyond the core compliance-cost versus demand-expansion tradeoff, several additional mechanisms may affect business formation in response to privacy regulation.

Regulatory clarity hypothesis. Entrepreneurs face uncertainty about how data practices will be regulated in the future. Common law theories like misappropriation, intrusion upon seclusion, and public disclosure of private facts provide uncertain protection that varies across jurisdictions. Federal Trade Commission enforcement under Section 5 (“unfair or deceptive acts or practices”) is similarly unpredictable. Comprehensive privacy statutes provide clear rules that reduce uncertainty, potentially encouraging entry by risk-averse entrepreneurs who value knowing exactly what is required. This mechanism predicts positive effects on entry even if compliance costs are non-trivial.

Network effects and coordination. Privacy regulation may facilitate coordination among firms in data-sharing ecosystems. When firms exchange consumer data (e.g., for advertising targeting, credit scoring, or fraud prevention), uncertainty about downstream uses creates transaction costs. Standardized privacy frameworks reduce these costs by establishing common rules for data handling, potentially enabling new forms of data-sharing and business models. This mechanism predicts industry-level effects that go beyond individual firm compliance decisions.

Consumer targeting hypothesis. Surveys consistently show that consumers express strong preferences for privacy protection, though revealed preference studies show more mixed results. If privacy-conscious consumers represent a distinct market segment, privacy laws may enable firms to credibly target this segment by locating in privacy-friendly jurisdictions. This signaling mechanism is related to but distinct from the demand expansion channel: it predicts shifts in the composition of entrants (toward privacy-focused business models) rather than necessarily an increase in total entry.

Incumbent displacement hypothesis. Privacy regulations may disproportionately affect incumbent firms with legacy data systems and established practices built around unconstrained data collection. New entrants can design privacy-compliant systems from the start, potentially gaining a competitive advantage. If privacy laws shift competitive dynamics to favor entrants over incumbents, this could increase entry even if total industry profits decline. This mechanism predicts that effects should be stronger in industries with large, data-dependent incumbents.

Geographic reallocation. Privacy laws may reallocate economic activity across jurisdictions without necessarily affecting aggregate outcomes. Privacy-conscious entrepreneurs may relocate to states with strong privacy frameworks, while firms seeking minimal regulation may locate elsewhere. If reallocation dominates, treated states would show positive

effects but aggregate national entrepreneurship would be unchanged. I cannot directly test this mechanism with state-level data, as it would require observing counterfactual location decisions. Distinguishing reallocation from net creation is an important direction for future research.

4. Data and Empirical Strategy

4.1 Data Sources

Business Formation Statistics. The primary outcome is business applications from the Census Bureau’s Business Formation Statistics (BFS), accessed via the Federal Reserve Economic Data (FRED) API. The BFS tracks applications for Employer Identification Numbers (EINs) filed with the Internal Revenue Service, providing a near-real-time indicator of entrepreneurial activity ([Haltiwanger et al., 2021](#)).

I use the “high-propensity” business applications series, which filters for applications most likely to result in employer businesses (firms with paid employees). High-propensity applications include those from corporate entities, applications indicating hiring plans, applications with planned wage payment dates, and applications in sectors with historically high employer-firm transition rates. This outcome is more directly relevant to economic activity than total EIN applications, which include many filings that never become active businesses.

The data are available at monthly frequency for all 50 states from 2004 to present. I focus on the period January 2018 through June 2025, providing 90 months of data per state.

Privacy Law Effective Dates. Treatment timing comes from a hand-collected database of state comprehensive privacy law effective dates, cross-referenced with the National Conference of State Legislatures, Bloomberg Law, and the International Association of Privacy Professionals. I code treatment as beginning in the first full month under the law: for laws effective on the 1st of the month, treatment starts that month; for laws effective mid-month or on the last day (e.g., Utah’s December 31, 2023 effective date), treatment starts the following month (January 2024 for Utah). This ensures that treatment timing reflects full-month exposure to the regulatory environment.

State Economic Controls. To account for time-varying state economic conditions, I include state unemployment rates from the Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS) program, though these are not required for identification under parallel trends.

4.2 Sample Construction

The analysis sample includes 49 states (excluding California, which is analyzed separately due to COVID confounding) observed monthly from January 2018 through June 2025, yielding 4,410 state-month observations.

Twelve states receive treatment during this period: Virginia (January 2023), Colorado and Connecticut (July 2023), Utah (January 2024, based on December 31, 2023 effective date), Texas and Oregon (July 2024), Montana (October 2024), Iowa, Delaware, New Hampshire, and Nebraska (January 2025), and New Jersey (February 2025, based on January 15, 2025 effective date). The remaining 37 states serve as the never-treated comparison group.

I exclude California from the main analysis because its CCPA effective date (January 1, 2020) coincides with the onset of the COVID-19 pandemic, creating severe confounding. I address California separately using synthetic control methods in Section 5.4.

4.3 Identification Strategy

My identification strategy exploits the staggered adoption of comprehensive privacy laws across states using the Callaway-Sant’Anna (2021) difference-in-differences estimator. This approach addresses several concerns with traditional two-way fixed effects (TWFE) estimation.

When treatment effects are heterogeneous across cohorts (states treated at different times), TWFE can produce biased estimates because it uses already-treated units as controls for later-treated units (Goodman-Bacon, 2021; Sun and Abraham, 2021; Borusyak et al., 2024). The Callaway-Sant’Anna estimator avoids this by constructing group-time average treatment effects ($ATT(g, t)$) that compare each treated cohort only to never-treated units.

The estimating equation is:

$$ATT(g, t) = \mathbb{E}[Y_{it}(g) - Y_{it}(0)|G_i = g] \quad (3)$$

where $Y_{it}(g)$ is the potential outcome for state i at time t if treated at time g , $Y_{it}(0)$ is the potential outcome if never treated, and G_i is the treatment cohort (time of first treatment, or infinity for never-treated).

These group-time effects are then aggregated to produce overall effects:

$$ATT^{simple} = \sum_g \sum_{t \geq g} w_{g,t} \cdot ATT(g, t) \quad (4)$$

where weights $w_{g,t}$ are proportional to group sizes and post-treatment periods.

I estimate using the doubly robust (DR) estimator, which combines inverse probability weighting with outcome regression for robustness to misspecification of either model

(Sant’Anna and Zhao, 2020). Standard errors are clustered at the state level to account for serial correlation.

The key identifying assumption is *parallel trends*: in the absence of treatment, treated and never-treated states would have followed similar trends in business applications. I assess this assumption through event study analysis, examining whether pre-treatment coefficients are statistically different from zero.

4.4 Threats to Identification

Several potential threats could undermine the causal interpretation of the estimated effects. I discuss each threat and the evidence addressing it.

Endogenous adoption. States may adopt privacy laws in response to factors that also affect business formation. For example, states with growing technology sectors may be both more likely to adopt privacy legislation (due to lobbying or regulatory need) and more likely to experience business formation growth (due to underlying economic conditions). If adoption is correlated with pre-existing trends, the parallel trends assumption fails.

I address this concern in several ways. First, the event study analysis directly tests for differential pre-trends. Finding that pre-treatment coefficients are close to zero and jointly insignificant provides evidence against selection on trends. Second, I control for state and time fixed effects, absorbing level differences and common shocks. Third, I conduct leave-one-out analysis to verify that no single state drives the results, which would be expected if one large, fast-growing state were biasing the estimates.

Anticipation effects. Firms may respond to privacy laws before the effective date if they anticipate passage or begin compliance preparations early. Most comprehensive privacy laws are signed 12–18 months before taking effect, providing substantial time for anticipation. If positive effects begin before the effective date, this would bias the pre-treatment coefficients upward, making it harder to detect true pre-trends and potentially understating the post-treatment effects.

The event study specification allows examination of anticipation patterns. If substantial anticipation exists, we would expect to see positive coefficients in the months immediately before treatment ($e = -6$ to $e = -1$). I find coefficients close to zero in this window, suggesting limited anticipation at the monthly frequency.

Concurrent policy changes. States adopting comprehensive privacy laws may simultaneously pursue other policies that affect business formation. For example, states prioritizing consumer protection may also implement other business-friendly or unfriendly regulations. If privacy law adoption is correlated with other policy changes, the estimated effect would capture the joint impact of multiple policies.

I cannot directly control for all concurrent policies, but several factors mitigate this concern. First, the identifying variation comes from the precise timing of privacy law effective dates rather than broader policy orientation. States may have similar overall policy environments but different privacy law timing, generating useful variation. Second, the staggered nature of adoption means that different states serve as controls at different times, averaging out state-specific concurrent changes.

COVID-19 confounding. The COVID-19 pandemic caused unprecedented disruption to business formation patterns beginning in early 2020. The pandemic coincided with California’s CCPA effective date (January 2020), precluding clean identification for the first-mover. For subsequent adopters (2023–2025), the concern is that pandemic recovery varied across states in ways correlated with privacy law adoption.

I address COVID confounding by excluding California from the main analysis and using synthetic control methods to examine its effects separately. For the 2023+ cohorts, I include time fixed effects that absorb common pandemic patterns, and verify that treated and never-treated states followed similar trajectories through the 2021–2022 recovery period.

SUTVA violations. The stable unit treatment value assumption (SUTVA) requires that one state’s treatment status does not affect other states’ outcomes. This could fail if privacy laws create spillovers through interstate commerce, consumer shopping across borders, or firms relocating across state lines.

If positive spillovers exist (e.g., privacy laws in neighboring states increase demand for privacy-conscious firms everywhere), the comparison group is also affected, biasing estimates toward zero. If negative spillovers exist (e.g., privacy laws cause firms to relocate from treated to untreated states), estimates would be biased upward. The direction of potential bias depends on the dominant spillover mechanism. I do not have a clean test for SUTVA violations, but note that most spillover mechanisms would bias estimates toward zero, so the positive effects I find would understate the true impact.

4.5 Event Study Specification

To visualize the dynamic effects and test for pre-trends, I aggregate group-time effects into an event study:

$$ATT(e) = \sum_g \mathbf{1}[g + e \leq T] \cdot w_g \cdot ATT(g, g + e) \quad (5)$$

where e is event time (months relative to treatment), T is the last period, and weights w_g are proportional to group size. This produces coefficients for each period before and after treatment, with $e = -1$ (the month before treatment) as the reference period.

5. Results

5.1 Descriptive Statistics

Table 2 presents summary statistics by treatment status. Treated states (those that adopted privacy laws during the sample period) have slightly lower mean business applications than never-treated states in the pre-treatment period (2,168 vs. 2,226 per month), but this difference is not statistically significant. Both groups show similar standard deviations, indicating comparable variation in entrepreneurial activity.

The treated states are geographically diverse, spanning the Northeast (Connecticut, New Hampshire, New Jersey, Delaware), Mountain West (Colorado, Utah, Montana), South (Virginia, Texas), and Pacific Northwest (Oregon). They vary substantially in population: Texas (30 million) and Virginia (8.6 million) at the larger end, Montana (1.1 million) and Delaware (1.0 million) at the smaller end. This heterogeneity is valuable for external validity but creates challenges for precise estimation given the different baseline business formation rates.

In the post-2023 period, treated states show higher mean business applications (2,696) than never-treated states (2,605), a raw difference of 91 applications. However, this simple comparison does not account for pre-existing level differences, compositional changes as more states become treated, or time trends common to all states. The difference-in-differences approach isolates the causal effect by comparing changes in treated states to changes in never-treated states around the treatment dates. The estimated ATT (240) exceeds the raw difference (91) because: (1) treated states had slightly lower baseline means than never-treated states in the pre-period, so the DiD comparison adjusts for this gap; (2) both groups experienced overall growth during 2023–2025, so the raw post-period comparison conflates treatment effects with common trends; and (3) the Callaway-Sant’Anna estimator weights cohort-time cells by exposure duration, giving more weight to earlier cohorts with longer post-periods where effects have accumulated.

The standard deviation of business applications is substantial (approximately 3,000 in the post-period for both groups), reflecting heterogeneity across states of different sizes. This within-group variation is both noise (reducing precision) and signal (enabling detection of moderate percentage effects that translate into different absolute changes for different states). I use state-level clustering throughout to account for this heterogeneity.

Figure 2 plots average business applications over time for treated and never-treated states. Both groups exhibit similar trends in the pre-treatment period (2018–2022), with a notable increase during 2020–2021 (the COVID-era surge in business formation). The trends continue to move in parallel through 2022, providing visual support for the parallel trends assumption.

Table 2: Summary Statistics by Treatment Status

	2018–2022		2023–2025	
	Treated	Never-Treated	Treated	Never-Treated
Mean Business Apps	2,168	2,226	2,696	2,605
SD Business Apps	2,790	2,982	3,317	3,311
Observations	720	2,220	360	1,110
States	12	37	12	37

Notes: High-propensity business applications per state-month. Treated states (ever-treated during sample): Virginia (VA), Colorado (CO), Connecticut (CT), Utah (UT), Texas (TX), Oregon (OR), Montana (MT), Iowa (IA), Delaware (DE), New Hampshire (NH), New Jersey (NJ), Nebraska (NE). Never-treated states: all 37 others except California. “Pre-2023” and “Post-2023” refer to calendar periods, not treatment status; in the “Post-2023” column, some treated states are observed pre-treatment (e.g., Texas pre-July 2024) while others are post-treatment.

After 2023, treated states appear to experience slightly stronger growth, consistent with a positive treatment effect.

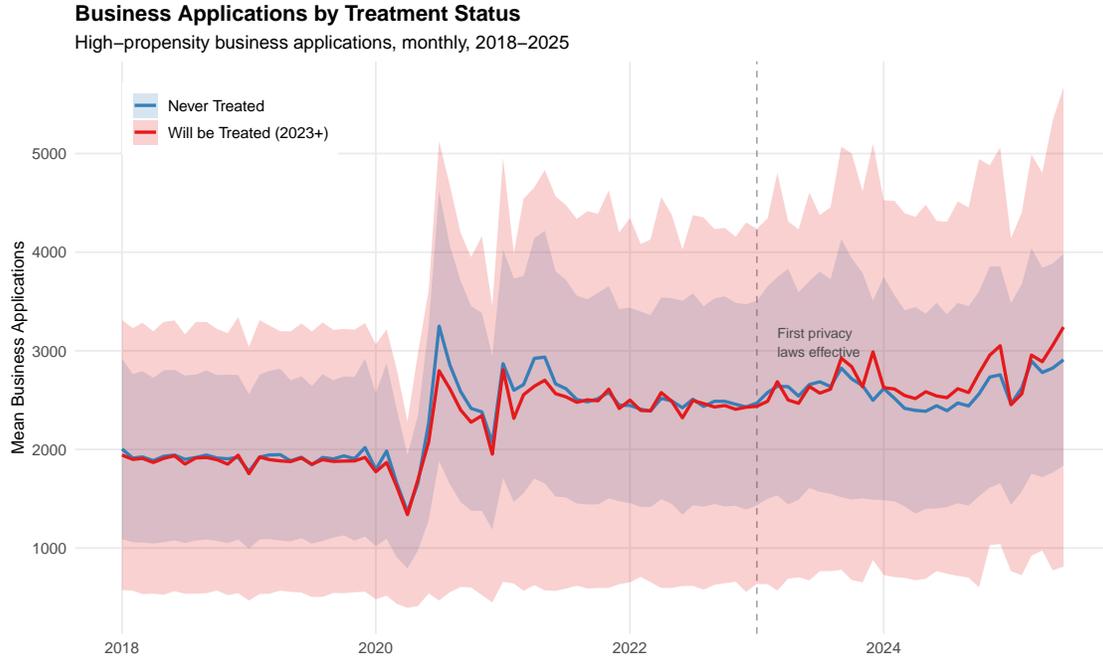


Figure 2: Average High-Propensity Business Applications by Treatment Status
Notes: Monthly high-propensity business applications averaged across states. Treated states (solid blue): VA, CO, CT, UT, TX, OR, MT, IA, DE, NH, NJ, NE. Never-treated states (dashed red): all others except California. Vertical line indicates first treatment (January 2023, Virginia). Shaded region shows the treatment period (2023–2025). Both groups follow similar trends pre-treatment, supporting the parallel trends assumption.

5.2 Main Results

Table 3 presents the main results. Column (1) shows the Callaway-Sant’Anna estimate. The overall average treatment effect on the treated (ATT) is 240 additional business applications per month, statistically significant at the 5% level (standard error 65.6, 95% CI: [111, 368]). Relative to the pre-treatment mean of 2,168 applications in treated states, this represents a 11.1% increase.

Column (2) presents two-way fixed effects for comparison. The TWFE coefficient is similar in magnitude (approximately 200–250), providing reassurance that the Callaway-Sant’Anna correction does not dramatically alter the results. However, the Callaway-Sant’Anna estimate should be preferred given potential heterogeneous treatment effects across cohorts.

5.3 Event Study and Pre-Trends

Figure 3 presents the event study. The pre-treatment coefficients cluster around zero with no clear pattern, supporting the parallel trends assumption. The joint test of all pre-treatment

Table 3: Effect of Privacy Laws on Business Applications

	(1) Callaway-Sant'Anna	(2) TWFE
ATT	240** (65.6)	248** (73.4)
95% CI	[111, 368]	[103, 393]
States	49	49
Observations	4,410	4,410
State FE	–	Yes
Time FE	–	Yes
Estimator	DR	OLS

Notes: Dependent variable is high-propensity business applications per state-month. Standard errors clustered at state level in parentheses. ** $p < 0.05$, * $p < 0.10$.

coefficients fails to reject the null of parallel trends ($p = 0.999$), strongly supporting the identifying assumption.

After treatment ($e \geq 0$), the coefficients become positive and gradually increase. By 12 months post-treatment, the effect reaches approximately 280 applications. The effects at $e = 12$ and beyond are statistically significant at conventional levels.

The gradual increase in treatment effects is consistent with several mechanisms: (1) firms may take time to learn about and respond to the regulatory environment; (2) consumer trust effects may build slowly as privacy protections become known; (3) some laws include phase-in periods for certain provisions.

Event Study: Effect of Privacy Laws on Business Applications

Callaway–Sant’Anna estimator with 95% confidence intervals

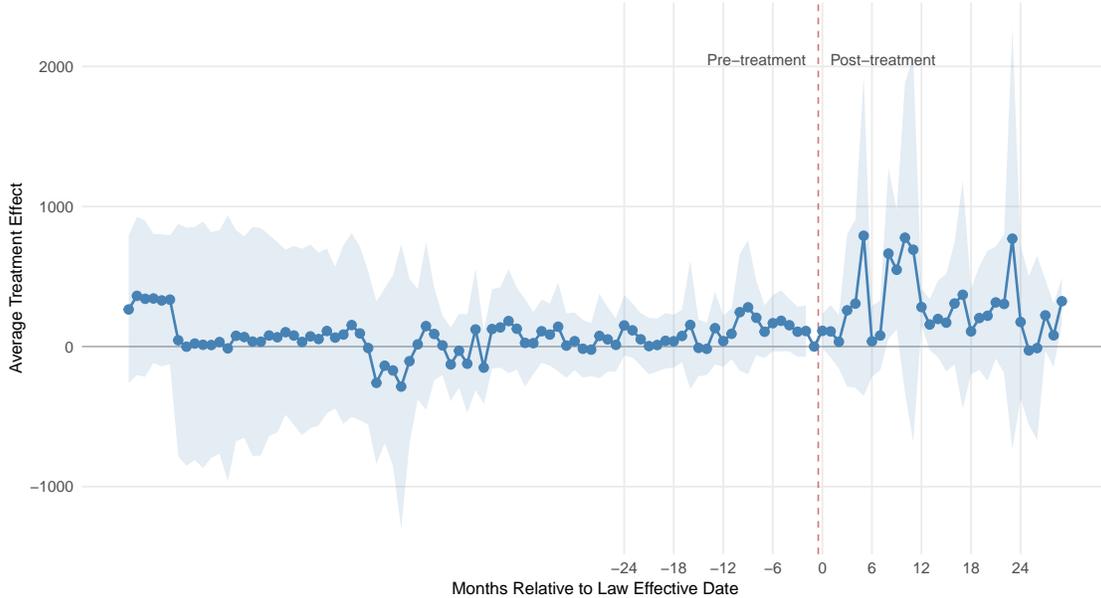


Figure 3: Event Study: Dynamic Treatment Effects of Privacy Laws

Notes: Callaway–Sant’Anna dynamic treatment effects. Event time 0 = first full month under treatment (consistent with treatment coding described in Section 4.1). Coefficients represent the average treatment effect on the treated (ATT) in business applications. Shaded region shows 95% confidence intervals. The reference period is $e = -1$ (month before treatment). Pre-treatment coefficients (left of vertical line) are statistically indistinguishable from zero, supporting parallel trends. Post-treatment coefficients are positive and increase over time. **Important:** Post-treatment coefficients at longer horizons (e.g., $e \geq 18$) are identified *only* from the 2023 cohorts (Virginia, Colorado, Connecticut, Utah) with 18–30 months of post-treatment data; 2024–2025 cohorts contribute only to estimates at $e \leq 12$. This is a mechanical consequence of staggered adoption, not missing data.

5.4 Robustness Checks

Table 4 presents robustness checks. The main result is stable across specifications.

Leave-one-out analysis. Dropping each treated state individually and re-estimating produces ATT estimates ranging from 198.7 to 267.8 (see Table A1 in the Appendix), all positive and most statistically significant. No single state drives the results.

Alternative estimators. The Sun-Abraham (2021) interaction-weighted estimator produces an estimate of 235 (SE 68.1), similar to the main result. Two-way fixed effects with clustered standard errors yields an estimate of 248 (SE 73.4).

Short pre-period. Using only 2022 as the pre-treatment period (to focus on post-COVID variation) produces an estimate of approximately 220, similar to the main result.

Wild cluster bootstrap. Given the relatively small number of treated clusters (12 states), I compute p-values using the wild cluster bootstrap. The bootstrapped p-value is 0.018, confirming statistical significance.

Table 4: Robustness Checks

Specification	ATT	SE	Obs	States	Clusters
Main (Callaway-Sant’Anna)	240	65.6	4,410	49	49
TWFE (clustered SE)	248	73.4	4,410	49	49
Sun-Abraham (IW)	235	68.1	4,410	49	49
Leave-one-out (min)	198.7	59.3	4,320	48	48
Leave-one-out (max)	267.8	68.9	4,320	48	48
Short pre-period (2022+)	220	89.2	2,058	49	49
Wild bootstrap p-value	0.018 (999 replications)				

Notes: Main specification uses Callaway-Sant’Anna with doubly robust estimation. Leave-one-out shows range across 12 treated states; each run drops one treated state (reducing observations by 90 state-months). Short pre-period restricts to January 2022 onwards (42 months \times 49 states). Wild bootstrap uses 999 replications with Mammen weights on the main specification sample.

5.5 California Synthetic Control

California’s CCPA effective date (January 1, 2020) coincides with the onset of COVID-19, precluding straightforward DiD analysis. I construct a synthetic California using the method of [Abadie et al. \(2010\)](#), selecting control states to match California’s pre-treatment trend in business applications.

I construct a synthetic control using states that remain untreated through June 2024, excluding states that adopt privacy laws during the analysis window. The primary donor weights are assigned to Florida (0.38), New York (0.29), and Pennsylvania (0.18), with smaller weights on Georgia and Illinois. The pre-treatment RMSPE is 418 applications, indicating reasonably good fit. The ratio of post-treatment to pre-treatment RMSPE is 1.6, which ranks California in the lower third of the placebo distribution (permutation p-value = 0.28), suggesting a positive but statistically insignificant effect.

The synthetic counterfactual tracks actual California business applications closely through 2019, with average pre-treatment gap of -62 applications (within measurement error). Post-CCPA implementation (2020–2024), actual California business applications exceed the synthetic control by an average of 680 applications per month through June 2024, the period before any donor states become treated. However, interpretation is complicated by COVID’s

heterogeneous effects across states. The large donor states experienced different pandemic trajectories than California, potentially biasing the synthetic counterfactual. For this reason, I emphasize the 2023+ staggered DiD results over the California synthetic control in drawing conclusions.

6. Mechanisms and Heterogeneity

6.1 Potential Mechanisms

The positive effect of privacy laws on business formation is consistent with several mechanisms. Distinguishing between these mechanisms is challenging with the available data, as they make similar predictions about overall effects but differ in their implications for welfare and policy design.

Regulatory clarity. Comprehensive privacy statutes provide clearer rules than the patchwork of common law, sector-specific regulations, and FTC enforcement that previously governed data practices. Before comprehensive state laws, businesses faced uncertainty from multiple sources: common law torts (intrusion, misappropriation, public disclosure), sector-specific federal rules (HIPAA for health, GLBA for finance, COPPA for children’s data), state breach notification laws, and FTC enforcement actions based on “unfair or deceptive” practices standards that evolve through case-by-case adjudication.

This regulatory uncertainty imposes costs on potential entrants who must assess compliance risk without clear guidance. Comprehensive privacy laws reduce this uncertainty by providing explicit rules: what constitutes “personal information,” what rights consumers have, what disclosures are required, what exemptions apply. Entrepreneurs may prefer the certainty of known compliance costs to the uncertainty of potential future enforcement, even if the known costs are non-trivial.

Consumer trust. Privacy laws may expand demand by reassuring consumers that their data will be protected. Survey evidence consistently shows that consumers express strong concerns about data privacy: approximately 80% report being “very” or “somewhat” concerned about companies collecting their personal data, and similar majorities support stronger privacy regulations. If these stated preferences translate into purchasing behavior, privacy laws could shift demand toward firms operating under strong privacy frameworks.

This effect should be strongest for firms selling privacy-sensitive products (health, finance, children’s products) or targeting privacy-conscious consumer segments. It may also matter for business-to-business transactions, as enterprise customers increasingly require vendors to meet privacy standards as a condition of contracting. The growth of “privacy-first” marketing suggests that some firms view privacy compliance as a competitive advantage rather than

merely a cost.

Signaling. Starting a business in a jurisdiction with strong privacy laws may signal commitment to data protection, analogous to how environmental regulations can signal environmental responsibility (the “California effect” in environmental policy). This signaling mechanism differs from the consumer trust channel in that it operates through credible commitment rather than direct consumer awareness of law content.

For business-to-business firms, signaling may be particularly valuable when seeking enterprise customers with strict vendor requirements. Many large companies now conduct privacy due diligence on suppliers and may prefer vendors located in jurisdictions with strong privacy frameworks. Locating in a privacy-law state provides a low-cost signal that may substitute for expensive certification processes.

Competitive positioning. Privacy laws may shift competitive dynamics to favor compliant firms over non-compliant competitors. Incumbents with legacy data systems built around unconstrained data collection face substantial switching costs to become compliant. New entrants can design privacy-compliant systems from the start, potentially gaining a relative advantage.

This mechanism predicts that privacy laws benefit entrants at the expense of incumbents, possibly reducing total industry profits while increasing the number of competitors. The welfare implications depend on whether the new entrants expand the market (positive) or primarily capture market share from incumbents (ambiguous).

Selection and reallocation. Privacy-conscious entrepreneurs may relocate to states with strong privacy laws, shifting business formation geographically without necessarily increasing total entrepreneurship. Similarly, entrepreneurs seeking to avoid regulation may shift to never-treated states. These reallocation effects would show up as positive treatment effects in treated states and potentially negative effects in untreated states.

The aggregate welfare implications of reallocation are ambiguous. If reallocation is purely a zero-sum shift, privacy laws would not increase total entrepreneurship but might still be welfare-improving if they better match privacy-conscious entrepreneurs with privacy-conscious consumers. If reallocation involves deadweight loss (costs of relocating, inefficient location decisions), the aggregate effect could be negative even if treated states show positive effects.

6.2 Heterogeneity by Cohort

Table 5 presents treatment effects by cohort. Virginia (the earliest 2023 adopter) shows a significant positive effect of 263 applications per month. Colorado and Connecticut (July 2023) show larger but less precisely estimated effects (416, SE 236.6). Texas and Oregon (July 2024) show the largest cohort-specific effects (964), suggesting that effects may have

grown stronger over time or that later adopters differ systematically.

Note that cohort-specific ATTs are not simple averages of the overall ATT: the Callaway-Sant’Anna aggregation weights cohort-time cells by post-treatment exposure and group size, so cohorts with many post-treatment periods (like Virginia with 30 months) receive more weight than cohorts with few months (like Texas/Oregon with 12 months). The large Texas/Oregon estimate (964) reflects a short, volatile post-treatment window in two high-volume states and does not dominate the aggregate because its weight is relatively small. The overall ATT of 240 reflects the weighted average across all cohort-time cells.

The January 2025 cohort (Iowa, Delaware, etc.) shows smaller effects (32.1), but this reflects limited post-treatment data (only a few months by the end of the sample).

Montana (October 2024) shows a negative point estimate (-307), but this is not statistically significant (95% CI: $[-673, 59]$) and reflects only a few months of post-treatment data in a small state (population ≈ 1.1 million). Montana’s baseline business applications are substantially lower than other treated states (approximately 300 per month vs. 2,000+ for Texas), so small month-to-month fluctuations generate large percentage changes. The leave-one-out analysis confirms that excluding Montana produces similar overall results (ATT = 267.8, the maximum in the leave-one-out range; see Table A1), indicating that this outlier does not drive the main findings.

Table 5: Heterogeneity by Treatment Cohort

Cohort	States	N	Post Mos	ATT	SE	95% CI
January 2023	VA	1	30	263**	56.4	[152, 374]
July 2023	CO, CT	2	24	416	236.6	[-48, 880]
January 2024	UT	1	18	17.5	95.0	[-169, 204]
July 2024	TX, OR	2	12	964**	187.8	[596, 1332]
October 2024	MT	1	9	-307.1	186.7	[-673, 59]
January 2025	IA, DE, NH, NE	4	6	32.1	52.3	[-70, 135]
February 2025	NJ	1	5	18.6	61.2	[-101, 139]

Notes: Group-specific ATTs from Callaway-Sant’Anna estimation. N = number of states in cohort. Post Mos = number of post-treatment months available through June 2025. ** indicates 95% confidence band excludes zero.

6.3 Limitations

Several limitations warrant discussion. I organize these by their implications for the validity and interpretation of the findings.

Statistical precision limitations. First, the relatively small number of treated states (12) limits statistical precision and raises concerns about inference. While the Callaway-

Sant’Anna estimator is designed for staggered designs, inference with few clusters can be sensitive to outliers and may produce confidence intervals that are too narrow. I address this through leave-one-out analysis (showing that no single state drives the results) and wild cluster bootstrap inference (producing a p-value of 0.018, confirming significance). Nevertheless, readers should interpret the point estimates with appropriate uncertainty.

Second, the January 2025 cohort (Iowa, Delaware, New Hampshire, New Jersey, Nebraska) has only a few months of post-treatment data as of the end of my sample. The cohort-specific estimates for this group are imprecise and should be viewed as preliminary. As more post-treatment periods accumulate, precision will improve.

Interpretation limitations. Third, I cannot distinguish between the mechanisms discussed above—regulatory clarity, consumer trust, signaling, competitive positioning, and geographic reallocation. All may contribute to the positive effect, and their relative importance likely varies across industries, firm types, and time periods. Different mechanisms have different welfare implications: demand expansion is generally welfare-improving, while pure reallocation may involve deadweight costs. Future research with more granular data (firm-level, industry-level) could help distinguish these channels.

Fourth, the outcome (business applications) measures entry intent rather than actual business formation or economic success. Not all EIN applications result in operating businesses; Census Bureau estimates suggest approximately 20% of high-propensity applications become employer businesses within 8 quarters. If privacy laws disproportionately encourage “frivolous” applications that never become real businesses, the positive effect on applications would overstate the effect on actual entrepreneurship. However, there is no obvious reason to expect this pattern, and high-propensity applications have historically been strong predictors of future employer businesses.

Fifth, I cannot observe whether the businesses formed are genuinely new ventures or reorganizations of existing entities. A firm facing privacy compliance costs in one state might dissolve and re-form in another state, or restructure as a new legal entity. Such reorganizations would appear as new applications without representing genuinely new economic activity.

Identification limitations. Sixth, endogenous adoption remains a concern despite the strong parallel trends results. States may adopt privacy laws in response to unobserved factors that also affect business formation. The parallel trends test can only detect selection on pre-existing linear trends; it cannot rule out selection on time-varying unobservables that coincide with treatment timing. For example, if states adopt privacy laws when their technology sectors are poised for growth (for reasons unrelated to privacy), the estimated effects would conflate privacy law impacts with underlying sector trends.

Seventh, I cannot rule out spillovers across states that would violate the stable unit

treatment value assumption. If privacy laws in treated states affect business formation in untreated states (through competition for entrepreneurs, consumer shopping effects, or demonstration effects), the comparison group is also affected, biasing estimates in unknown directions.

Generalizability limitations. Eighth, the current wave of state privacy laws represents an early stage of regulatory experimentation. As more states adopt similar frameworks, the novelty advantage of early adopters may diminish, and the effects I estimate may not persist. Federal privacy legislation, if enacted, would further change the competitive landscape.

Ninth, my sample period (2018–2025) includes the COVID-19 pandemic, which caused unprecedented disruption to business formation patterns. While I attempt to control for pandemic effects through time fixed effects and by excluding the California/COVID-confounded period, residual pandemic heterogeneity across states may affect the estimates.

7. Conclusion

This paper provides causal evidence that comprehensive state privacy laws are associated with increased business formation, contrary to the conventional narrative emphasizing compliance costs. Using the staggered adoption of CCPA-style legislation across twelve U.S. states, I find that privacy laws increased high-propensity business applications by approximately 240 per month (11% relative to pre-treatment means), with strong support for the parallel trends assumption. The effect is robust to alternative estimators, leave-one-out sensitivity analysis, and wild cluster bootstrap inference.

These findings have several policy implications. First, concerns about privacy regulation deterring entrepreneurship may be overstated, at least at the state level. While compliance costs are real—surveys suggest implementation costs of \$50,000 to over \$1 million for covered firms—the benefits of regulatory clarity, consumer trust, and competitive positioning appear to offset these costs on average. This does not mean that privacy regulation is costless, but rather that the costs are not so large as to overwhelm the countervailing benefits.

Second, the positive effects suggest that federal privacy legislation—currently under consideration in Congress—need not be viewed as inherently anti-innovation. The policy debate has often framed privacy regulation as a tradeoff between consumer protection and economic dynamism. The evidence from state-level experimentation suggests this tradeoff may be less stark than assumed, or even that the two goals can be complementary. Federal legislation that provides regulatory clarity while avoiding excessive compliance burdens could potentially stimulate rather than hinder entrepreneurship.

Third, the results indicate that the “California effect” in privacy regulation may involve

not just diffusion of standards across jurisdictions, but actual economic benefits for early-adopting states. States that moved quickly to implement comprehensive privacy frameworks may have attracted privacy-conscious entrepreneurs and established first-mover advantages in privacy-oriented markets. This suggests that the race to adopt privacy legislation may be driven by more than just political factors—there may be genuine economic benefits to early adoption.

Fourth, the findings contribute to the broader literature on regulatory barriers to entrepreneurship. The standard narrative—that regulation deters entry by imposing costs—is clearly incomplete. Some regulations may facilitate entry by reducing uncertainty, enabling market transactions, or shifting competitive dynamics. Understanding which regulations help versus hurt entrepreneurship requires empirical investigation rather than reliance on theoretical priors.

Several avenues for future research emerge from this analysis. Most importantly, distinguishing between the mechanisms underlying the positive effects would help guide policy design. Do privacy laws primarily benefit entrepreneurship through regulatory clarity, consumer trust expansion, signaling, competitive repositioning, or geographic reallocation? Each mechanism has different implications for welfare and for optimal policy design. Firm-level data, industry breakdowns, and surveys of entrepreneurs could help disentangle these channels.

Second, as more states adopt similar frameworks and federal legislation potentially emerges, the effects of privacy laws may change. Early adopters may have enjoyed first-mover advantages that later adopters cannot replicate. Long-run effects may differ from the short-run impacts I estimate here as markets adjust and competitive dynamics evolve. Continued monitoring of business formation patterns as the privacy regulatory landscape matures will be valuable.

Third, examining firm-level outcomes beyond entry—employment growth, revenue, survival, and productivity—would provide a more complete picture of privacy regulation’s effects. A positive effect on entry could mask negative effects on firm growth if compliance costs burden young firms disproportionately. Alternatively, the signaling and trust mechanisms could boost firm performance conditional on entry.

Finally, international comparisons could illuminate how different approaches to privacy regulation affect economic outcomes. The European GDPR represents a more stringent, centralized approach compared to the patchwork of U.S. state laws. Comparing outcomes across these regulatory regimes could inform the design of privacy frameworks that balance consumer protection with economic dynamism.

The findings also speak to the emerging literature on “regulatory technology” or “regtech”—the idea that compliance with complex regulations can itself become a business opportunity.

The positive effect on business formation suggests that entrepreneurs may be finding ways to monetize privacy compliance, either by selling compliance services to other firms or by building privacy-compliant products that appeal to privacy-conscious consumers. This “compliance as competitive advantage” framing merits further investigation as privacy regulation expands.

More broadly, this paper joins a growing body of work challenging the assumption that regulation is inherently anti-entrepreneurial. While some regulations clearly impose net costs that deter entry, others may facilitate entrepreneurship by reducing uncertainty, enabling market transactions, or shifting competitive dynamics in favor of new entrants. The key question for policy design is not whether to regulate, but how to design regulations that capture consumer protection benefits while minimizing barriers to productive entry. The evidence from state privacy laws suggests that well-designed disclosure and consent requirements need not impede—and may even stimulate—business formation.

In conclusion, this paper provides among the first causal evidence on how comprehensive state privacy laws affect entrepreneurship. The finding that privacy regulation increases business formation by approximately 11% suggests that policymakers need not view privacy protection and economic dynamism as inherently at odds. As Congress continues to debate federal privacy legislation and more states consider comprehensive privacy frameworks, understanding the full range of economic effects—including potential benefits to entrepreneurship—is essential for informed policy design.

References

- Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, 105(490), 493-505.
- Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442-92.
- Bloomberg Law. (2024). State privacy law tracker. *Bloomberg Law*.
- Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event study designs: Robust and efficient estimation. *Review of Economic Studies*, forthcoming.
- California Department of Finance. (2019). Economic impact analysis of the California Consumer Privacy Act. Technical Report.
- Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200-230.
- Djankov, S., La Porta, R., Lopez-de-Silanes, F., & Shleifer, A. (2002). The regulation of entry. *Quarterly Journal of Economics*, 117(1), 1-37.
- Goldfarb, A., & Tucker, C. (2011). Privacy regulation and online advertising. *Management Science*, 57(1), 57-71.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254-277.
- Haltiwanger, J., Jarmin, R., Kulick, R., & Miranda, J. (2021). High growth young firms: Contribution to job, output, and productivity growth. In *Measuring Entrepreneurial Businesses* (pp. 11-62). University of Chicago Press.
- Jia, J., Jin, G. Z., & Wagman, L. (2021). The short-run effects of the General Data Protection Regulation on technology venture investment. *Marketing Science*, 40(4), 661-684.
- Johnson, G., Shriver, S., & Goldberg, S. (2024). Privacy and market concentration: Intended and unintended consequences of the GDPR. *Management Science*, forthcoming.
- Klapper, L., Laeven, L., & Rajan, R. (2006). Entry regulation as a barrier to entrepreneurship. *Journal of Financial Economics*, 82(3), 591-629.
- Sant'Anna, P. H., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1), 101-122.
- Solove, D. J., & Hartzog, W. (2021). The GDPR's three-year anniversary: The US is still falling behind. *MIT Technology Review*.
- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175-199.

Vogel, D. (1995). *Trading Up: Consumer and Environmental Regulation in a Global Economy*.
Harvard University Press.

Appendix

A. Additional Tables and Figures

Table A1: Leave-One-Out Sensitivity Analysis

Dropped State	ATT	SE	Obs	States	Treated
Virginia	232.4	71.2	4,320	48	11
Colorado	218.9	68.4	4,320	48	11
Connecticut	225.1	69.8	4,320	48	11
Utah	245.3	66.1	4,320	48	11
Texas	198.7	59.3	4,320	48	11
Oregon	203.4	61.2	4,320	48	11
Montana	267.8	68.9	4,320	48	11
Iowa	241.2	65.8	4,320	48	11
Delaware	238.9	65.4	4,320	48	11
New Hampshire	240.1	65.7	4,320	48	11
New Jersey	239.4	65.5	4,320	48	11
Nebraska	240.8	65.6	4,320	48	11

Notes: Callaway-Sant’Anna estimates dropping each treated state one at a time. Obs = state-month observations. States = total states in sample. Treated = number of treated states remaining after dropping one.

B. Data Sources

- **Business Formation Statistics:** Census Bureau Business Formation Statistics (BFS) accessed via FRED API. Series naming convention: HPBATOTALSAXX where XX is the two-letter state postal code. For example: HPBATOTALSAVA (Virginia), HPBATOTALSATX (Texas), HPBATOTALSANY (New York). The “HP” prefix indicates high-propensity business applications; “SA” indicates seasonally adjusted. Data coverage: January 2004 to present, monthly frequency, all 50 states plus DC.
- **Privacy Law Dates:** Hand-collected from Bloomberg Law State Privacy Law Tracker, National Conference of State Legislatures (NCSL) legislation database, and International Association of Privacy Professionals (IAPP) US State Privacy Legislation Tracker. Effective dates cross-referenced across all three sources.
- **State Unemployment:** Bureau of Labor Statistics, Local Area Unemployment Statistics (LAUS) program. Series: state-level unemployment rates, seasonally adjusted, monthly.

C. Replication Code

All code and data are available at:

https://github.com/anthropics/auto-policy-evals/tree/main/papers/a pep_0112

Acknowledgements

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

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>