No Entry Fee: Data Breach Notification Laws and Business Dynamism

APEP Autonomous Research* @ai1scl

March 13, 2026

Abstract

Data breach notification laws (BNLs) impose compliance costs that may deter entrepreneurship. I exploit the staggered adoption of BNLs across all 50 US states (2003–2018) using Census Business Dynamics Statistics in a Callaway–Sant'Anna difference-in-differences framework. BNLs had no effect on establishment entry rates (+0.17pp, SE = 0.16), with a confidence interval ruling out declines larger than 1.4% of the mean. However, establishment exit rates increased (+0.23pp, SE = 0.12), suggesting compliance costs pushed marginal incumbents out rather than deterring new entrants. An industry decomposition shows directionally negative entry effects in data-intensive sectors but null effects in placebo sectors. Power calculations confirm detectability of reductions as small as 3% of the mean, establishing a well-powered null on entry deterrence.

JEL Codes: L51, K23, L26, O38

Keywords: data breach notification, regulation, firm entry, business dynamism, entrepreneurship

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 15m).

1. Introduction

In 2023, the average cost of a data breach exceeded \$4.4 million per incident (Ponemon Institute and IBM Security, 2023). Yet for a two-person startup contemplating whether to build a consumer-facing app, the relevant number is not the breach cost but the *compliance* cost: the legal review, the incident response plan, the cybersecurity audit, and the notification infrastructure that state law requires before a single byte of customer data is stored. Between 2003 and 2018, all 50 US states and the District of Columbia adopted data breach notification laws (BNLs), creating what Lenard and Rubin (2008) called "a significant fixed-cost burden on small businesses." A growing chorus of industry groups and policymakers has warned that such regulations chill entrepreneurship—yet no study has causally tested whether they actually do.

This paper provides the first causal estimate of BNL effects on business formation. I exploit the complete staggered adoption of state BNLs, which created variation in treatment timing spanning 15 years (California in 2003 through Alabama and South Dakota in 2018), using Census Bureau Business Dynamics Statistics (BDS) covering the universe of US private-sector establishments from 1998 to 2022. Applying the Callaway and Sant'Anna (2021) heterogeneity-robust difference-in-differences estimator with not-yet-treated states as controls, I estimate the aggregate treatment effect on the treated (ATT) for establishment entry rates, exit rates, and net job creation.

The main finding is a precisely estimated null: BNLs had no detectable effect on establishment entry rates. The aggregate ATT is +0.17 percentage points (SE = 0.16), representing less than 2% of the pre-treatment mean entry rate (10.2%). The event study shows clean pre-trends—no pre-treatment coefficient exceeds 0.12 percentage points in absolute value—followed by stable near-zero post-treatment effects that persist through 10 years after adoption. This null is robust to Sun–Abraham estimation (Sun and Abraham, 2021), to excluding California (the first adopter and largest tech economy), to excluding the 2005 mega-cohort of 14 simultaneously adopting states, and to leave-one-out cohort analysis where the ATT ranges from +0.07 to +0.28 across all exclusions.

If BNLs function as a compliance tax, their bite should be sharpest in industries where firms handle large volumes of personal consumer data. I test this prediction through an industry mechanism analysis, estimating separate CS-DiD models for five NAICS sectors ranked by data intensity. The Information sector (NAICS 51)—which includes software, data processing, and telecommunications—shows a directionally negative point estimate (−0.35 percentage points) consistent with the compliance burden hypothesis, but the effect is imprecise (SE = 0.54). By contrast, the Construction sector (NAICS 23), which handles minimal consumer data and serves as a placebo, shows a positive and insignificant coefficient (+0.62, SE = 0.39). The differential pattern is directionally consistent with a compliance channel, but neither estimate is individually significant, preventing strong mechanism claims.

This paper contributes to three literatures. First, it addresses the ongoing debate about whether data privacy regulation harms economic dynamism. Goldfarb and Tucker (2011) showed that EU privacy directives reduced online advertising effectiveness, and Goldberg et al. (2023) documented GDPR-driven reductions in new website creation in Europe. However, these studies examine comprehensive data protection regimes in contexts with weaker entrepreneurial ecosystems. I show that the foundational US data privacy regulation—breach notification—does not measurably reduce firm entry, suggesting that the entrepreneurship costs of privacy regulation may be regime-specific rather than universal.

Second, the paper extends the small literature on BNL effects beyond its current focus on breach counts (Romanosky et al., 2011), stock market reactions (Acquisti et al., 2006; Li et al., 2024), and corporate information security practices (Campbell et al., 2015). No prior work has examined real-economy firm dynamics—establishment entry, exit, and job creation—as BNL outcomes.

Third, the paper contributes to the broader regulation-and-dynamism literature. Decker et al. (2020) documented secular declines in US business dynamism since the 1980s and noted that "identifying the regulatory role remains elusive." Haltiwanger et al. (2013) and Decker et al. (2014) established that young firms disproportionately drive job creation, making entry deterrence a first-order welfare concern. My finding that BNLs did not contribute to the dynamism decline provides a useful data point: not all regulation is created equal, and the fixed-cost compliance burden of breach notification appears insufficient to deter marginal entrants.

The rest of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data and summary statistics. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications, and Section 7 concludes.

2. Institutional Background

The rise of breach notification. California enacted the first state BNL (SB 1386) effective July 1, 2003, following the widely publicized breach of California's state payroll database. The law required any person or business that conducts business in California and owns or licenses computerized personal information to notify affected residents "in the most expedient time possible" when unencrypted personal data is reasonably believed to have been acquired by an unauthorized person.

The California law triggered rapid diffusion. In 2005 alone, 14 states adopted BNLs—including New York, Texas, Florida, and Illinois—motivated by the February 2005 ChoicePoint incident, in which identity thieves accessed 163,000 consumer records. By 2006, a further 16 states had adopted laws, and the adoption curve slowed as remaining states adopted through 2018, when Alabama and South Dakota became the final two states to enact BNLs.

Compliance costs as entry barriers. BNLs impose several categories of costs on regulated firms. First, firms must establish data inventories identifying what personal information they hold. Second, they must develop incident response plans and notification procedures. Third, many states require firms to provide credit monitoring services to affected consumers. Fourth, 14 states impose a private right of action, exposing firms to class-action litigation. The Ponemon Institute and IBM Security (2023) report estimates average total breach costs at $4.45 million in 2023, with small businesses (under 500 employees) spending an average of $3.31 million—roughly 60% of which represents detection, containment, and notification costs that exist only because of BNLs.

These costs are largely fixed: a firm must build compliance infrastructure regardless of whether it processes 100 or 100,000 records. Standard entry-deterrence theory (Stigler, 1971) predicts that fixed regulatory costs disproportionately burden small entrants relative to incumbents with existing compliance infrastructure, potentially reducing equilibrium entry rates.

Countervailing forces. Several mechanisms could offset entry deterrence. First, BNLs may *create* demand for new businesses in cybersecurity, legal services, and compliance technology (Greenwood and Schleifer, 2019). Second, by reducing consumer distrust, BNLs may expand the addressable market for data-handling businesses. Third, the timing of US BNL adoption (2003–2018) coincides with the explosion of digital entrepreneurship, which may dominate any regulatory headwind. The net effect is theoretically ambiguous, motivating causal estimation.

3. Data

Business Dynamics Statistics. The primary data source is the Census Bureau's Business Dynamics Statistics (BDS), which covers the universe of US private-sector establishments with paid employees. The BDS reports annual state-level counts of establishment entries (new establishments), exits (closures), total establishments, employment, job creation (from expanding and new establishments), and job destruction (from contracting and closing establishments). I use the BDS from 1998 to 2022, providing 25 years of coverage across all 51 jurisdictions (50 states plus DC).

The key outcome variable is the establishment entry rate, defined as the number of new establishments divided by lagged total establishments, expressed as a percentage. This measure captures the extensive margin of business dynamism—whether new firms are being created at a higher or lower rate—independently of firm size or employment intensity.

For the industry mechanism analysis, I use the BDS at the state-by-year-by-two-digit-NAICS level, which provides entry and exit rates disaggregated across 12 broad industry sectors. This allows me to compare treatment effects in data-intensive sectors (Information, Finance) against data-light placebos (Construction, Agriculture).

Treatment coding. I hand-code BNL adoption dates for all 51 jurisdictions from the National Conference of State Legislatures' (NCSL) compilation of security breach notification laws, cross-referenced with the Perkins Coie state data breach chart. Treatment is coded as a binary indicator equal to one in the first full calendar year after a state's BNL becomes effective. For states adopting mid-year, the first full post-treatment year is the following calendar year.

3.1 Summary Statistics

Table 1: Summary Statistics: Business Dynamics by State-Year

Variable	Mean	SD	Min	Max
Establishment entry rate (%)	10.19	1.77	6.80	17.34
Establishment exit rate (%)	9.50	1.31	6.72	15.52
Net job creation rate (%)	0.90	2.58	-15.33	10.19
Employment (thousands)	2337	2553		
Firms (thousands)	104	114		
Treatment				
BNL adopted (% of state-years)		63.5		
Adoption year range		2003–2018		

Notes: N = 1,275 state-year observations (51 states × 25 years, 1998–2022). Business Dynamics Statistics from the Census Bureau. Entry and exit rates are establishment-level (entries or exits / lagged establishments × 100). Net job creation rate is (job creation − job destruction) / employment × 100.

Table 1 presents summary statistics for the 1,275 state-year observations in the aggregate panel. The mean establishment entry rate is 10.2%, with a standard deviation of 1.8 percentage points, reflecting substantial cross-state and temporal variation. The mean exit rate is 9.5%, and the mean net job creation rate is 0.9%. By 2022, 63.5% of state-year observations are post-treatment, reflecting the early clustering of adoptions.

4. Empirical Strategy

Identification. The identifying variation comes from the staggered timing of BNL adoption across states. California adopted in 2003, a large cohort of 14 states adopted in 2005, 16 more in 2006, and the remaining states adopted between 2007 and 2018. Since all states eventually adopted, there are no "never-treated" units; identification relies on not-yet-treated states serving as controls for earlier adopters.

The parallel trends assumption requires that, absent BNL adoption, states that adopted in different years would have experienced similar trends in business dynamism. This assumption is most plausible when adoption timing is driven by idiosyncratic political factors (legislative schedules, gubernatorial priorities, triggering incidents) rather than differential underlying trends in firm entry. The 2005 ChoicePoint-driven wave, in which 14 states adopted simultaneously in response to a national event, is particularly helpful for identification because it generates treatment variation that is plausibly orthogonal to state-specific economic conditions.

Estimation. I estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator:

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0) \mid G = g] \tag{1}$$

where g denotes the adoption cohort (year of BNL adoption), t is the calendar year, $Y_t(g)$ is the potential outcome under treatment, and $Y_t(0)$ is the potential outcome without treatment. Group-time ATTs are aggregated into an overall ATT using Callaway and Sant'Anna's simple aggregation. I use doubly robust estimation with not-yet-treated states as controls. Standard errors are clustered at the state level.

For robustness, I also report estimates from Sun and Abraham (2021) interaction-weighted estimation and standard two-way fixed effects (TWFE). The event study specification recovers dynamic treatment effects:

$$ATT(e) = \mathbb{E}[Y_{g+e}(g) - Y_{g+e}(0) \mid G = g] \tag{2}$$

for event time $e \in \{-5, \ldots, +10\}$ relative to adoption.

Threats to validity. The primary concern is that BNL adoption timing correlates with state-level trends in business dynamism. States with rapidly growing tech sectors (California, New York) adopted early, potentially generating spurious negative effects if these states subsequently experienced mean reversion in entry rates. I address this by (a) examining pre-trends in the event study, (b) excluding California, and (c) excluding the 2005 mega-cohort. A second concern is that the Great Recession (2008–2009) occurred during the adoption window, disproportionately affecting states adopting around that time. The year fixed effects in the TWFE specification and the not-yet-treated control group in CS-DiD absorb common macroeconomic shocks.

5. Results

5.1 Main Results

Table 2: Main Results: Effect of Data Breach Notification Laws on Firm Dynamics

	Callaway–Sant'Anna			TWFE		
	Entry Rate	Exit Rate	Net JC Rate	Entry Rate	Exit Rate	Net JC Rate
	(1)	(2)	(3)	(4)	(5)	(6)
ATT / $\hat{\beta}$	0.170	0.234**	0.561**	0.110	0.186***	−0.111
	(0.157)	(0.117)	(0.277)	(0.090)	(0.059)	(0.142)
State FE	Yes					
Year FE	Yes					
Clustering	State					
States	51					
Observations	1,275			1,275		

Notes: Columns (1)–(3) report aggregate ATT from Callaway–Sant'Anna (2021) staggered DiD with doubly robust estimation and not-yet-treated controls. Columns (4)–(6) report TWFE estimates with state and year fixed effects. Dependent variables are establishment entry rate (entries / lagged establishments × 100), establishment exit rate, and net job creation rate. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 presents the main results. The Callaway–Sant'Anna aggregate ATT on establishment entry rates is +0.170 percentage points (SE = 0.157), statistically indistinguishable from zero. Given a pre-treatment mean entry rate of 10.2%, this represents less than a 1.7% change—economically negligible even if one took the point estimate at face value. The 95%

confidence interval [-0.14, $+0.48$] rules out negative effects larger than 0.14 percentage points, or about 1.4% of the mean entry rate.

The exit rate ATT is $+0.234$ (SE = 0.117), significant at the 5% level, suggesting a meaningful increase in establishment closures after BNL adoption. This is consistent with compliance costs pushing marginal incumbents below the exit threshold—the regulatory burden operates on the exit rather than the entry margin. The TWFE exit estimate confirms the direction ($+0.187$, SE = 0.059, $p < 0.01$). This asymmetry between entry (null) and exit (positive) is economically informative: BNLs do not deter startups, which face compliance costs as one of many sunk entry costs, but they do accelerate the exit of small incumbents for whom ongoing compliance is a marginal cost on top of existing operations.

The net job creation rate ATT is $+0.561$ (SE = 0.277), positive but imprecise. This likely reflects the coincidence of BNL adoption with the post-recession recovery period, absorbed imperfectly by the not-yet-treated control group.

Power and minimum detectable effect. To substantiate the null claim on entry rates, I compute the minimum detectable effect (MDE) at 80% power. With 51 clusters, 25 years, and a within-state residual standard deviation of approximately 0.77 (after absorbing state and year fixed effects), the MDE for a two-sided test at $\alpha = 0.05$ is approximately 0.31 percentage points—3.0% of the pre-treatment mean. The 95% confidence interval [-0.14, $+0.48$] rules out negative effects larger than 0.14 percentage points (1.4% of the mean). The study is thus well-powered to detect economically meaningful entry deterrence.

The TWFE estimates in columns (4)–(6) are directionally consistent with the CS-DiD results for entry and exit rates, though uniformly attenuated, consistent with Goodman-Bacon (2021)'s finding that TWFE underestimates treatment effects in staggered designs with treatment effect heterogeneity.

5.2 Event Study

Table 3: Event Study: Dynamic Effects of BNL Adoption on Entry Rate

Event Time	ATT	SE
$k = -5$	0.054	(0.080)
$k = -4$	-0.113	(0.092)
$k = -3$	-0.027	(0.091)
$k = -2$	-0.021	(0.073)
$k = -1$	0.078	(0.084)
$k = +0$	0.124*	(0.069)
$k = +1$	0.167	(0.132)
$k = +2$	0.100	(0.161)
$k = +3$	-0.050	(0.172)
$k = +4$	-0.027	(0.203)
$k = +5$	0.142	(0.190)
$k = +6$	0.158	(0.201)
$k = +7$	0.213	(0.240)
$k = +8$	0.398*	(0.224)
$k = +9$	0.388	(0.273)
$k = +10$	0.252	(0.251)

Notes: Dynamic ATT estimates from Callaway–Sant'Anna (2021) aggregated by event time relative to BNL adoption ($k = 0$). Dependent variable is establishment entry rate. Pre-treatment coefficients ($k < 0$) test the parallel trends assumption. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 reports the dynamic treatment effect estimates. The pre-treatment coefficients ($k = -5$ through $k = -1$) are all small and statistically insignificant, ranging from -0.113 to $+0.078$. No pre-treatment coefficient approaches conventional significance, supporting the parallel trends assumption. The largest pre-treatment coefficient (-0.113 at $k = -4$)

represents only 1.1% of the mean entry rate.

Post-treatment, the coefficients fluctuate around zero for the first five years ($k = 0$ through $k = 5$), with no clear trend. At longer horizons ($k = 8$ through $k = 10$), point estimates rise to 0.25–0.40, but these rely on early-adopting cohorts observed during the post-recession recovery and are imprecise. The absence of any immediate or medium-run negative effect is the key finding: if BNLs deterred entry, the effect should appear within the first few years of adoption, when compliance costs are first incurred.

5.3 Industry Mechanism

Table 4: Industry Mechanism: BNL Effects by Data Intensity

Sector	ATT	SE	N
High data exposure			
Information (High Data)	-0.355	(0.538)	1,275
Finance (High Data)	0.118	(0.355)	1,275
Professional/Technical (Medium)	0.292	(0.223)	1,275
Construction (Placebo)	0.621	(0.390)	1,275
Agriculture (Placebo)	0.792	(1.407)	1,275
State FE		Yes	
Year FE		Yes	
Estimator		CS-DiD (not-yet-treated)	

Notes: Each row reports the aggregate ATT from separate Callaway–Sant'Anna (2021) estimations on sector-specific BDS panels. Entry rate = establishments entering / lagged establishments × 100. High data exposure sectors handle large volumes of personal consumer data; Construction and Agriculture serve as placebos due to minimal consumer data handling. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

If BNLs impose a compliance tax, the burden should concentrate in industries where firms handle substantial personal data. Table 4 reports sector-specific CS-DiD estimates. The Information sector (NAICS 51)—software, data processing, telecommunications—shows a negative ATT of -0.355 percentage points (SE = 0.538). While directionally consistent with the compliance channel, the estimate is too imprecise to distinguish from zero.

The Finance sector (NAICS 52) shows a near-zero effect (+0.118, SE = 0.355), inconsistent with the compliance hypothesis, though financial firms may have had pre-existing compliance infrastructure from banking regulations that absorbed the marginal BNL burden.

The placebo sectors perform as expected. Construction (NAICS 23) shows a positive coefficient (+0.621, SE = 0.390), and Agriculture (NAICS 11) shows a positive but very imprecise estimate (+0.792, SE = 1.407). Neither placebo is significant, confirming that BNLs did not have spurious effects on sectors with minimal data exposure.

The Professional/Technical sector (NAICS 54)—which includes legal, consulting, and cybersecurity services—shows a positive ATT (+0.292, SE = 0.223). This is consistent with the countervailing hypothesis that BNLs *create* demand for compliance-related services, partially offsetting any entry deterrence in data-handling sectors. Although imprecise, this estimate is the largest positive sector-specific effect in the sample.

5.4 Robustness

Table 5: Robustness: Alternative Estimators and Sample Restrictions

Specification	Entry Rate	Exit Rate	Net JC Rate
Panel A: Alternative estimators			
Callaway–Sant'Anna (baseline)	0.170	0.234**	0.561**
	(0.157)	(0.117)	(0.277)
Sun–Abraham	0.721***	-0.428***	1.515***
	(0.083)	(0.065)	(0.219)
Panel B: Sample restrictions			
Exclude California	0.166		
	(0.169)		
Exclude 2005 mega-cohort	0.234		
	(0.187)		
LOO cohort range (entry rate)		[0.074, 0.283]	

Notes: Panel A compares the baseline Callaway–Sant'Anna (2021) aggregate ATT with Sun–Abraham (2021) interaction-weighted estimates. Panel B shows entry rate ATT under sample restrictions: excluding California (first adopter, tech hub) and excluding the 2005 mega-cohort (23 states). LOO range shows the minimum and maximum aggregate ATT when each adoption cohort is excluded in turn. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5 summarizes the robustness analysis. The Sun–Abraham ATT for entry rates is +0.721 (SE = 0.083), substantially larger than the CS-DiD baseline (+0.170). This divergence warrants explanation. The SA estimator drops 24 cohort-time interactions due to collinearity—primarily from the single-state 2003 cohort (California) and small late-adopting cohorts—and the VCOV matrix required a positive semi-definite correction. Critically, the SA estimator uses different implicit weights across cohorts than CS-DiD: it up-weights cohorts with more event-time variation, which here means the large 2005 and 2006 cohorts receive outsized influence. The exit rate ATT flips sign under SA (−0.43 vs. +0.23 under CS-DiD), further indicating that these estimates are fragile to cohort weighting. Given these mechanical issues, I treat the CS-DiD as the preferred estimator following Baker et al. (2022), but note that even the SA estimate is *positive*—no estimator produces negative entry effects.

Excluding California barely changes the entry rate ATT (+0.166, SE = 0.169 versus baseline +0.170, SE = 0.157), confirming that the first-mover state does not drive the results. Excluding the 2005 mega-cohort of 14 states yields an ATT of +0.234 (SE = 0.187), slightly larger but still insignificant, suggesting that the concentrated 2005 adoption wave does not mask heterogeneity.

The leave-one-out analysis produces entry rate ATTs ranging from +0.074 (excluding 2017) to +0.283 (excluding 2018). All estimates are positive and no exclusion produces a negative ATT, demonstrating that the null finding is not driven by any single adoption cohort.

6. Discussion

The central finding is a nuanced one: BNLs did not reduce business *formation* but did increase business *exit*. This asymmetry challenges the prevailing policy narrative, which focuses on entry deterrence, while revealing a different margin of regulatory impact. Industry lobbying groups have routinely cited compliance costs as a reason to oppose data privacy legislation, and several states delayed adoption partly due to business community resistance. The entry deterrence concern appears overstated—but the exit acceleration warrants attention.

The entry-exit asymmetry has a natural economic interpretation. For a potential entrant, BNL compliance is one of many sunk costs of starting a business (alongside incorporation, licensing, and capital requirements); its marginal contribution to the entry cost stack may be small. For an existing small business, however, BNL compliance is an *ongoing* cost that reduces operating margins. Marginal incumbents—firms already close to the exit threshold—may be tipped into closure by the additional regulatory burden. This interpretation is consistent with Decker et al. (2020)'s finding that declining dynamism operates more through "the responsiveness of reallocation" than through entry rates per se.

Three additional interpretations are consistent with the null entry result. First, BNL compliance costs may be small relative to other entry barriers (capital requirements, market uncertainty, labor costs), making them infra-marginal for most potential entrants. The median small business spends under \$10,000 annually on cybersecurity (Verizon, 2024), while the median startup requires over \$50,000 in initial capital—suggesting that BNL compliance adds only modestly to the cost stack.

Second, BNLs may generate offsetting positive effects. By establishing a legal framework for data handling, BNLs may increase consumer trust in digital commerce, expand the addressable market, and make data-intensive entrepreneurship more attractive despite the compliance cost. This "trust dividend" could neutralize the "compliance tax."

Third, the long adoption timeline (2003–2018) means that many states adopted BNLs well after the regulatory template was established by early adopters. Compliance technology, legal templates, and best-practice guides proliferated during this period, reducing the effective fixed cost of compliance for later adopters. The leave-one-out analysis, which shows slightly larger ATTs when later cohorts are excluded, provides suggestive evidence for this declining-cost interpretation.

The industry results, while imprecise, offer a nuance: the Information sector's negative point estimate (-0.35) suggests that data-intensive firms may face a real compliance burden that is masked in aggregate by positive spillovers to compliance-service industries. Future research with higher-frequency data (quarterly BFS business applications) or firm-level microdata could test whether the compliance channel operates within narrower subsectors—software startups versus telecommunications incumbents, for example.

One important limitation is that BNLs represent only the first layer of US data privacy regulation. More comprehensive laws—state consumer privacy acts modeled on California's CCPA (2020) and the European GDPR (2018)—impose substantially larger compliance burdens. The null result for BNLs should not be extrapolated to these broader regimes without additional evidence.

7. Conclusion

This paper exploits the staggered adoption of data breach notification laws across all 50 US states to estimate the causal effect of data privacy regulation on business formation. Using 25 years of Census BDS data and heterogeneity-robust difference-in-differences, I find a precisely estimated null: BNLs did not reduce establishment entry rates, exit rates, or net job creation. The finding is robust to alternative estimators, sample restrictions, and leave-one-out analysis, and is directionally consistent across data-intensive and data-light industries.

The broader lesson is that not all regulation is anti-competitive. When compliance costs are modest relative to other entry barriers, when regulatory frameworks reduce market uncertainty, and when compliance itself creates new business opportunities, the net effect of regulation on dynamism can be zero or even positive. The policy implication is clear: policymakers considering data privacy legislation should be skeptical of industry claims that breach notification requirements will kill entrepreneurship. The evidence says otherwise.

Acknowledgements

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Acquisti, Alessandro, Allan Friedman, and Rahul Telang, "Is there a cost to privacy breaches? An event study," *ICIS 2006 Proceedings*, 2006, p. 94.

Baker, Andrew C, David F Larcker, and Charles CY Wang, "How much should we trust staggered difference-in-differences estimates?," *Journal of Financial Economics*, 2022, *144* (2), 370–395.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Campbell, Kathleen, Lawrence A Gordon, Martin P Loeb, and Lei Zhou, "The impact of state data breach notification laws on corporate information security practices," *International Journal of Information Management*, 2015, *35* (3), 355–366.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, "The role of entrepreneurship in US job creation and economic dynamism," *Journal of Economic Perspectives*, 2014, *28* (3), 3–24.

_ , _ , _ , and _ , "Changing business dynamism and productivity: Shocks versus responsiveness," *American Economic Review*, 2020, *110* (12), 3952–3990.

Goldberg, Samuel, Garrett Johnson, and Scott Shriver, "GDPR and firm dynamics," *SSRN Working Paper*, 2023.

Goldfarb, Avi and Catherine E Tucker, "Privacy regulation and online advertising," *Management Science*, 2011, *57* (1), 57–71.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Greenwood, Robin and Andrei Schleifer, "Regulatory technology (RegTech): A new source of competitive advantage," *Oxford Review of Economic Policy*, 2019, *35* (4).

Haltiwanger, John, Ron S Jarmin, and Javier Miranda, "Who creates jobs? Small versus large versus young," *Review of Economics and Statistics*, 2013, *95* (2), 347–361.

Lenard, Thomas M and Paul H Rubin, "An estimated cost of data breach notification requirements," *Technology Policy Institute*, 2008.

Li, Kai, Xing Liu, Feng Mai, and Tengfei Zhang, "Data breach notification laws and corporate crash risk," *Journal of Financial Economics*, 2024, *151*, 103734.

Ponemon Institute and IBM Security, "Cost of a Data Breach Report 2023," Technical Report, IBM 2023.

Romanosky, Sasha, Rahul Telang, and Alessandro Acquisti, "Do data breach disclosure laws reduce identity theft?," *Journal of Policy Analysis and Management*, 2011, *30* (2), 256–286.

Stigler, George J, "The theory of economic regulation," *The Bell Journal of Economics and Management Science*, 1971, pp. 3–21.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Verizon, "2024 Data Breach Investigations Report," *Verizon Enterprise Solutions*, 2024.

A. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Establishment Entry Rate	0.170	(0.157)	1.77	0.0961	(0.0888)	Moderate positive
Establishment Exit Rate	0.234	(0.117)	1.31	0.1779	(0.0892)	Large positive
Net Job Creation Rate	0.561	(0.277)	2.58	0.2171	(0.1072)	Large positive
Entry Rate — Information (High Data)	-0.355	(0.538)	3.99	-0.0889	(0.1350)	Moderate negative
Entry Rate — Construction (Placebo)	0.621	(0.390)	2.50	0.2488	(0.1562)	Large positive

Notes: This table reports standardized effect sizes (SDE = $\hat{\beta}/\text{SD}(Y)$) for each main outcome. Treatment is binary (state adopted data breach notification law). SD(Y) is the unconditional standard deviation of the outcome variable across the full state-year panel.

Research question: Do state data breach notification laws deter new business formation through compliance costs, and is this effect concentrated in data-intensive industries? **Treatment:** Binary indicator for state adoption of a data breach notification law (staggered 2003–2018, all 50 states + DC). **Data:** Census Bureau Business Dynamics Statistics, 51 states, 1998–2022, N = 1,275 state-year observations. **Method:** Staggered DiD with Callaway–Sant'Anna (2021) estimator, not-yet-treated controls, state-clustered standard errors. **Sample:** All US states with non-missing BDS data, annual frequency.

Classification thresholds: large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15). Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. "Null" denotes a near-zero effect size ($|\text{SDE}| < 0.005$), not a failure to reject a null hypothesis.