Fifty years of nearly passing a privacy law

American privacy law was built one sector at a time, for a reason. This is the history of that choice, the bills that tried to reverse it, and why the closest attempt still failed.

> The United States does not lack a federal privacy law by accident. It has repeatedly chosen sectoral rules, and the cost of that choice now falls on data the sectors never anticipated.

HOW IT GOT HERE

1970 **Fair Credit Reporting Act**
The first modern data-rights statute: access to your own credit file, and a route to dispute what is in it. Sectoral by design.

1974 **Privacy Act**
Binds federal agencies rather than companies, following the Watergate-era concern about government files. The private sector is left out.

1986 **ECPA and CFAA**
Electronic communications get statutory protection, on a model of stored versus in-transit that predates cloud storage and ages badly.

1996 **HIPAA**
Health privacy attaches to covered entities - providers, plans, clearinghouses - not to health data itself. The distinction becomes load-bearing.

1998 **COPPA**
Children under 13 get a consent regime. The age line and the actual-knowledge standard shape platform design for two decades.

2001 **USA PATRIOT Act**
Post-9/11 expansion of collection authority, running against the direction of the preceding thirty years.

2018 **CCPA**
California legislates where Congress has not, and the practical effect is a national standard set by one state.

2022 **ADPPA**
Reported out of committee 53 to 2, the furthest a comprehensive federal bill has advanced. It never receives floor time.

SECTORAL BY CHOICE, NOT BY ACCIDENT

American privacy law is usually described as fragmented, as though the fragments were an oversight. They were not. Each statute was written to answer a specific scandal in a specific industry: credit files that people could not see, government dossiers assembled without limit, health records moving between insurers, children targeted by early web advertising. In each case Congress legislated the sector and stopped.

The approach has a real advantage, which is why it persisted. A rule written for credit reporting can be specific about credit reporting. It can name the actors, define the record, and set a dispute process that fits how the industry actually works. Comprehensive law has to be abstract enough to cover everything, and abstraction is where enforcement goes to die.

The disadvantage only became structural later. Sectoral rules attach to actors, not to data. Health privacy law binds providers and their business associates. It does not bind an app that infers a pregnancy from a shopping pattern, because that app is not a provider - and the inference is not, legally, a health record.

THE GAP THE INFERENCE OPENED

Machine learning made the actor-based structure much harder to defend, because it made sensitive categories derivable from insensitive inputs. Location traces yield a clinic visit. Purchase histories yield a diagnosis. Typing cadence yields a neurological condition. None of the inputs are protected, and the output is not held by a covered entity.

This is the specific failure the last decade of proposals has been trying to answer, and it explains why the serious bills abandoned consent as the organising principle. Consent regulates collection. Inference is not collection: nothing was collected that the person could have refused, because the sensitive fact was never collected at all. It was computed.

Data minimisation is the answer the strongest proposals converged on. Rather than asking whether the person agreed, ask whether the processing was necessary for the thing they actually requested. It moves the burden onto the company and away from a consent dialogue nobody reads, and it reaches inferences, because an inference either was or was not necessary to deliver the service.

WHY THE CLOSEST ATTEMPT STILL FAILED

The American Data Privacy and Protection Act came nearer than anything before it: reported out of the House Energy and Commerce Committee 53 to 2, with a genuinely bipartisan structure built on minimisation, civil-rights protection in data use, and algorithmic impact assessment.

It did not fail on the substance of those provisions. It failed on two questions that are about federalism rather than privacy. The first is pre-emption: a federal standard that displaces state law is worth less than the state law it displaces if it sets a lower bar, and California had by then set a higher one. The second is the private right of action: whether individuals may sue, which determines whether the statute is enforced at the scale of the harm or at the scale of an agency’s budget.

Both are distributive fights about who holds power afterwards, not disagreements about whether people should have privacy. That is worth stating plainly, because it predicts the shape of the next attempt. A comprehensive bill will pass when a pre-emption formula is found that a state with strong law can accept, and not before.

WHAT THE STATES DID IN THE MEANTIME

In the absence of a federal standard, one state's law became the operative national one. Compliance is not economically separable by state for most services, so a company building to the strictest applicable regime builds to it everywhere. The result is a national privacy standard set through a process no national electorate participated in.

That is an argument for federal legislation rather than against state legislation, and it is the strongest one available. It is also why the pre-emption fight is so hard: the states that legislated first have the most to lose from a federal floor set below where they already are.

The more recent state frontier-AI statutes suggest the same pattern repeating one level up. Obligations on model developers are being written in individual states, on thresholds those states set, with the practical effect of national rules. It is the privacy story again, with less time to notice it happening.

WHERE WE THINK THE WORK IS

Minimisation over consent. Consent is the mechanism that failed, and it failed in a way that cannot be repaired by better dialogs. Necessity is testable after the fact; agreement is not.

Coverage that attaches to data rather than to actors, so that health-adjacent inference is governed like health data regardless of who holds it.

A pre-emption formula that sets a floor rather than a ceiling. This is the actual blocker, and treating it as a drafting detail is why three Congresses have not cleared it.

Enforcement with a private route. Agency-only enforcement scales to an agency's capacity, which is not the scale at which this harm occurs.

SOURCES

01 American Data Privacy and Protection Act, H.R.8152 - U.S. Congress, 117th

02 Fair Credit Reporting Act - U.S. Congress, 1970

03 Privacy Act of 1974 - U.S. Congress

04 Health Insurance Portability and Accountability Act - U.S. Congress, 1996

05 California Consumer Privacy Act - State of California