PRIVACY IN AMERICA

What Every American Needs to Know

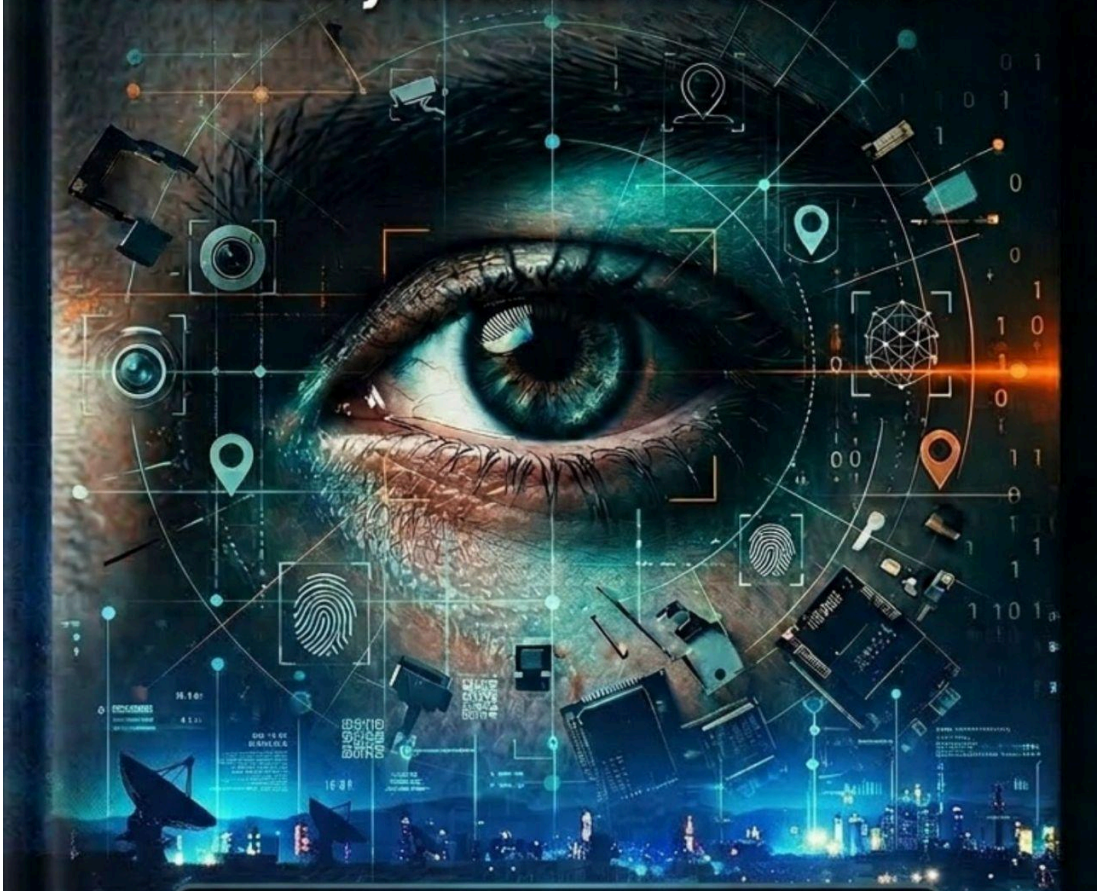

Mitch Jackson, Esq.

Introduction: They Already Have Everything

> *Here's the thing. Every day you do not understand how your privacy is being stripped away, you and your family are exposed to risks you cannot even see yet. That is exactly why I tore down the paywall on new book, "Privacy in America" and made it free and available to you 24/7. This information is too important to sit on a shelf collecting dust while the laws and technology keep shifting under your feet. This is not a static book. It is a living, real-time resource I update the moment new cases and laws drop, so you are never caught off guard again.*

In January 2025, a ransomware gang broke into the servers of a company called Conduent. Most of the 25 million Americans whose lives got turned upside down by this attack had never heard of Conduent. The company is a government contractor. Conduent processes Medicaid claims, food assistance payments, unemployment benefits, and child support disbursements for state agencies across the country. You never signed up for Conduent. You never agreed to their terms of service. Conduent collected your information because a government agency you trusted handed your information over.

The attackers spent nearly three months inside the system, from October 2024 through January 2025. They walked out with more than eight terabytes of stolen files containing names, home addresses, dates of birth, Social Security numbers, health insurance details, and medical records. By February 2026, state filings confirmed more than 25 million victims, making this one of the largest thefts of personal data from a single company in American history.

The notification letters arrived this year. Cold, form letters from a company most recipients had never dealt with, offering credit monitoring and a phone number. No explanation for why a private company held the personal information of tens of millions of Americans with security so poor a criminal gang camped inside the system for months before anyone noticed.

How did we get to a point where 25 million Americans lose their most intimate data to a company they never chose, never contacted, and never knew existed? This book answers the question.

Why I Wrote This Book

My name is Mitch Jackson, and I have spent decades as a trial lawyer in California, representing people who had their rights violated. For most of my career, those violations happened in courtrooms, in contract disputes, in personal injury cases, and in corporate negligence claims. The legal system, for all its flaws, gave ordinary people a fighting chance to hold powerful institutions accountable.

Over the last several years, I have watched a new kind of violation take shape, one happening silently and constantly inside the devices you carry, the homes you live in, and the cars you drive. Clients now walk through my door dealing with stolen identities built from data broker profiles, fraudulent accounts opened with information purchased for pennies, reputations destroyed by AI generated deepfakes, and children whose digital footprints formed before they learned to walk.

I wrote this book because I am worried. I am worried about Americans who do not realize how much of their private lives companies have collected and sold. I am worried about my kids and grandkids growing up in a world where surveillance starts before they take their first steps. I am worried about a democracy where citizens lose the ability to think, search, move, or make personal decisions without generating a record someone else owns. And I know something most Americans do not: the legal protections you assume exist do not actually exist. The United States stands alone among major democracies as the only country without a strong federal privacy law. Americans have a patchwork. Some states offer real protections. Most do not.

What You Do Not Know Is Costing You

In 2025, the Identity Theft Resource Center recorded 3,322 data breaches in the United States, sending nearly 279 million victim notices to affected Americans. Eighty percent of Americans surveyed said they received at least one breach notification letter in the previous twelve months. Seventy percent of those letters did not even disclose how the breach happened, leaving recipients unable to assess their own risk.

Behind the breach headlines, data brokers maintain detailed profiles on more than 250 million American adults, containing your name, address, phone number, email, estimated income, political affiliation, religious identity, health conditions, and hundreds of additional data points. More than 750 data broker companies registered in California alone. You did not sign up for any of them. An investigation by CalMatters and The Markup found 35 data brokers intentionally hid their opt out pages from search results, making removal nearly impossible.

Think about an ordinary Tuesday. Your phone sent your location to advertisers and data aggregators. Your car logged your route, your speed, and whether you buckled your seatbelt. Your credit card purchase at a coffee shop joined a stream of transaction data flowing to marketing companies. Your search engine queries became data points in a profile tied to your device. Your smart speaker sent audio to a cloud server. Your television logged every program on your screen. Your health app stored a symptom search in a database no health privacy law governs. If someone followed you around all day writing down every place you went, every purchase you made, and every question you typed, you would call the police. The digital version of this surveillance happens every day, and the companies doing the tracking already have a device in your pocket sending the information for them.

The AI Accelerant

Artificial intelligence has made every dimension of this crisis faster and more dangerous. Voice cloning now creates a near perfect replica of any person's voice for less than five dollars and a few seconds of audio. Parents have received calls from what sounded exactly like their children begging for money. Cloned voices have fooled bank representatives during verification checks. The FTC reported Americans lost 12.5 billion dollars to consumer fraud in 2024, a 25 percent increase over the prior year, with AI powered scams growing fast. Romance bots running entirely on artificial intelligence now conduct long term relationships with victims, building emotional bonds over weeks before requesting money, operating around the clock in multiple languages at a scale no human scammer has ever achieved.

Researchers at Stanford found all six major AI companies train their models on user conversations by default. When you describe a medical condition or a family conflict or a legal problem to a chatbot, the conversation becomes training data. The models themselves learned by scraping billions of web pages, pulling in personal

photos, resumes, blog posts, and financial records belonging to people who never consented. Deleting your data from a trained model is technically impossible. Millions of records containing personally identifiable information sit inside AI training sets, and American citizens have no clear legal path to removal.

Your Body, Your Health, Your DNA, Your Children

Facial recognition systems now operate in airports, stadiums, retail stores, and public streets. TSA installed facial recognition at more than 80 airports. Clearview AI scraped billions of social media photos for a law enforcement database. Unlike a password or credit card number, you do not get to request a new face. A breach of biometric data lasts forever. Illinois passed the Biometric Information Privacy Act. Most states have no similar law.

Most Americans assume HIPAA protects all health related information. HIPAA does not. The fitness tracker on your wrist, the symptom checker on your phone, the wellness app tracking your sleep and exercise fall entirely outside the law. Reproductive health data has become uniquely dangerous in post Dobbs America, where period tracking apps and location visits to reproductive health clinics generate potential evidence in states where abortion carries criminal penalties. A federal court in June 2025 struck down a rule protecting reproductive health records under HIPAA, and the administration chose not to appeal.

When 23andMe filed for bankruptcy, the genetic data of approximately 15 million customers entered a legal limbo no privacy law addresses. Your DNA does not expire, does not change, and does not belong to you alone. The information reveals details about your parents, your siblings, your children, and generations not yet born. State attorneys general issued urgent alerts urging users to delete their data before a sale transferred ownership to an unknown buyer.

Children face surveillance from their earliest digital moments. Games, apps, educational platforms, and school systems harvest data on children long before those children understand what data means. The FTC brought enforcement actions against education technology companies storing records on more than ten million students. Age verification systems create new privacy problems by requiring biometric data to prove age. The tools marketed to protect children expose children to entirely new forms of data collection.

The System Built to Fail You

The absence of a federal privacy law is not an oversight. Powerful industries have spent decades lobbying Congress to keep privacy legislation from reaching the finish line. Cookie banners are engineered to exhaust you into clicking Accept All. Privacy policies average more than 7,000 words. The FTC has documented opt out buttons placed on websites knowing those buttons did not function. California has brought enforcement actions against businesses ignoring Global Privacy Control signals. The entire architecture of digital consent exists to ensure you agree without understanding what you are agreeing to.

Behind consent theater, companies use your browsing history and location and income estimates to charge you different prices for the same products. Tenant screening companies compile reports determining whether you get an apartment. Background check firms generate dossiers influencing whether you get a job. Errors are common. Most people never learn these reports exist until a decision has already gone against them. And federal agencies including ICE, the FBI, and the DEA routinely purchase location and behavioral data from

commercial brokers, bypassing the Fourth Amendment entirely. The Fourth Amendment Is Not For Sale Act passed the House and died in the Senate. And someone is still tracking you.

What This Book Will Give You

This book takes you from understanding to action. The early chapters reveal the structural failures making everything else possible. The middle chapters walk you through biometric surveillance, government surveillance, deepfake technology, AI powered fraud, chatbot data collection, the collapse of health data protections, reproductive privacy dangers, the genetic privacy crisis, and the surveillance tracking your children. The later chapters expose phone number hijacking, data breaches, government data purchases, algorithmic pricing discrimination, and consent theater. Each chapter includes real stories. Each chapter explains the mechanism. Each chapter tells you what the law protects and what the law leaves exposed. Each chapter ends with concrete steps you take right now. The final chapter gives you a thirty minute privacy action plan for you and your family.

I did not write this book for privacy professionals, Silicon Valley engineers or politicians. I wrote this book for the parent checking their kid's phone at the dinner table, for the grandmother who got a call from someone who sounded exactly like her grandson asking for money, for every American who has ever clicked Agree without reading what they agreed to, which is all of us.

I have spent my career fighting for people who were told to read the fine print, to accept the terms, to trust the process. The process is broken. The fine print is a trap. You deserve to know what is happening to your information. You deserve to understand who profits from the collection and sale of your personal life. You deserve a guide who will tell you the truth in plain language, without hedging and without pretending the situation is less urgent than the situation truly is.

The companies profiting from your data are counting on you to stay confused, to stay overwhelmed, to stay passive. Every chapter you are about to read replaces confusion with clarity. Every action step replaces passivity with power.

You did not agree to be the product. You do not have to remain the product. Start reading. Start fighting back. Your privacy is worth the fight, and so are you.

Copyright and Disclaimer

COPYRIGHT NOTICE

DISCLAIMER

Let me be real with you before we go any further.

I wrote this book because I believe every American deserves to understand what is happening to their personal information and what they are able to do about it. I have spent decades as a trial lawyer protecting consumers, and I have seen firsthand how much damage happens when people do not know their rights. This book is my way of putting that knowledge into your hands.

Here is what you need to understand.

This book is for education and general information only. Full stop. I am a lawyer, and a pretty good one, and I am proud of the work my firm does every single day. I am not, as of this moment, your lawyer. Nothing in this book, on any associated website, podcast, social media post, digital resource, Web3 platform, spatial computing environment, or any other format connected to this work creates an attorney-client relationship between you and me, between you and my law firm, or between you and anyone on my team. You do not become my client by reading this book, attending an event, sending a message, or interacting with me or my team online or in person.

This is not legal advice. This is not financial advice. This is not medical advice. This is not tax advice. And this is not professional counsel tailored to your specific situation. Every person's circumstances are different, and privacy law is changing fast at the federal, state, and international levels. What applies in California might not apply in Texas. What is accurate the day this book goes to print might shift by the time you read it.

Here is what I need you to do. Before you make any legal, financial, or personal decision based on something you read in these pages, talk to a qualified professional who knows your situation. Hire a lawyer in your jurisdiction. Consult a financial advisor. Speak with a cybersecurity specialist. Get advice that is specific to you.

The information in this book is provided "as is," without warranties of any kind, either express or implied. I have done my best to make sure everything here is accurate and current as of the date of publication, and my team and I put in serious work to get the details right. Laws change. Regulations shift. Technology evolves. Court decisions get reversed. I do not guarantee that every piece of information will remain accurate after publication, and I am not responsible for errors, omissions, or outcomes that result from actions you take based on this material.

By continuing to read this book, by using any associated resources, by interacting with me or my team through any platform or medium—including websites, social media, podcasts, live events, digital communities, Web3 environments, virtual worlds, blockchain applications, AI-powered tools, or any technology that has not been invented yet—you acknowledge and accept full responsibility for your own decisions and actions. That includes the risks associated with rapidly evolving technology, artificial intelligence, smart contracts, virtual assets, digital privacy tools, and global regulatory changes that no single book is able to predict.

If you are not comfortable with any of that, put this book down. Seriously. No hard feelings.

If you are still here, good. That tells me you are ready to learn, ready to take action, and ready to protect yourself and the people you love.

Let's go.

Chapters

Chapter 01: The Privacy You Trust Isn’t Real

Your Zip Code Decides Your Rights. And Nobody Told You.

Here is something most Americans have never been told, and the people profiting from your personal data are counting on you never finding out. The United States of America, the country founded on individual liberty, is the only G7 nation on Earth without a comprehensive federal privacy law. One hundred and forty four countries protect their citizens' personal data with national legislation. Ninety eight percent of developed nations guarantee their people baseline digital privacy rights. And the United States, in 2026, does not.

Your privacy rights in this country depend entirely on where you live. If you live in California, you hold a set of enforceable digital privacy rights in your hands right now. You can find out exactly what a company knows about you. You can order them to delete every piece of personal information they have collected. You can stop them from selling your data. And if they lose your data in a breach because they failed to protect your information, you can sue them.

If you live in Alabama, Mississippi, Georgia, New York, Pennsylvania, Massachusetts, or any of thirty one other states plus the District of Columbia, you have none of those rights. Zero. You have no legal way to find out what companies collect about you, no legal mechanism to tell them to delete your data, no legal authority to stop them from selling your most intimate details, and no ability to hold them accountable when they fail you.

Two people living in two different states can download the exact same app, hand over the exact same personal information, and face the exact same data breach. One of them has legal recourse. The other has nothing. The only difference between them is their address.

This did not happen by accident. Congress has debated, drafted, negotiated, and thrown away comprehensive privacy legislation for more than twenty five years. Every single time, the same forces killed the bill. The tech industry, which spent more than $250 million on federal lobbying in 2024 alone, fought to keep the system exactly as broken as it is. And here you are, clicking "I agree" hundreds of times a year on privacy policies nobody reads, believing you made a choice, when in reality the system was designed to make sure you never had one.

Twenty Five Years of Broken Promises

The Federal Trade Commission formally asked Congress to pass a comprehensive federal privacy law in May 2000. That was twenty six years ago. Every president since has acknowledged the problem. Not one has signed a solution.

President Obama released a Consumer Privacy Bill of Rights in 2012 with seven principles, including individual control, transparency, security, and accountability. The proposal was voluntary. Nobody in Congress even introduced the formal draft bill when the White House released legislative language in 2015. The Cambridge Analytica scandal exploded in 2018 and members of Congress introduced multiple privacy bills in the aftermath. Not a single one of those bills made it to a floor vote in either chamber.

The closest the country has ever come was the American Data Privacy and Protection Act in 2022. A bipartisan group of legislators from both parties sponsored the bill, and it passed the House Energy and Commerce Committee by a vote of 53 to 2. That was the first time in American history that a comprehensive federal privacy bill cleared a congressional committee. The bill included data minimization requirements, individual rights to access, correct, and delete personal data, algorithmic impact assessments, and the right for individuals to sue companies for violations.

Three forces killed the bill before it reached a full vote. The chair of the Senate Commerce Committee blocked the bill from advancing because she believed the enforcement provisions were too weak. The Speaker of the House prevented the bill from reaching the House floor because she wanted California to maintain its stronger protections. And the California delegation, including the governor, the attorney general, and the state's own privacy agency, opposed the bill because a federal law would have overridden California's pioneering protections.

A new version called the American Privacy Rights Act appeared in 2024 with sponsors who had been on opposite sides of the previous fight. For a moment, progress seemed possible. Then controversial revisions stripped out the bill's civil rights and algorithmic accountability provisions, major civil rights organizations withdrew their support, the scheduled committee vote was canceled, and the bill died when the 118th Congress ended.

The current Congress has not even introduced a comprehensive privacy bill. In early 2025, the House Energy and Commerce Committee created a privacy working group composed entirely of members from one party and issued a public request for information. The overwhelming majority of responses came from industry groups supporting the position that federal law should cap privacy protections at a single level and prevent states from going further. As of March 2026, we are right back where we started.

The three disputes killing every bill have stayed identical for a quarter century. Should federal law set a ceiling that prevents states like California from offering stronger protections, or should federal law set a floor that allows states to go further? Should individual Americans be able to sue companies for privacy violations, or should enforcement be left exclusively to regulators? How strictly should companies be limited in the data they collect about you? The tech industry has a preferred answer to each of these questions. And for twenty five years, the industry's preferred answer has won.

The Zip Code Lottery

Because Congress has refused to act, states have stepped in. Nineteen states have now passed comprehensive consumer data privacy laws. Five states led the way through 2022, including California, Virginia, Colorado, Connecticut, and Utah. Seven more followed in 2023. Another seven passed laws in 2024. 2025 was the first year since 2020 that no new state passed a comprehensive privacy law, though nine states updated their existing statutes.

These laws are not the same. They are not close to the same. They create a confusing patchwork where your rights depend on which side of a state border you happen to live on, and the differences between them matter enormously.

California stands alone with the broadest protections. California covers employee and business data, runs the country's only dedicated privacy enforcement agency, and gives residents a private right of action for data breaches with statutory damages of $100 to $750 per person per incident. That means a breach affecting a million people creates potential liability exceeding $375 million. California also created the DELETE Act, which launched a centralized platform in January 2026 called DROP. This platform allows Californians to submit a single deletion request that reaches all 545 registered data brokers at once. Before DROP existed, deleting your data from data brokers meant contacting hundreds of companies one by one. Nobody was doing that.

A handful of states occupy a second tier with notably strong provisions. Maryland prohibits companies from collecting data beyond what is strictly necessary for the service you requested, and Maryland bans the sale of sensitive data outright. Connecticut requires businesses to tell you if your data was used to train artificial intelligence models. Minnesota mandates data inventories and gives consumers the right to challenge profiling decisions.

The majority of states with privacy laws follow a more business friendly model. Attorney general enforcement only. Thirty day cure periods giving companies a month to fix violations before facing consequences. Standard consumer rights without the distinctive features California created. No private right of action for any violation.

Here is what this patchwork means in real life. In California, you can file a request to find out what a company knows about you, and that company must respond within 45 days with specific details about every piece of your personal information, where they got it, why they have it, and who they shared it with. In Alabama, you have no legal mechanism to compel any company to tell you anything.

In California, you can submit a deletion request that forces a company to delete your data and notify every third party that received your data to do the same. You can use DROP to wipe your information from the entire data broker industry with one request. In Alabama, you have no right to request deletion from anyone.

In California, if a company loses your data in a breach because they failed to maintain reasonable security, you can sue for statutory damages without having to prove you suffered a specific financial loss. In Alabama, you must pursue traditional negligence claims that require you to prove concrete injury. That bar is one most breach victims cannot clear.

Thirty one states plus the District of Columbia have no comprehensive consumer privacy law at all. New York, Illinois, Pennsylvania, Massachusetts, Washington, and Georgia are among the major states without one. Residents of those states have no legal right to find out what companies collect about them, no right to demand deletion, no right to opt out of data sales, and no way to hold companies accountable when their data is misused.

The narrow federal laws that exist, covering health data through HIPAA, children's data through COPPA, education records through FERPA, and financial data through the Gramm Leach Bliley Act, are riddled with gaps. They cover only specific sectors and specific types of data held by specific types of companies. They leave enormous categories of personal information entirely unprotected. And they do not come close to providing the kind of comprehensive coverage that citizens of 144 other countries already have.

The Grand Illusion of Clicking I Agree

Every time you sign up for an app, visit a website, or download software, a box pops up asking you to agree to a privacy policy. You click the button. You move on. You believe you made a choice.

You did not.

The entire system of "notice and consent" that supposedly protects your privacy is built on a fiction. The idea is straightforward enough. Companies tell you what they do with your data in a privacy policy. You read the policy. You understand the terms. You agree. Every single element of that idea is false.

The average American encounters approximately 1,462 unique websites per year, and each one has its own privacy policy. Those policies now routinely run 4,000 to 5,000 words or more. Reading every privacy policy you encounter in a year would take approximately 244 hours. That is 76 full work days. If every American internet user actually read every privacy policy they encountered, the aggregate time cost would be worth approximately $781 billion per year. That number exceeds the gross domestic product of the entire state of Florida.

Even if you tried to read them, you would struggle to understand them. Privacy policies are consistently written at a college reading level, with social media platform policies scoring at the difficulty level of academic journal articles. The average American reads at an eighth grade level. The policies are not written for you to understand. They are written for a company's lawyers to defend.

In practice, 78% of Americans routinely click agree without reading privacy policies. Among people aged 18 to 34, a staggering 97% consent to terms and conditions without reading them. In one academic experiment, 98% of participants agreed to a consent form that included a clause granting the researchers naming rights to their firstborn child. Even among the 11% of participants who claimed they usually read terms carefully, virtually all of them agreed to give away naming rights to their own child.

More than half of American adults do not even understand what the phrase "privacy policy" means. Many people believe a privacy policy is a promise that the company will keep their data private. One leading researcher in the field concluded plainly that informed consent at scale is a myth.

The academics and privacy law scholars who study this field have taken the system apart piece by piece. The cognitive problems include the fact that human beings cannot rationally evaluate thousands of separate privacy decisions across hundreds of companies over years of accumulated data collection. The structural problem goes deeper. Privacy harms come from the gradual accumulation of data over time by many different companies. No individual can assess the total risk of saying yes to any single consent screen because the danger comes from the combination of all of them over months and years.

One Princeton professor described the situation as "privacy theater." We pretend to have read the privacy policies. The companies pretend we understood and agreed. Everyone plays their part. And the data keeps flowing.

The Federal Trade Commission itself has admitted the system does not work. A former FTC Chair called the notice and consent framework "outdated and insufficient." Another commissioner stated that click through

consent does not present a meaningful choice because consumers often have no real alternative. The FTC's 2022 rulemaking proceeding directly questioned whether consumer consent is an effective way to determine if a practice is unfair or deceptive.

While regulators publicly acknowledge the failure of notice and consent, companies are actively making it worse through what privacy experts call dark patterns. These are manipulative design elements engineered to trick you into giving up your data. A major international review found that 76% of websites and apps employed at least one dark pattern and 67% used multiple dark patterns. One company's cookie banner required two clicks to reject advertising cookies and only one click to accept them. Amazon's internal name for its Prime cancellation process was "Iliad," a reference to the long and grueling Trojan War epic. California's law now explicitly states that agreement obtained through dark patterns does not count as consent. The rest of the country has no such protection.

The Privacy Tool You Have Never Heard Of

Buried inside your browser settings is the single most practical privacy tool available to Americans right now. Almost nobody knows about it. And the companies collecting your data are perfectly happy to keep it that way.

Global Privacy Control, or GPC, is a web standard launched in 2020 that automatically tells every website you visit that you do not want your personal data sold or shared. You turn it on once, and from that moment forward, every website you visit receives an automatic signal communicating your opt out preference. The signal travels with every page you load. You never have to think about it again.

GPC was created by a coalition that included privacy researchers, the Electronic Frontier Foundation, Consumer Reports, major newspapers, and browser makers. In 2024, the World Wide Web Consortium adopted GPC as an official work item on the formal internet standards track. GPC carries legal force in a growing number of states. That is the critical difference between GPC and the failed Do Not Track standard from years ago. Do Not Track was voluntary. Companies ignored it. GPC has teeth.

Three browsers currently support GPC out of the box. Brave enables it by default with no way to turn it off. DuckDuckGo enables it by default. Firefox makes it available in your settings. Google Chrome, Apple Safari, and Microsoft Edge do not support it at all. Extensions like Privacy Badger and DuckDuckGo Privacy Essentials can add GPC to Chrome and other browsers. About 50 million people use GPC today, spread across the browsers and extensions that support it. That represents less than 10% of web users, because over 90% of people use browsers that do not offer GPC.

That is about to change. In October 2025, California's governor signed the Opt Me Out Act, which requires every browser operating in California to include built in opt out signal functionality by January 1, 2027. Chrome holds approximately 65% of global browser market share. Safari and Edge hold substantial shares of the rest. This law will force the biggest browsers in the world to give their users the ability to send a GPC signal or an equivalent with a single toggle.

Twelve states already require businesses to honor GPC signals. California, Colorado, Connecticut, Montana, Texas, Delaware, Oregon, Nebraska, New Hampshire, New Jersey, Minnesota, and Maryland all mandate compliance. Four of those states have named GPC specifically in their regulations.

Enforcement has been real. The first major California privacy enforcement action, a $1.2 million settlement with a beauty retailer in 2022, centered on the company's failure to honor GPC. When consumers activated GPC on the company's website, nothing happened. Data kept flowing to advertisers. In 2025, the California Privacy Protection Agency fined three companies a combined total exceeding $2.3 million for failing to honor opt out signals. A joint enforcement sweep involving the attorneys general of California, Colorado, and Connecticut marked the first multi state action targeting GPC noncompliance.

Here is the fact that should make you stop and think. When researchers offered people the option to enable GPC during browser setup, 94% chose to turn it on. The problem is not that people do not want privacy. The problem is that the most popular browsers in the world have not given people the option. The Opt Me Out Act changes that equation entirely.

If you are reading this, go to your browser settings right now and look for Global Privacy Control. If your browser supports it, turn it on. If your browser does not support it, download one that does, or install an extension that adds GPC. This is the single easiest step you can take today to exercise your privacy rights.

144 Countries Got There Before Us

At least 144 countries have enacted comprehensive national data protection laws. The European Union treats privacy as a fundamental human right and enforces data protection through independent agencies in every member state with fines reaching up to 4% of a company's global revenue. Cumulative fines under Europe's General Data Protection Regulation have reached approximately 6 billion euros, and eight of the ten largest fines have been imposed on American companies. One American social media company alone has been fined more than 2.2 billion euros.

Brazil passed a comprehensive data protection law in 2018. South Korea has one of the strictest data protection regimes in the world with penalties up to 3% of global revenue. Japan earned the first data adequacy decision from the European Union of any Asian country. India passed its national data protection law in 2023 as the nineteenth G20 country to do so. Australia enacted sweeping privacy reforms in December 2024.

The United States stands alone among its peers. And American companies are paying the price internationally. Roughly 63% of total European data protection fine value has been imposed on companies headquartered in the United States. The absence of strong domestic protections makes American companies easier targets for international enforcement.

The European Union has twice invalidated the legal framework for transferring personal data from Europe to the United States, each time concluding that American law does not adequately protect individuals from government surveillance and does not provide meaningful recourse. A current framework adopted in 2023 survived its first legal challenge in 2025, and further challenges are expected.

Every time an American company gets hit with a massive European privacy fine, that is a direct consequence of the fact that our own government has not required those companies to protect your data in the first place.

Forty People Standing Between You and an Entire Industry

Without a federal privacy law, the Federal Trade Commission has become the country's default privacy regulator by using a single provision of the FTC Act, Section 5, that does not even contain the word "privacy." The FTC has used its authority to pursue unfair and deceptive practices to bring privacy cases against some of the biggest companies in the world, including a $5 billion settlement with a major social media platform and a $520 million case against a video game company.

About 40 full time FTC staff members work on privacy issues. The United Kingdom's privacy enforcement office employs more than 700 people to serve a country one fifth the size of the United States. The FTC brings roughly 20 privacy cases per year across a digital economy that touches every single American, every single day. The agency cannot impose civil penalties for first time violations. It can only get a consent order and then penalize future violations.

Under the current administration, the FTC has pulled back from aggressive privacy enforcement. Staffing has been cut. The commercial surveillance rulemaking process that the agency launched in 2022 to create comprehensive privacy rules appears to have been shelved. The Privacy and Civil Liberties Oversight Board lost its quorum when its independent members were removed.

The Supreme Court provided a narrow form of digital privacy protection in 2018 when it ruled that the government needs a warrant to access historical cell phone location data. That decision recognized that such data reveals a detailed chronicle of a person's movements. The ruling was deliberately limited and did not address many other forms of digital surveillance.

The largest privacy enforcement victories in the country have come not from federal action, but from state laws. A $1.4 billion settlement against a major social media company for facial recognition violations came under Texas state law. A $51.75 million settlement against a facial recognition company for scraping 60 billion images from the internet came under an Illinois state law. These outcomes are possible because a few states decided to act. They are not available to the 190 million Americans living in states without those protections.

What You Want, What You Deserve, and What You Are Not Getting

Eighty one percent of Americans say they are concerned about how companies use their data. Seventy three percent say they have little or no control over what companies do with their personal information. Seventy two percent want more government regulation of corporate data practices, a position supported by majorities of both Republicans and Democrats. Eighty four percent support stricter federal data privacy laws. This is one of the most bipartisan issues in the country.

A federal privacy law based on the bills Congress has already drafted and debated would change your daily life in ways you would feel immediately. Companies would be prohibited from collecting more data than they need to provide the service you asked for. Every American would have the right to find out what data companies hold, to correct mistakes, to delete their information, and to stop their data from being sold to third parties and advertisers. A single setting in your browser would legally require every company in the country to stop selling your data. If an algorithm denied you a loan or insurance, you would have the right to know why and to challenge the decision. Data brokers would have to register with the government and comply with deletion requests. If a company lost your data because they failed to maintain reasonable security, you could hold them accountable in court no matter what state you live in.

Health apps would be barred from selling your prescription information to advertisers. Companies could not collect your face scan, fingerprint, or voiceprint without your explicit permission. The right to sue for privacy violations would belong to every American, not only people who happen to live in California. Data brokers, the invisible companies that compile detailed profiles about you including your home address, your income, your health conditions, and your political views, would finally face real accountability.

This is not a wish list. Every one of these provisions has already appeared in a bill that received bipartisan support in Congress. They did not become law because the industries profiting from the current broken system spent hundreds of millions of dollars making sure they did not.

What Comes Next Is Up to You

You did not choose this. Nobody asked you whether you wanted a country where your privacy depends on your zip code. Nobody asked you whether you wanted to live under a system where 78% of people click "I agree" without reading privacy policies because the alternative would take 76 work days per year. Nobody asked you whether companies should be allowed to collect, sell, and lose your most personal information with no meaningful consequences.

The absence of a federal privacy law is not an oversight. For twenty five years, the same industries that profit from unregulated data collection have spent billions of dollars to block the same legislation that 84% of Americans say they want. Every year that Congress fails to act, your personal data becomes more exposed, more collected, more bought and sold, and more at risk.

One hundred and ninety million Americans live in states with no comprehensive privacy protection at all. If you are one of them, you currently have no legal right to know what companies collect about you, no right to stop the sale of your personal information, and no way to hold them accountable when things go wrong.

You can start protecting yourself today. Turn on Global Privacy Control in your browser. If your browser does not support it, switch to one that does. If you live in California, use the DROP platform to delete your data from all registered data brokers at once. If you live in a state without a privacy law, call your state representative and ask them why you have fewer rights than someone living one state over.

The most important thing you can take from this chapter is the knowledge that the system was not built to protect you. It was built to extract from you. The privacy policies you click through are not choices. They are rituals designed to make you feel like you agreed when you never had a meaningful option. The patchwork of state laws is not a solution. It is a symptom of a Congress that has chosen industry donors over the people it represents for a quarter of a century.

This fight is about more than your data. It is about whether you have the right to live a life that is not recorded, profiled, scored, and sold without your knowledge. It is about whether your children and grandchildren grow up in a country where their personal information is their own, or where their entire digital existence belongs to whoever can afford to buy it. And it is about whether our democracy can function when the people who are supposed to represent us keep choosing the interests of the companies harvesting our lives over the clear and overwhelming will of the American public.

You deserve the same protections that citizens of 144 other countries already have. Congress knows this. The data proves this. Eighty four percent of your fellow Americans agree.

Now the question is what you are going to do about it.

Chapter 02: Data Brokers- They Sold You for a Penny

Inside the Hidden Industry That Knows Your Secrets, Sells Your Identity, and Answers to No One

Right now, as you read this sentence, a company you have never heard of owns a file on you. That file contains your home address, your phone number, your Social Security number, your estimated income, your medical concerns, your religious affiliation, your political leanings, and a detailed record of everywhere your phone has traveled in the past year. That company did not ask your permission to collect any of it. And sometime today, that company will sell your file to a stranger for less than a dollar.

This is not a hypothetical. This is not something that could happen in the future. This is happening to you, to your parents, to your kids, and to roughly 250 million American adults every single day of the year.

The companies doing this are called data brokers. Most Americans have never heard that term. Most Americans do not know these companies exist. And that is exactly how data brokers want it, because the less you know, the more money they make. We are talking about a $300 billion global industry built on one foundational principle: your most personal information can be collected without your knowledge, packaged without your consent, and sold to anyone with a credit card. Advertisers buy it. Insurance companies buy it. Scammers buy it. Stalkers buy it. Foreign governments buy it. Your own government buys it, sometimes to avoid the hassle of getting a warrant.

In this chapter, I am going to show you exactly how this industry works, who the major players are, what they know about you, and what it costs to buy your life story. I am going to tell you about real people who were harmed, and in some cases killed, because a data broker sold their personal information to the wrong person. I am going to walk you through the strongest law in the country designed to fight back, California's Delete Act, and show you the tools that exist right now to start taking your privacy back. And I am going to explain why, despite everything you are about to read, almost everything the data broker industry does remains perfectly legal under federal law.

That last part should make you angry. By the time you finish this chapter, you will understand exactly why it should.

The Middlemen Who Know Everything

A data broker is a company that collects personal information about people it has no relationship with and sells that information to third parties. Think of it this way. You have never signed up for anything with Acxiom. You have never visited the Epsilon website. You have probably never typed the words "LexisNexis Risk Solutions" into a search bar. And yet each of these companies maintains a detailed profile on you that would make your closest friends uncomfortable.

California law defines a data broker as a business that knowingly collects and sells personal information about consumers with whom it has no direct relationship. That definition matters because it captures the core of the problem. These are not companies you chose to do business with. These are companies that built a business around you without ever telling you.

The data broker business model has three steps. First, collection. Brokers pull your information from public records like court filings, property deeds, voter registration, and DMV records. They scrape social media. They buy purchase histories from retailers, credit card transaction data, loyalty program records, and survey responses. They collect data from the invisible software tucked inside your phone apps. And they harvest enormous volumes of personal data from the online advertising system itself, which broadcasts your information to thousands of companies every time a web page loads on your screen or an ad appears in your app.

Second, aggregation. Brokers take all of these scattered data points and stitch them together into a single profile tied to your identity. They match records using your email address, your phone number, your Social Security number, and your physical address. They use algorithms that infer when two different data points belong to the same person, even when no shared identifier exists. Companies like LiveRamp maintain identity graphs that connect your online cookies, your device IDs, your email accounts, and your home address into one unified picture of who you are.

Third, sale. Once the profile is built, brokers sell it. They sell through subscriptions, through per-search pricing, through bulk data feeds, and through cloud marketplaces where buyers can browse available datasets the way you browse products on Amazon.

The scale is staggering. Acxiom, now owned by Omnicom through its $13 billion acquisition of IPG, maintains data on approximately 260 million Americans across 162 million households, with up to 10,000 attributes per person. Epsilon, owned by Publicis Groupe after a $4.4 billion acquisition, claims to hold data on every marketable household in the United States and processes 50 trillion data transactions annually. Experian holds records on 1.4 billion people worldwide. LexisNexis Risk Solutions sits on more than 78 billion records drawn from over 10,000 sources. Spokeo alone claims 6 billion consumer records.

How many data brokers are there? California's official registry lists 545 registered brokers as of January 2026. Vermont's registry shows 257. The Privacy Rights Clearinghouse and the Electronic Frontier Foundation built a combined database in 2025, cross referencing five state registries, and found more than 750 unique data broker groups. Hundreds of companies registered in one state had not registered in others, which suggests massive noncompliance with the law. Industry estimates put the global total at 4,000 or more. The global data broker market was valued between $270 billion and $323 billion in 2024, with the American share alone estimated at $131.7 billion.

Your Social Security Number Costs Less Than a Cup of Coffee

The depth of what data brokers know about you is difficult to overstate. A single broker examined by the Federal Trade Commission in 2014 maintained 3,000 data segments on nearly every American consumer, spanning 44 categories of health conditions. Acxiom currently advertises more than 1,500 attributes per person through its InfoBase product. With integrated partner data, that number climbs above 10,000. These are not vague generalities. The categories documented by federal investigations, academic research, and investigative journalism include Social Security numbers, income and net worth estimates, credit scores and bankruptcy status, specific health conditions including depression, anxiety, cancer, diabetes, erectile dysfunction, and pregnancy status, medications, religious affiliation, political party registration and voting history, sexual orientation, ethnicity, daily location patterns, purchasing habits, and web browsing history.

The labels brokers assign to the people in their databases are revealing in ways that should disturb every American. Documented marketing segment names include "Suffering Seniors," "Rural and Barely Making It," "Ethnic Second City Strugglers," "Retiring on Empty: Singles," and "Paycheck to Paycheck Consumers." In Congressional testimony, the World Privacy Forum documented lists of people with HIV/AIDS, people undergoing cancer treatment, people dealing with drug and alcohol addictions, and seniors suffering from dementia, all available for purchase. A leaked spreadsheet from Microsoft's ad platform Xandr, obtained by The Markup in 2023, revealed 650,000 audience segments including "Heavy Purchasers of Pregnancy Tests," "Depression Prone," and "Brain Tumor" interest categories, supplied by 93 different data companies.

And the price. The price is what makes the entire system possible, because personal data is almost unbelievably cheap. A Duke University study published in November 2023 purchased data on nearly 50,000 active duty military service members, including names, home addresses, health conditions, financial data, children's names, and religious affiliations, for approximately $10,000. That works out to somewhere between twelve cents and thirty two cents per person. Larger purchases dropped the per person cost to a single penny. The World Privacy Forum documented a list of "Rape Sufferers" available for purchase at seven point nine cents per name. Mental health data costs as little as five and a half cents per record. Researchers at Duke found that 11 of 37 data brokers they contacted agreed to sell mental health records, covering depression, anxiety, ADHD, bipolar disorder, and PTSD, with almost no vetting of who was buying. In May 2022, a journalist purchased a full week of location data covering more than 600 Planned Parenthood locations across the country for just over $160.

So who is buying? Advertisers represent the largest buyer category, using broker data for targeted advertising and what the Federal Trade Commission calls "surveillance pricing," which means adjusting what you pay for a product based on your personal data profile. Insurance companies buy data to inform underwriting and premium decisions. Employers and landlords run background checks powered by broker data. Political campaigns use commercially available data for voter targeting. Law enforcement agencies buy location data to track people without getting a warrant. Scammers buy lists of elderly and vulnerable Americans to run fraud schemes. Anti abortion organizations have purchased data showing which phones visited reproductive health clinics. Foreign entities have purchased detailed records on American military personnel. The Duke researchers found that purchasing data using an Asian domain name and a Singaporean IP address triggered virtually no additional screening from the brokers, which tells you exactly how easily a foreign adversary could exploit this system.

When a Data Broker Becomes an Accomplice

The cost of this industry is measured in human lives. On October 15, 1999, a man named Liam Youens drove to the workplace of Amy Boyer, a 20 year old dental hygiene student in Nashua, New Hampshire. He called her name so she would look up. Then he shot her through her car window and killed himself. Youens had been stalking Boyer since the eighth grade. He paid a data broker called Docusearch.com $45 for her Social Security number and $109 for her workplace address, which Docusearch obtained by making pretextual phone calls. Youens himself had written on his personal website: "It's actually obscene what you can find out about a person on the Internet." The New Hampshire Supreme Court later ruled that an investigator who obtains someone's work address through deception and sells it to a third party can be held liable for the harms that follow.

Twenty one years later, in July 2020, a self proclaimed anti feminist attorney named Roy Den Hollander appeared at the home of United States District Judge Esther Salas disguised as a FedEx driver and opened fire, killing her 20 year old son Daniel and critically wounding her husband Mark. The FBI found that Hollander kept detailed files on judicial targets. Their home addresses, as Judge Salas later testified, could "be purchased online for just a few dollars, including photos of our homes and the license plates on our vehicles." Congress eventually responded with the Daniel Anderl Judicial Security and Privacy Act in December 2022, specifically prohibiting data brokers from selling judges' personal information.

Then came the most devastating case yet. On June 14, 2025, a man named Vance Luther Boelter, 57 years old, wearing body armor and a silicone face mask, went to the home of Minnesota State Representative Melissa Hortman and assassinated her and her husband Mark. He then drove to the home of State Senator John Hoffman and shot him nine times and his wife Yvette eight times. Both survived. FBI agents found notebooks listing 45 elected officials, all from one political party, with home addresses, family members' names, and tactical surveillance notes. A separate list documented 11 different people search and data broker websites, with notations about which sites offered free trials and how to obtain home addresses from each one. Senator Ron Wyden responded publicly: "Every single American's safety is at risk until Congress cracks down on this sleazy industry."

Financial devastation runs parallel to the physical violence. Epsilon Data Management, one of the world's largest marketing companies, paid $150 million in 2021 to settle federal charges for knowingly selling data on more than 30 million Americans to scammers running fraudulent mass mailing schemes that targeted elderly victims from 2008 to 2017. A single fraudster client used Epsilon's lists to defraud more than 218,000 people out of $23.7 million. Epsilon's own internal records showed that more than 12,000 victims were defrauded more than 20 times each by this one scheme. Two Epsilon executives went to prison, one for ten years and the other for four.

And then there was the National Public Data breach of 2024, which exposed just how fragile the entire system is. National Public Data was a background check company that most Americans had never heard of. It was run by a single person, Salvatore Verini Jr., from a home office with two desktop computers, a laptop, and five Dell servers. When hackers broke in, they extracted 2.9 billion records, including 272 million unique Social Security numbers, which amounts to roughly 60 percent of all Social Security numbers ever issued in the United States. The stolen database, 277 gigabytes in size, was initially offered on the dark web for $3.5 million. Then it was released for free. National Public Data filed for bankruptcy with total assets between $25,000 and $75,000. One person, operating from a home office, had collected the most sensitive data on hundreds of millions of Americans. That person had no meaningful security infrastructure, no ability to notify victims, and no resources to face the resulting lawsuits.

California Fights Back and the Country Is Watching

On October 10, 2023, Governor Gavin Newsom signed Senate Bill 362, authored by Senator Josh Becker. It is called the California Delete Act, and it represents the most ambitious data broker regulation ever enacted in the United States. The law requires every data broker operating in California to register annually with the California Privacy Protection Agency, known as CalPrivacy. Registration costs $6,000 per year. And those fees fund something that has never existed before: a centralized system that lets you delete your data from every registered data broker in the state with a single request.

That system is called the Delete Request and Opt Out Platform, or DROP. It launched on January 1, 2026, right on schedule. CalPrivacy built it in partnership with the California Department of Technology at a cost of approximately $4.4 million. Here is how it works. If you are a California resident, you go to the DROP website, verify your identity through the state's Identity Gateway or Login.gov, provide your name, email, phone number, ZIP code, and optionally your mobile advertising ID or vehicle identification number, and hit submit. The entire process takes less than ten minutes. Your deletion request then goes out simultaneously to all 545 registered data brokers.

Starting August 1, 2026, brokers must pull the hashed consumer identifier lists from DROP at least every 45 days, process deletion requests within 45 days of retrieval, and report back with a status for each request. The possible responses are "Deleted," "Exempt," "Opted Out," or "Record Not Found." And here is the part that makes this law different from everything that came before it. After your initial deletion goes through, brokers must continue deleting any new data they collect about you every 45 days going forward. They are banned from selling new information about you after you submit your request.

Early results have exceeded expectations. CalPrivacy Executive Director Tom Kemp reported more than 215,000 consumer registrations by mid February 2026, with an average of roughly 7,000 new signups every day.

The enforcement structure gives this law real teeth. Failure to register as a data broker carries a penalty of $200 per day. Failure to process a deletion request carries $200 per request, per day. A broker that ignores one million pending deletion requests faces theoretical exposure of $200 million for every 45 day cycle. CalPrivacy launched a Data Broker Enforcement Strike Force in November 2025 and has brought at least nine enforcement actions against noncompliant brokers. One of those actions hit S&P Global with a $62,600 fine for 313 days of failing to register. Another required a company called Datamasters to stop selling all Californians' personal data after investigators found Datamasters was selling names, addresses, and phone numbers of millions of people with Alzheimer's disease, drug addiction, bladder incontinence, and other health conditions.

In October 2025, Governor Newsom signed SB 361, the Defending Californians' Data Act, which expanded the Delete Act's requirements even further. Brokers must now disclose whether they collect data on sexual orientation, union membership, citizenship or immigration status, biometric identifiers, and government issued IDs. They must also disclose whether they have sold data to foreign actors, to U.S. government or law enforcement agencies, or to developers of generative AI systems.

No formal legal challenge to the Delete Act has been filed as of March 2026, and the early enrollment numbers suggest that consumers are ready for this kind of tool. The question now is whether other states will follow California's lead.

What You Can Do Right Now to Start Taking Your Privacy Back

Before the DROP platform existed, the burden of opting out of data brokers fell entirely on you. And it was crushing. Multiple estimates put the time required for a full manual opt out at 100 hours or more per year, because each of the 500 plus known broker websites has its own verification procedures, its own confusing interfaces, and its own deliberately misleading design patterns meant to keep you from completing the process. CalPrivacy Executive Director Tom Kemp estimated 20 to 30 minutes of interaction with each

individual broker, adding up to what he described as "as much as 10 days' worth of work." Even after you did all of that, the results were temporary. Research shows that 73 percent of data brokers re add consumer information within 90 days of removal, with an average re listing time of just 23 days. The people search site Spokeo reportedly re adds data in 89 percent of cases within 60 days.

A growing number of commercial data removal services have entered the market to take on this burden for you. Here are the leading options as of early 2026. Optery, priced between $3 and $20 per month depending on the tier, was the top performer in Consumer Reports' 2024 study, achieving a 68 percent removal rate after four months across 13 tested people search sites. It also provides before and after screenshot proof of removals. EasyOptOuts, at $19.99 per year, achieved a 65 percent removal rate and is the cheapest option tested. Incogni by Surfshark, at $7.99 per month, covers more than 420 brokers with automated recurring requests and has been verified by a Deloitte audit. DeleteMe, at $129 per year, is the oldest service in the space with more than 15 years of operation and combines automation with human privacy experts, though it achieved only about 27 percent removal in the Consumer Reports study. Kanary, Privacy Bee, and Privacy Duck round out the market at various price points, with Privacy Duck offering a white glove, fully human approach at $197.50 per year.

If you live in California, the DROP platform is your best starting point, and it is free. For everyone else, four states currently maintain data broker registries. California has 545 registered brokers with a $6,000 annual fee. Vermont has 257 active registrants with a $100 fee and was the first state in the nation to create a registry back in 2018. Texas imposes $500 per day penalties for nonregistration and sent warning letters to more than 100 companies in June 2024. Oregon's registry took effect in January 2024 with $500 penalties capped at $10,000 annually. New Jersey, Delaware, Michigan, and Alaska have passed data broker registration bills at various stages of implementation. Nineteen states now have some form of comprehensive data privacy law that includes data broker implications, with Indiana, Kentucky, and Rhode Island laws taking effect in 2026. And Montana became the first state to explicitly prohibit law enforcement from purchasing citizens' data from brokers when they would otherwise need a warrant.

The structural challenge remains real. Data brokers refresh their databases every two to four weeks using public records, social media, purchase data, and partnerships with other brokers. You can delete your information today and find it back online within a month. California's Delete Act addresses this through its ongoing deletion obligation, requiring brokers to re check the suppression list every 45 days and maintain it permanently. For the rest of the country, the choice is between paid commercial services with modest success rates and a manual process so time consuming it amounts to a second job.

Legal Does Not Mean Right

Here is the hardest truth in this entire chapter: almost everything the data broker industry does is technically legal under current federal law. The United States has no comprehensive federal privacy statute. The laws we do have, the Fair Credit Reporting Act for credit data, HIPAA for health provider data, the Gramm Leach Bliley Act for financial institution data, were written for specific sectors and leave enormous gaps. Data brokers that collect health related information from your location data, your app usage, or your purchasing behavior fall outside HIPAA's reach. Brokers that claim they are not producing "consumer reports" for credit, employment, or insurance purposes argue they fall outside the Fair Credit Reporting Act's scope. The Federal Trade Commission's authority to go after "unfair or deceptive practices" has been useful in individual

enforcement actions, and even FTC Chair Andrew Ferguson has acknowledged that his agency's authority "is not a comprehensive privacy law."

Federal legislative efforts to close this gap have failed repeatedly. The American Privacy Rights Act, introduced in April 2024 by the chairs of both the Senate Commerce Committee and the House Energy and Commerce Committee, would have created data minimization requirements, established individual rights to access and delete data, and specifically defined data brokers under federal law. It collapsed after a contentious markup in June 2024 and died when the 118th Congress ended in January 2025 without being reintroduced. The Fourth Amendment Is Not For Sale Act, which would have stopped government agencies from buying data they would otherwise need a warrant to obtain, passed the House 219 to 199 in April 2024 and then died in the Senate. The only data broker related federal law that actually made it across the finish line was the Protecting Americans' Data from Foreign Adversaries Act, signed in April 2024 as part of the TikTok legislation. That law prohibits data brokers from selling sensitive data to China, Russia, Iran, North Korea, Cuba, and Venezuela. It does nothing about the sale of that same data within the United States.

The Consumer Financial Protection Bureau tried a different approach. In December 2024, Director Rohit Chopra proposed a rule that would have brought data brokers under the Fair Credit Reporting Act by redefining key terms in the statute. The rule would have required accuracy standards, consumer access rights, and limits on what data could be used for. Acting Director Russell Vought withdrew the rule on May 15, 2025, stating that "rulemaking is not necessary or appropriate at this time." That withdrawal was part of a broader pullback at the agency that included rescinding more than 60 regulatory guidance documents and reducing staff from over 1,600 to approximately 300.

The Department of Justice's Bulk Data Rule, implementing a February 2024 executive order, took effect in April 2025 and represents the most significant new federal restriction. It broadly prohibits data brokerage transactions involving countries of concern, including China, Russia, Iran, North Korea, Cuba, and Venezuela. It covers bulk sensitive data including geolocation from 1,000 or more devices, personal identifiers for 100,000 or more persons, health data, financial data, and biometric identifiers. The rule stayed in place under the current administration. Its scope, though, is limited to transactions involving foreign adversaries. Domestic data brokerage remains entirely unaffected.

The European Union treats this differently. Under the GDPR, the General Data Protection Regulation, any processing of personal data requires a lawful basis, typically informed and freely given consent. There are no carve outs for third party collectors. Fines reach 20 million euros or 4 percent of global annual revenue. The EU treats privacy as a fundamental human right. The United States treats personal data as a commercial asset that anyone can collect unless a specific law says otherwise.

Sit with this for a moment. The federal government recognized that the sale of Americans' bulk personal data is so dangerous to national security that it banned the practice when foreign adversaries are the buyers. The same data remains freely available to domestic scammers targeting people with Alzheimer's, stalkers tracking domestic violence survivors, and government agencies buying information they cannot constitutionally compel through a warrant. The danger is the same. The legality is completely different depending on who is writing the check.

Your Tax Dollars Are Buying Your Own Surveillance

The connection between data brokers and government surveillance represents one of the most constitutionally significant issues of our time. Federal agencies have systematically purchased Americans' personal data, especially location records, from commercial data brokers as a way to get around the Fourth Amendment's warrant requirement. A declassified report from the Office of the Director of National Intelligence, released in June 2023, stated this plainly: "The government would never have been permitted to compel billions of people to carry location tracking devices on their persons at all times, to log and track most of their social interactions, or to keep flawless records of all their reading habits. Yet smartphone and internet technologies have had this effect without government action."

Read that again. The intelligence community itself admitted that the surveillance apparatus now running through commercial data brokers is something the government could never have legally built on its own.

The purchasing is well documented. Immigration and Customs Enforcement has held a $9.75 million contract with LexisNexis and in 2025 resumed buying location data through a no bid contract with the surveillance company PenLink. Customs and Border Protection spent more than $2 million on Venntel contracts in 2019 and 2020. A report from the Department of Homeland Security Inspector General found that CBP, ICE, and the Secret Service all violated federal law through warrantless data purchases. The FBI purchased location data from Venntel and renegotiated its contract around the time of the 2020 Black Lives Matter protests. The Centers for Disease Control spent $420,000 on location data to track people's movement during COVID. Defense contractors purchased location data harvested from Muslim prayer and dating apps and sold it to the U.S. military. As of March 2026, more than 70 members of Congress have written to the DHS Inspector General demanding an investigation into ICE's warrantless location data purchases.

The advertising ecosystem is the pipeline that makes all of this possible. Every time a web page loads or an app shows you an ad, your personal data is broadcast through a system called real time bidding to thousands of companies simultaneously. The Irish Council for Civil Liberties calculated that this system shares the intimate characteristics of internet users 178 trillion times per year across the United States and Europe. A typical American's data is broadcast 747 times per day. Google alone allows 4,698 companies to receive this data. Only one advertiser wins each auction. Every company that participates receives the personal information in the bid request. Surveillance companies have exploited this by pretending to be advertisers, harvesting the data from the bidding stream, and reselling it as tracking tools to governments around the world. The Belgian Data Protection Authority called the real time bidding system "the biggest data breach ever recorded. And it is repeated every day."

Consumer awareness remains dangerously low. Pew Research Center found that 67 percent of Americans say they understand "little to nothing" about what companies do with their personal data. Only 6 percent of American adults have ever used a data removal service. About 40 percent had no idea that data brokers sometimes sell information to the United States government. And in one study that perfectly captures the absurdity of the "notice and consent" model we all live under, 98 percent of survey respondents agreed to a fake consent form that included language granting naming rights to their firstborn child.

This Is Not About Having Something to Hide

I want to leave you with something that cuts through the noise, because I know this chapter has thrown a lot at you.

The data broker industry has a body count. From Amy Boyer in 1999 to Representative Melissa Hortman in 2025, the same failure has repeated itself for over a quarter century: the unrestricted sale of home addresses through people search sites, available to anyone for a few dollars, with no meaningful federal law to stop it.

The industry's scale is almost impossible to wrap your mind around. Profiles with thousands of attributes on 250 million Americans. Sold for fractions of a penny. Regenerating within weeks of deletion. Feeding a surveillance system that the government itself uses to sidestep constitutional protections it would otherwise be required to respect.

California's Delete Act and the DROP platform represent the first serious attempt to shift the burden from you, the individual, to the institutions that profit from your data. The early results, more than 215,000 enrollments in the first six weeks, active enforcement with meaningful financial penalties, suggest the model works. Other states need to follow. And we need a federal law.

This is not about having something to hide. This is about whether you get to decide who knows the most personal details of your life. It is about whether a company you have never heard of gets to sell your Social Security number for a penny. It is about whether your government gets to buy a record of everywhere your phone has been without asking a judge for permission. It is about your kids growing up in a country where their most private information was collected before they could spell the word privacy.

I am asking you to do three things. First, if you live in California, register for DROP today. It is free and it takes ten minutes. If you do not live in California, look at the data removal services listed in this chapter and pick one. Second, check whether your state has a data broker registry or a comprehensive privacy law. If it does not, call your state representative and ask why. Third, share what you have learned. Tell your family. Tell your friends. Tell the people you love.

The data broker industry survives on one thing above all else: your silence. The moment you understand what is happening, the moment you start talking about it and taking action, their business model starts to crack. You deserve to control your own information. Your family deserves it. And our democracy depends on it.

Chapter 03: Your Phone Sends Your Location 747 Times a Day

Seven hundred and forty seven times. That is how many times your phone broadcasts your exact location to strangers every single day. More than once every two minutes during your waking hours, an invisible auction fires inside the device in your pocket or your purse, and it sends your GPS coordinates, your device ID, and a bundle of personal information to thousands of companies you have never heard of. These companies do not ask for your permission. They do not send you a notification. They collect this data in the time it takes to blink, and then they keep it forever.

You did not sign up for this. None of us did.

You downloaded a weather app, a game for your kid, a coupon app for gas, maybe a Bible app or a prayer app. And buried inside those apps, invisible lines of code are recording where you sleep, where you worship, where you see your doctor, where your kids go to school, and where you stop on Friday night. That information gets packaged, sold, resold, and sold again to advertisers, data brokers, hedge funds, insurance companies, and federal law enforcement agencies that would otherwise need a warrant to get the same data.

This chapter is going to show you exactly how this machine works. You are going to learn how a simple ad on your phone screen triggers a data broadcast that reaches thousands of companies in milliseconds. You are going to meet the data brokers who built billion dollar businesses packaging your movements into dossiers. You are going to read about the massive breach that dumped the precise locations of millions of Americans onto the open internet in January 2025. You are going to see how your car reports your location every three seconds, how license plate readers scan 20 billion plates every month, and how all of this data flows to government agencies buying their way around the Constitution. And at the end, you are going to learn exactly what to do about it on your phone, in your car, and in your browser, starting today.

Let me be direct with you. I am a trial lawyer. I have spent my career protecting people, and what I am about to show you in this chapter makes me angry. It should make you angry too. Because the surveillance system I am about to describe is not something that might happen to you someday. It is happening to you right now, as you read these words. Your phone is doing it. Your car is doing it. The camera on the pole at the end of your street is doing it. And nobody asked you.

The Invisible Auction Inside Your Phone

Every time you open an app that shows an ad, a high speed auction fires in milliseconds. Your phone packages up your GPS coordinates, your device ID, your IP address, the name of the app you are using, and your browsing history. All of that gets sent to an ad exchange, which blasts it out to dozens or hundreds of companies, each one representing advertisers who want to reach you. The highest bidder wins and their ad appears on your screen. The whole thing takes less time than a heartbeat.

Here is what most people do not realize. Every company that received your data during that auction keeps it. Whether they won the bid or not, they walk away with your GPS coordinates, your device ID, and all the other information that came with the bid request. There are no technical safeguards to stop this. There is no system that deletes the data after the auction ends. The companies simply take it.

The Irish Council for Civil Liberties put a number on this system in a May 2022 report. Real time bidding systems broadcast personal data 294 billion times per day in the United States alone. That is 178 trillion broadcasts per year across the U.S. and Europe combined. Google allows 4,698 companies to receive this bidding data about American users. Microsoft allows 1,647. Dr. Johnny Ryan, the former ad tech insider who led the research, called it the biggest data breach ever recorded because the data flows to companies around the world, including in Russia and China, and nobody controls what happens to it after that.

The number 747 comes from that same research. It means that your online activity and your physical location get exposed to the advertising system 747 times every single day. That is your life, minute by minute, broadcast to an industry that profits from knowing exactly where you are.

Your Favorite Apps Are Tracking Beacons

A French Navy officer went for a run on the deck of the Charles de Gaulle aircraft carrier and, like millions of people do, uploaded the workout to Strava. In doing so, he unintentionally revealed the precise location of a nuclear powered warship as it moved toward the Middle East. This was not the result of hacking or espionage. It was simply the consequence of sharing location data without thinking through the risks.

This incident is not isolated. It reflects a broader and ongoing issue with how fitness tracking apps handle privacy. Strava, by design, encourages sharing, and many accounts are set to public by default. That means each recorded run or workout can broadcast a user's location and movement patterns to anyone who looks. Over time, this kind of data has been used to identify sensitive locations, including military bases around the world.

A similar situation occurred in 2024 involving Emmanuel Macron. His movements were pieced together not through any direct breach, but by analyzing the publicly shared workout data of his security personnel as they traveled with him. Their activity effectively mapped his location.

The larger point is straightforward. You do not need to be a public figure or part of the military for this to matter. If your apps are sharing your location data by default, you may be exposing far more than you realize. Taking a few minutes to review and adjust your privacy settings is not optional anymore. It is a necessary step in protecting your personal security.

The advertising auction is another pipeline. Data brokers pay app developers to embed special code called Software Development Kits into their apps. These SDKs silently harvest your GPS coordinates, Wi-Fi signals, Bluetooth beacon data, and cell tower connections. They cross reference all of those signals at once, and they can pinpoint your location to within one to five meters. That is accurate enough to know which room of a building you are standing in.

The apps that carry this tracking code are the ones you use every day without a second thought. Prayer apps like Muslim Pro, which has 98 million downloads. Dating apps like Tinder and Grindr. Games like Candy Crush and Temple Run. Fitness apps like MyFitnessPal. Family safety apps like Life360. Shopping apps like GasBuddy. Weather apps. Navigation apps. Even VPN apps, the ones people download specifically because they want more privacy. One data broker, Venntel, claimed to pull location data from over 80,000 apps. When Gravy Analytics got hacked in January 2025, a single data sample exposed tracking code in 3,455 Android apps.

The money involved is staggering. The location data market was estimated at 12 to 16 billion dollars globally in 2021. Data brokers pitch app developers with offers of $12,000 to $1 million per year just for embedding the tracking code. Gravy Analytics alone claimed to process 17 billion signals daily from approximately one billion mobile devices. The New York Times obtained a single location dataset from the 2016 to 2017 era that contained 50 billion location pings from 12 million Americans. That was nearly a decade ago. The system has grown exponentially since.

From Your Phone to the Federal Government

The data travels through a supply chain that works like a set of nesting dolls. Your phone sends data to the app. The app passes it to the SDK or ad network inside it. That data flows to an aggregator like Gravy Analytics or Mobilewalla. From there, it reaches a government facing subsidiary like Venntel or Babel Street's Locate X product, which repackages the data for federal agencies. Mobilewalla disclosed that 60 percent of its location data came from ad exchange auctions it did not even win. The company simply collected your data by participating in the bidding process.

The buyers include the FBI, ICE, Customs and Border Protection, the IRS Criminal Investigation unit, the DEA, the Secret Service, and U.S. Special Operations Command. The FBI signed a $27 million contract with Babel Street, which purchased all of its data from Venntel. Commercial buyers include hedge funds, insurance companies, real estate firms, and anti abortion organizations. A company called Near Intelligence licensed location data to a group that geofenced 600 Planned Parenthood clinics across 48 states and served targeted ads to the women who walked through the doors. Near Intelligence boasted of tracking 1.6 billion people across 44 countries before filing for bankruptcy in late 2023.

The Federal Trade Commission has moved against several of these companies. In January 2025, the FTC banned Gravy Analytics and Venntel from selling sensitive location data. The FTC banned X Mode from selling raw location data in January 2024 after it sold military location data sourced from Muslim prayer apps. In December 2024, the FTC prohibited Mobilewalla from collecting consumer data from ad exchanges for non advertising purposes, a first of its kind order, after finding that the company had created audience targeting categories like "pregnant women," "Hispanic churchgoers," and "LGBTQ+ community" members. The FTC reached a settlement with Kochava in February 2026 for selling data that tracked people to reproductive health clinics.

The Gravy Analytics Breach: A Nightmare Scenario Becomes Real

On January 4, 2025, a hacker accessed the Amazon Web Services cloud storage of Gravy Analytics using a stolen access key. The hacker posted samples on a Russian language cybercrime forum and demanded a ransom. Pay within 24 hours or the full database goes public. The stolen haul included an estimated 17 terabytes of location data. A 1.4 gigabyte sample posted publicly contained roughly 30 million location data points. If the full dataset scales proportionally, it could contain over 200 billion records. The stolen files also included Gravy's entire customer list of more than 1,000 companies, internal emails, business strategies, and plaintext passwords.

The breached data included precise GPS coordinates, timestamps, device advertising IDs, the names of apps associated with each location ping, and historical movement patterns. Researchers demonstrated they could follow a single device ID traveling from New York to a home in Tennessee. The apps represented in the leak

included Tinder, Grindr, Candy Crush, Subway Surfers, Call of Duty Mobile, MyFitnessPal, Flightradar24, Muslim prayer apps, Christian Bible apps, pregnancy trackers, period tracking apps, and multiple VPN apps. Sensitive locations in the data included the White House, the Kremlin, Vatican City, and military bases around the world.

The most important revelation was how Gravy obtained its data. A senior threat analyst explained publicly for the first time that Gravy appeared to be collecting its data from the advertising bidstream itself, not from code embedded directly in the apps. That means the app developers often had no idea their users' data was being harvested. Tinder, Grindr, Flightradar24, and Muslim Pro all denied any direct relationship with Gravy. They all acknowledged displaying ads. The advertising pipeline itself was the collection mechanism.

The FTC had announced its enforcement action against Gravy and Venntel on December 3, 2024, with a unanimous 5 to 0 vote. The order was finalized on January 15, 2025, days after the breach became public. The order bans the companies from selling sensitive location data, requires them to delete all historic location data, and mandates they notify all customers from the past three years that the data must be destroyed. FTC Commissioner Alvaro Bedoya put the stakes in plain English. He said you may not know anything about Gravy Analytics, but Gravy Analytics may know whether you eat breakfast at McDonald's, whether you buy CBD oil, whether you recently bought lingerie, whether you are pregnant, and whether you sit in the pews every Sunday in Charlotte. He then pointed out the constitutional problem. The Supreme Court's Carpenter decision said you need a warrant to get this data. Venntel let the government get it without one.

When Location Data Exposes Your Most Private Moments

In July 2021, a small Catholic publication called The Pillar showed the world exactly how easy it is to strip away the fiction of "anonymized" location data. Using commercially available app signal data purchased from a data broker, the publication matched a mobile device identifier to Monsignor Jeffrey Burrill, the general secretary of the United States Conference of Catholic Bishops. The method was devastatingly simple. They matched the device's locations to places uniquely connected to Burrill. His office. His residence. A family lake house. A Wisconsin apartment. Once the device was identified, its history revealed visits to gay bars and a gay bathhouse in Las Vegas over parts of 2018, 2019, and 2020.

Burrill resigned the same day the story published. He later sued Grindr, saying his reputation had been destroyed.

A 2013 study at MIT analyzed 15 months of mobility data for 1.5 million people and found that just four location data points are enough to uniquely identify 95 percent of individuals. Four points. Fewer than you need to match a fingerprint. A 2021 follow up study found that 93 percent of people would be uniquely identified even in a dataset of 60 million using just four data points. Separate research found that 99.98 percent of Americans could be correctly re identified using 15 demographic attributes. The word "anonymized" is a lie. It is a marketing term designed to make you feel safe while companies sell the coordinates of your life to anyone with a credit card.

Abortion Clinic Visits for $160

In May 2022, days after the leaked Supreme Court draft overturning Roe v. Wade, a reporter at Vice demonstrated that SafeGraph was selling location data about visits to abortion providers on the open market.

He purchased a week's worth of data covering more than 600 Planned Parenthood locations across the country for $160. The data showed where groups of visitors came from by census block, how long they stayed, and where they went afterward. Some records contained as few as four or five devices, making individual identification straightforward.

After the Dobbs decision in June 2022, 22 states criminalized or severely restricted abortion. Location surveillance became an urgent threat. Washington State passed the My Health My Data Act, which specifically prohibits geofencing around healthcare facilities. California, New York, Connecticut, Illinois, and other states passed reproductive health data protections. The Biden Administration issued a HIPAA rule in April 2024 prohibiting the disclosure of reproductive health records for prosecution. A federal judge vacated that rule in June 2025, leaving the protection entirely up to the states.

Prayer Apps That Fed Your Location to the Pentagon

In November 2020, journalists revealed that Muslim Pro, a prayer time and Qibla direction app with 98 million downloads, was sending granular GPS coordinates to a data broker called X Mode Social. X Mode then sold that data to U.S. military contractors, and it ultimately reached the U.S. Special Operations Command. At least five additional Muslim prayer apps were linked to the same broker. Muslim Pro did not mention X Mode in its privacy policy. The Council on American Islamic Relations called for congressional hearings, calling it one of the first documented cases of the U.S. military purchasing the movement and location data of Muslim app users.

Location data has also been used to track people seeking addiction treatment. An analysis of 12 virtual opioid treatment websites found that all 12 used ad trackers capable of identifying visitors, and half of them embedded Meta Pixel, meaning people searching for help with substance abuse were being tracked by the advertising surveillance system, completely outside the protections of federal substance abuse confidentiality rules.

During the George Floyd protests in 2020, location and social media surveillance tools were pointed at demonstrators. An AI startup relayed protest related posts directly to police departments including the NYPD, LAPD, and Chicago PD. Documents obtained by the Brennan Center for Justice revealed that D.C.'s Metropolitan Police collaborated with federal law enforcement to surveil racial justice protesters through 2020 and 2021, with Capitol Police tracking groups including Black Lives Matter chapters with no evidence of violence.

Your Car Reports on You Every Three Seconds

On January 14, 2026, the FTC finalized its first enforcement action targeting connected vehicle data. The target was General Motors and its OnStar subsidiary. The complaint said GM collected precise geolocation data as frequently as every three seconds, along with hard braking events, acceleration data, speeds over 80 miles per hour, trip timing and distance, seatbelt usage, and even which radio stations drivers listened to. GM then sold this data to LexisNexis Risk Solutions and Verisk, which are consumer reporting agencies. Those companies compiled the data into driver risk profiles that insurance companies used to raise rates, deny coverage, or cancel policies.

The enrollment process was deceptive. GM marketed its Smart Driver feature as a free driving improvement tool. There was no disclosure that the data would end up in the hands of companies that score you for insurance purposes. Dealership salespeople were incentivized to sign buyers up, and in some cases penalized for not doing so. They filled in enrollment screens with preset answers. A New York Times reporter discovered that her husband's Chevrolet Bolt had been enrolled in the program even though neither of them remembered consenting. GM blamed a "bug."

The real world consequences hit fast. Kenn Dahl, a 65 year old Seattle driver with a clean record, saw his insurance premium jump 21 percent. When he requested his LexisNexis report, a document he did not know existed, it ran 258 pages and cataloged 640 trips with every instance of hard braking, rapid acceleration, and speeding. He told the Times it felt like a betrayal. Romeo Chicco, a Cadillac driver in Florida, had his premium doubled and his Liberty Mutual application denied based on his LexisNexis driving report.

The FTC's 20 year consent order bans GM from sharing geolocation and driver behavior data with consumer reporting agencies for five years. It requires the company to get clear, affirmative consent before collecting connected vehicle data. It mandates that consumers can disable all data collection without losing vehicle functionality. GM must delete all previously retained driver data within 180 days. No financial penalty was imposed.

Texas Goes After the World's Largest Driving Database

On January 13, 2025, Texas Attorney General Ken Paxton filed the first enforcement action by any state attorney general under a comprehensive data privacy law. The target was Allstate and its subsidiary Arity. Starting in 2017, Arity paid app developers millions to embed its driving behavior tracking code into popular apps like Life360, GasBuddy, and Fuel Rewards. That code captured data every 15 seconds or less, including geolocation, accelerometer readings, gyroscopic data, and derived events like speeding and distracted driving. Arity called what it built the "world's largest driving behavior database." It contained trillions of miles of driving data from over 45 million consumers.

The tracking code could not tell the difference between a driver and a passenger. A passenger's "bad driving" data could be attributed to them and raise their insurance rates. One consumer who had been using Life360 since 2020 reported that her premiums nearly tripled between 2022 and a December 2024 quote, despite having no accidents or tickets.

The Mozilla Foundation's 2023 "Privacy Not Included" report delivered a verdict that should terrify every car owner in America. After 600 hours of research across 25 car brands, every single brand earned a privacy warning label. A 100 percent failure rate. The worst of any product category in the program's seven year history. Eighty four percent of car brands can share your data with data brokers. Seventy six percent can sell it. Fifty six percent can share it with the government on an informal request. Nissan's privacy policy admits to collecting "sexual activity" and "health diagnosis data." Tesla warns that opting out of data collection may cause "reduced functionality, serious damage, or inoperability." Subaru claims that passengers consent to data collection simply by sitting in the car. The car data monetization industry is projected to reach $750 billion by 2030.

License Plate Readers Are Watching 20 Billion Times a Month

Automated License Plate Readers are high speed cameras equipped with optical character recognition software that capture every passing vehicle's plate number, GPS coordinates, timestamp, and photographs of the vehicle itself, including make, model, color, bumper stickers, and roof racks. A single ALPR can scan up to 2,000 plates per minute. They sit on poles and overpasses running around the clock. They ride on patrol cars. They hide inside speed enforcement trailers and construction barrels. The DEA once disguised one as a cactus.

Flock Safety, founded in 2017 in Atlanta, has become the dominant force in this space. The company has a $7.5 billion valuation as of September 2025, more than 80,000 cameras across 49 states, and over 5,000 law enforcement agencies as clients. Flock scans more than 20 billion license plates every month. Its "Vehicle Fingerprint" technology can track a car even without a readable plate, using make, model, color, and distinguishing features. More than 75 percent of agencies using Flock opt into a shared national database, meaning a single search can query over 83,000 cameras spanning nearly the entire country.

Here is the number that tells you everything. In Piedmont, California, the police department's own data showed that 99.97 percent of ALPR scans will never be used for a public safety purpose. The Brennan Center for Justice confirmed that an extremely small percentage of scanned vehicles, far below 1 percent, are connected to any crime. This means almost every scan captures the movements of innocent people going about their daily lives. And unlike your phone, you cannot turn off your license plate. Covering or tampering with it is a crime.

The data flows to federal agencies through overlapping networks. Customs and Border Protection regularly searched more than 80,000 Flock cameras through most of 2025. ACLU documents show that ICE officers maintained years long relationships with fusion center detectives who ran database searches on their behalf. Documented misuse cases include a Kansas police officer who used ALPR data to stalk his estranged wife. In Texas in 2025, deputies searched Flock's nationwide database ostensibly for a "missing person" case, when they were actually tracking a woman who had a self administered abortion. In Colorado in 2020, Brittney Gilliam was pulled over at gunpoint, and her daughter and nieces were handcuffed and forced face down on the pavement, because an ALPR misidentified her SUV as a stolen motorcycle with plates from a different state.

There is no federal law governing ALPR use. Eighteen states have some form of regulation. Legal challenges have produced split results. A Virginia federal judge dismissed a constitutional challenge to Norfolk's Flock system in January 2026. A Virginia state court judge reached the opposite conclusion in 2024, finding that ALPR data does constitute a Fourth Amendment search requiring a warrant. The EFF and ACLU filed suit against San Jose in November 2025. At least 30 localities have rejected or dropped Flock since 2025, including Denver, Austin, Cambridge, Mountain View, Eugene, and Olympia. Community resistance is growing.

What You Can Do About It Right Now

No single step will eliminate all location tracking. Your cell carrier always knows your approximate position when your phone is powered on, and you cannot turn off your license plate. Every step you take here reduces the digital breadcrumbs that feed the surveillance economy. Here are the highest impact actions you can take today.

Kill the Advertising ID on Your Phone

The single most important thing you can do right now is neutralize the advertising identifier on your phone. This is the master key that ties all of your app activity together into a single profile. On iPhone, go to Settings, then Privacy and Security, then Tracking, and toggle off "Allow Apps to Request to Track." This zeros out your IDFA, the unique identifier that data brokers use to build cross app dossiers on you. On Android version 12 or later, go to Settings, then Security and Privacy, then Privacy, then Ads, and select "Delete advertising ID." This permanently removes the identifier that companies like Gravy Analytics and Mobilewalla used to track a billion devices.

Tighten Your Location Permissions

Both iPhone and Android give you granular location controls that most people never touch. On iPhone, go to Settings, then Privacy and Security, then Location Services. For each app, choose "Never," "While Using the App," or "Ask Next Time." For any app that does not need turn by turn navigation, turn off Precise Location. This limits the app to knowing you are somewhere within a region of roughly 4 to 20 kilometers instead of your exact GPS coordinates. Android 12 and later offers the same precise versus approximate toggle. Set most apps to "While Using" or "Don't Allow." Then go to Settings, then Location, and turn off Wi-Fi scanning and Bluetooth scanning, both of which allow background positioning even when those radios appear to be off.

Understand the Five Layers of Location Tracking

Turning off location services on your phone does not stop all location tracking. There are five layers. GPS is the most precise at 3 to 50 meters, and it is blocked for apps when location services is off. This is the layer you control most directly. Cell tower triangulation is the second layer, accurate to 300 meters to a kilometer or more, and your carrier always knows where you are as long as the phone is on. Airplane mode or powering off are the only ways to stop it. Wi-Fi positioning is the third layer at 15 to 40 meters, and phones scan nearby networks even with Wi-Fi toggled off from the quick settings panel. You have to disable Wi-Fi scanning in your full settings menu. Bluetooth beacons are the fourth layer at 1 to 3 meters, used by retail stores and public spaces. IP geolocation is the fifth layer, which operates at city level and can be masked by a VPN.

On iPhone, Apple quietly collects a log of places you frequently visit, even with many app permissions disabled. Check Settings, then Privacy and Security, then Location Services, then System Services, then Significant Locations. Review this data and clear it. On Android, Google can figure out your location from search queries, Maps usage, and weather requests even with GPS off. Disable Location History and Web and App Activity in your Google Account settings.

Turn On the Global Privacy Control in Your Browser

The Global Privacy Control is a browser signal that automatically tells every website "Do not sell or share my personal data." It is legally enforceable under California's CCPA and CPRA. Sephora was fined $1.2 million in 2022 for ignoring it. Healthline paid $1.55 million in July 2025. As of January 2026, 12 state privacy laws require businesses to honor this signal. California's Opt Me Out Act, signed October 2025, will require all browsers to include built in GPC functionality by January 1, 2027. Brave and DuckDuckGo browsers have it on by default. Firefox users can enable it in Settings under Privacy and Security. Chrome users should install the Privacy Badger extension from the Electronic Frontier Foundation.

Use the California Delete Act

California's Delete Act created the Delete Request and Opt Out Platform, known as DROP, which opened to consumers on January 1, 2026 at privacy.ca.gov/drop. A single request tells all registered data brokers to delete your personal information. Brokers must begin processing by August 1, 2026, and they are required to check the platform every 45 days. Penalties for noncompliance run $200 per request per day. When you register, include your mobile advertising ID alongside your name, email, and address to maximize the effectiveness of the deletion.

If you live outside California or want immediate action, paid removal services like Incogni and DeleteMe automate the tedious process of sending opt out requests to hundreds of brokers. You can also request your LexisNexis consumer report, which may include your driving data if your car has been reporting on you. Check consumer.risk.lexisnexis.com and Verisk's portal at fcra.verisk.com.

The Five Minute Privacy Reset

For iPhone: First, disable tracking in Settings under Privacy and Security, then Tracking, and toggle off "Allow Apps to Request to Track." Second, audit your location settings and set most apps to "While Using" or "Never," and turn off Precise Location for apps that do not need navigation. Third, disable Apple's ad targeting in Settings under Privacy and Security, then Apple Advertising, and toggle off "Personalized Ads." Fourth, enable Private Relay if you subscribe to iCloud Plus. Fifth, install Brave or Firefox as your browser.

For Android: First, delete the advertising ID in Settings under Privacy, then Ads, then "Delete advertising ID." Second, audit location permissions and change most apps to "While Using" or "Don't Allow." Third, disable background scanning by going to Settings, then Location, and turning off Wi-Fi scanning and Bluetooth scanning. Fourth, disable Google tracking in your Google Account under Data and Privacy by turning off Location History and Web and App Activity. Fifth, install Brave or Firefox.

On all phones: Delete apps you have not used in the past month. Set up NextDNS, which has a free tier and takes about five minutes, for network level ad and tracker blocking. Enable the Global Privacy Control in your browser. Register at California's DROP portal if you are eligible.

This Is Not a Future Problem. It Is Happening Right Now.

The surveillance system described in this chapter is not a theoretical risk. It is a fully operational, commercially mature industry that tracks billions of devices every day and generates tens of billions of dollars in revenue. The Gravy Analytics breach proved that the worst case scenario is real. A single company's careless security exposed the movements of millions of people across the most sensitive places in their lives, from the White House to addiction treatment centers to houses of worship.

The FTC has taken five enforcement actions against location data brokers since 2022. Those actions matter. They are also a handful of orders against an industry that broadcasts your location 294 billion times a day. The architecture that makes all of this possible, the system in which your phone sends your coordinates to thousands of companies every time an ad loads, remains fully operational.

What sets this moment apart from every previous era of surveillance is scale and permanence. The 50 billion location pings obtained by the New York Times were from nearly a decade ago. Today, Flock Safety's 80,000 cameras scan 20 billion plates every month. Your connected car may report your position every three seconds. All of this data gets stored, aggregated, and resold. And as the Gravy breach showed, it is all vulnerable to theft.

The protective steps in this chapter are real and meaningful. Take them today. Share them with your family. Teach them to your children. Every setting you change and every identifier you delete shrinks the digital dossier that this industry has built about your life.

But I want to leave you with a bigger question. A society that generates 294 billion surveillance broadcasts a day has already made a choice about privacy. The question is whether any of us actually got to vote on it. I wrote this book because I believe the answer is no, and I believe it is time for that to change. The steps in this chapter protect you. The chapters that follow will show you how to protect your country.

Chapter 04: The House You Live In Is a 24/7 Recording Studio

Right now, as you read this, there is a good chance that something inside your house is listening to you. And watching you. And collecting information about what you do, when you do it, who you do it with, and how often you do it. It is sending all of that information to a corporation you have never met, where employees you will never know can access it, review it, and sell it.

I need you to sit with that for a second.

The smart speaker on your kitchen counter. The doorbell camera on your front porch. The television mounted on your living room wall. The thermostat in your hallway. These devices were sold to you as tools of convenience. Ask for a recipe. See who is at the door. Adjust the temperature from bed. Simple. Helpful. Modern. And every single one of them is generating a continuous stream of personal data about you and your family that feeds a multibillion dollar surveillance economy you never agreed to join.

Here is what almost nobody tells you. Nearly half of American homes now have at least one of these devices. And the companies that make them have been caught, fined, sued, and sanctioned for secretly recording people in their own homes, for letting employees watch private camera feeds of women in their bedrooms and bathrooms, for retaining voice recordings of more than 800,000 children even after their parents demanded deletion, and for building tracking technology into your television that captures what is on your screen twice every second and sells that information to advertisers. The combined fines, settlements, and legal judgments against these companies have reached hundreds of millions of dollars. And those numbers keep climbing.

This chapter is going to walk you through exactly what is happening inside your connected home. You are going to learn which devices collect the most data about you, what those companies do with that data once they have it, what happened when they got caught, and most importantly, what you can do right now to take back control of the most private space in your life.

Every Device in Your Home Is Collecting More Than You Think

Let me start with the device that sits in more American homes than any other smart product. Amazon Alexa.

A 2024 privacy study found that Alexa collects 28 out of 32 possible categories of personal data, making it the most data hungry smart home device on the market. That is more than three times the average. Those categories include every voice recording you make after the wake word, full text transcripts, your location, your shopping habits, your payment information, how you use every connected device in your home, and even ambient sounds picked up by features that listen for things like glass breaking and smoke alarms.

Until September 2019, Amazon stored all of those recordings forever. Today, the default setting is still "save until I delete them," which means most people never delete them. And here is something that should concern every parent, grandparent, and caregiver reading this. The FTC discovered that Amazon gave approximately 30,000 employees access to users' voice recordings. Thirty thousand people could listen to what you said in your living room.

In March 2025, Amazon made things worse. The company eliminated the option to process voice commands locally on your Echo device. Every single thing you say to Alexa now travels through Amazon's cloud servers.

The end result is that you lost the ability to keep your voice inside your own house.

Google Home and Nest devices collect 22 out of 32 data categories. Your Nest thermostat knows when you wake up, when you leave, when you come home, and when you go to sleep, because it tracks occupancy patterns, ambient light, and even syncs with your phone's location. Nest cameras collect video and audio, run facial recognition through a feature called Familiar Faces, and log every person and object that passes through the frame. Google says it does not sell your personal data to third party marketers. Your data still flows through Google's own massive advertising system.

Apple takes a different approach with its HomePod and HomeKit products. Siri requests get associated with a random identifier instead of your Apple ID. Audio recordings are not stored by default. Many commands get processed right on the device without ever leaving your home. And Apple's HomeKit Secure Video analyzes camera footage locally on your home hub, encrypts clips with military grade encryption before uploading to iCloud, and uses end to end encryption that Apple itself cannot break. After the 2019 listening scandal, Apple made its human review program something you have to choose to participate in, and limited reviewers to Apple employees only.

Ring doorbells sit on the front porches of more than 10 million American homes. They collect HD video, audio, motion data, and facial recognition information. The FTC found that before September 2017, Ring employees could access every single customer video with zero restrictions. One employee spent months watching thousands of private video recordings from female users' cameras positioned in bathrooms and bedrooms. Think about that. The camera you bought to protect your home was used by a company employee to spy on women in the most intimate spaces of their lives.

Wyze cameras, which are popular because they cost so little, have experienced a cascade of security disasters. A 2019 data leak exposed 2.4 million customers. A vulnerability that lasted from 2019 to 2022 allowed anyone to bypass the login screen and access camera feeds. And in February 2024, about 13,000 Wyze users opened their apps and saw thumbnail images from other people's cameras due to a system error. Strangers were looking into other families' homes.

Even your thermostat plays a part. Ecobee smart thermostats use built in radar sensors that detect which rooms in your house are occupied, with up to 32 sensors covering a 60 foot range per sensor. To their credit, Ecobee reportedly refused when Amazon demanded access to passive device data from Alexa enabled thermostats when customers were not actively using the voice assistant. That refusal is the exception, not the rule.

The Super Bowl Ad That Showed America What Surveillance Really Looks Like

On February 8, 2026, more than 120 million Americans watched Ring air a Super Bowl commercial for its new Search Party feature. The ad told the heartwarming story of a lost yellow Labrador named Milo. A family uploaded Milo's photo to the Ring app, blue circles bloomed across a neighborhood map as interconnected Ring cameras activated one after another, and a tracking box locked onto the dog with a green confirmation.

Tearful reunion. Amazon announced a million dollar donation to animal shelters. It was, on its face, a puppy rescue story.

And it terrified people.

Within hours, the backlash went viral. Matt Nelson of WeRateDogs, who has roughly 20 million followers, posted a video calling the ad an attempt to manufacture consent for mass surveillance. The Electronic Frontier Foundation said Americans should feel unsettled about the potential loss of privacy, pointing out that Ring already runs facial recognition through its Familiar Faces feature and that applying that same technology to scan an entire neighborhood of cameras was the obvious next step. The ACLU published an analysis that asked a question every American should think about. Today this technology searches for puppies. Tomorrow it could search for immigrants. Or people wearing political t-shirts. Or anyone a government agency decides to target.

Senator Ed Markey wrote directly to Amazon's CEO urging the company to stop using facial recognition technology in Ring products. And the controversy deepened when reporters revealed Ring's recent partnerships with Flock Safety, a company that makes automated license plate readers used by law enforcement, and Axon Enterprises, which makes body cameras and Tasers. These partnerships had relaunched police access to Ring footage through a feature called Community Requests. Investigative journalists found that local police departments were already searching Flock's camera network on behalf of federal immigration agencies, giving ICE and CBP what amounted to a side door into neighborhood surveillance networks.

An internal email from Ring's founder revealed he planned to expand Search Party far beyond lost pets. He said the technology would help, in his words, zero out crime in neighborhoods. The ACLU noted that achieving that goal would require extraordinarily pervasive surveillance of American life.

Four days after the Super Bowl, Ring canceled its Flock Safety partnership. The company said it was about resources, not public pressure. The cancellation did nothing to change Ring's facial recognition features, the Community Requests tool, or the Search Party feature itself. The infrastructure remains in place.

They Were All Listening and They All Got Caught

In the summer of 2019, the biggest names in technology got exposed for doing exactly what millions of Americans feared most. They were listening to private conversations inside people's homes.

On April 10, 2019, a major business news outlet reported that thousands of Amazon workers around the world, in Boston, Costa Rica, India, and Romania, were listening to, transcribing, and annotating Alexa voice recordings. Each worker processed up to 1,000 audio clips during a nine hour shift. They heard people asking for recipes and setting timers. They also heard a woman singing in the shower, a child screaming for help, and what two reviewers believed was a sexual assault. The recordings were linked to account numbers, device serial numbers, and each user's first name. A follow up report two weeks later revealed the auditing team could also access users' home addresses.

Amazon's response followed a predictable pattern. Acknowledge. Minimize. Add an opt out option buried deep in the settings menu. The company created a way to disable human review inside the Alexa app. They

made it opt out, meaning you had to go find it and turn it off. It was on by default. Amazon's devices chief admitted the company considered making human review something you would have to opt into. They decided not to.

Three months later, a Belgian news outlet revealed that Google contractors were systematically listening to Google Assistant recordings. A whistleblower shared more than 1,000 recordings, and 153 of those recordings were never intentionally triggered by the user. Google had claimed the recordings were anonymous. Reporters proved that was not true. They were able to identify specific users and locate their homes from spoken addresses captured in the recordings. Google suspended human review and later brought it back on an opt in basis.

Two weeks after the Google story, a major newspaper revealed that Apple contractors were hearing confidential medical conversations, drug deals, and recordings of couples having sex through a Siri quality program Apple called grading. Apple suspended the program, issued an apology, and changed the review process to require users to opt in. They also limited the program to Apple employees.

The legal fallout was enormous. In May 2023, the FTC brought two enforcement actions against Amazon on the same day. The first hit Amazon with a 25 million dollar civil penalty for violating children's privacy law. Amazon had kept voice recordings from more than 800,000 children indefinitely and used them to train its algorithms, even when parents specifically asked for those recordings to be deleted. The second action resulted in a 5.8 million dollar settlement against Ring for failing to prevent employees from watching customer video feeds and for security failures that allowed hackers to access roughly 55,000 user accounts. In 2024, the FTC sent refund checks to Ring customers affected by those failures.

In the class action arena, the numbers are staggering. Apple agreed to a 95 million dollar settlement over Siri's recording practices, with payments starting in January 2026. Google proposed a 68 million dollar settlement for its Assistant recording practices, filed in January 2026. Amazon faces an active nationwide class action in which a federal judge certified a class of Alexa users in July 2025, making it one of the most expansive privacy cases involving voice devices in American legal history.

Your Television Is Watching You

Your smart TV does not just show you content. It watches what you watch. And it does this through a technology called Automatic Content Recognition, or ACR. Think of it as a system that works like a music recognition app that never turns off. Your television periodically captures small pixel or audio samples from whatever is on your screen, converts those samples into compact digital fingerprints, sends them to the manufacturer's servers, and matches them against a library of known content. Your Samsung TV does this approximately every 60 seconds. Your LG TV does it every 15 seconds. A 2024 university study confirmed that ACR stays active even when you are watching content through an external device connected by HDMI. That means your TV is tracking what you watch on your gaming console, your streaming stick, and your laptop.

The most aggressive offender was Vizio. Starting in February 2014, Vizio installed tracking software on all new TVs and then remotely downloaded that same software onto TVs people had already purchased. The system captured what was on your screen on a second by second basis, collecting more than 100 billion data points per day across 11 million televisions. Vizio combined your viewing data with your IP address,

demographic information including your age, sex, income, marital status, and education level, and sold all of it to third parties for advertising and targeting. ACR was turned on by default. A pop up notification about the change timed out after just one minute and never mentioned data collection.

In February 2017, the FTC and the New Jersey Attorney General hit Vizio with a 2.2 million dollar settlement. A follow up class action settlement reached 17 million dollars, covering about 16 million televisions. And here is the part that tells you everything you need to know about the economics of smart home surveillance. By the third quarter of 2021, Vizio's advertising and data division generated 57.3 million dollars in gross profit. That was more than double the 25.6 million dollars Vizio earned from selling the televisions themselves. The TV was never the product. You were the product. The TV was just the delivery mechanism.

In December 2025, Texas Attorney General Ken Paxton filed suit against five TV manufacturers, Samsung, Sony, LG, Hisense, and TCL, under the Texas Deceptive Trade Practices Act. The complaint alleged that ACR captures screenshots every 500 milliseconds, monitors all content sources including anything connected through HDMI, and that companies bury their disclosures behind more than 200 clicks spread across four or more menus. Against the Chinese manufacturers Hisense and TCL, Paxton raised national security concerns, noting that Chinese law can compel companies to share data with the Chinese government. A Texas court issued the first ever temporary restraining order against a TV maker for ACR, blocking Hisense from collecting data from Texas consumers in December 2025. Samsung reached a settlement in February 2026, agreeing to stop all ACR collection without express consent.

Your Home Has a Fingerprint and Advertisers Already Know It

Here is something that most people have never heard of, and it matters enormously. The collection of smart devices in your home creates a unique digital fingerprint that identifies your household, even without your name or address attached. Researchers have shown that internet service providers can identify specific homes just by analyzing the smart devices connected to their network. Because every device in your home transmits from the same IP address, that address becomes a cross device identifier that the advertising industry uses to target your household as a single unit.

The privacy risks go deeper than advertising. Academic researchers built a system that can listen to the encrypted network traffic coming from smart home devices and classify specific events, like which lights turn on, when doors open, and whether anyone is home, with 89 percent accuracy. They did this without cracking the encryption. Another study from 2025 showed that someone in an adjacent apartment could learn personal details about a household using three inexpensive antennas placed along a shared wall, detecting when devices change state and mapping daily routines.

The advertising industry has built massive infrastructure around this concept. Companies maintain what they call device graphs that stitch together identifiers from your web browsing, your apps, your connected TV, and your smart home devices into clusters that belong to the same household. One major media company claims its identification system covers 110 million households and 260 million device IDs. When you combine your TV's viewing data, your smart speaker's voice recordings, your thermostat's occupancy data, and your camera's visitor logs, all linked through a shared IP address, you get a portrait of your daily life that is more detailed than anything a private investigator could assemble in a month of surveillance.

Hackers Can Take Your Home Hostage

In January 2026, Experian published its annual fraud forecast and named smart home hijacking as one of the top five threats Americans will face this year. The warning was blunt. Devices like virtual assistants, smart locks, security systems, and smart appliances will be targeted by criminals to access personal data, monitor household activity, and take control of physical access points. Experian specifically predicted new forms of ransomware that could lock people out of their own homes and create opportunities for account hijacking on a massive scale.

This is not speculation. It has already started happening.

In December 2019, a hacker broke into a Ring camera in an eight year old girl's bedroom in Mississippi and told her he was Santa Claus. In another incident, a hacker used a Ring device to demand a ransom of 50 bitcoin, worth about 400,000 dollars. An Illinois couple watched helplessly as a hacker took over their Nest camera, screamed profanities at their child through the speaker, and cranked their thermostat up to 90 degrees at the same time. A podcast called NulledCast featured hosts breaking into Ring and Nest devices live on air for entertainment.

At a major hacking conference in 2016, a researcher demonstrated the first smart thermostat ransomware, locking the device at 99 degrees, displaying a payment demand on the screen, and changing the unlock PIN every 30 seconds. The researcher completed the entire attack in two evenings. Security researchers who tested 16 popular Bluetooth smart locks found that 75 percent had serious vulnerabilities, including locks that transmitted passwords as plain text. One researcher was able to change a lock's administrator password remotely, permanently locking the actual owner out of their own home.

The numbers tell a sobering story. Global IoT attacks now reach 820,000 per day. More than half of all smart home devices have critical security weaknesses. One in every three data breaches involves a connected device. And a 2024 consumer investigation found that cheap smart doorbells sold on major retail websites had security flaws so severe that anyone could access photos and footage from the cameras, and most of those doorbells lacked required federal identification numbers, making them effectively illegal products.

You Can Take Your Home Back

The single most important decision you can make for your family's privacy right now is choosing devices that keep your data inside your house. When a device processes your voice commands, your video footage, and your sensor data locally, on the device itself, that information never leaves your home network. When a device sends everything to the cloud, your data ends up on corporate servers where companies can store it forever, analyze it, sell it to advertisers, let employees review it, or hand it to law enforcement, sometimes without a warrant.

Apple currently offers the strongest mainstream privacy protections. HomeKit Secure Video analyzes your camera footage locally on a HomePod or Apple TV hub, identifies people, pets, vehicles, and packages using on device artificial intelligence, and encrypts everything before sending it to iCloud with end to end encryption. Since 2021, Siri speech processing happens on your device. Apple's newer Private Cloud Compute system makes sure that any data that does reach the cloud is never stored and can be independently audited.

For the highest level of privacy, an open source platform called Home Assistant runs entirely on a small piece of hardware you keep in your home, a dedicated hub that costs about $129. It supports more than 1,000 device integrations, processes all your automations locally, and sends zero data to any company. This is significant because multiple smart home cloud services have shut down over the years, turning purchased hardware into paperweights overnight. When your system runs locally, no company decision can take it away from you.

If you want security cameras that do not feed a corporate cloud, you have real options. Reolink cameras store footage locally on microSD cards or network video recorders with zero recurring costs. Lorex is the only major camera brand that does not even offer a cloud storage plan, shipping a microSD card with every camera. Frigate is an open source video management system that runs on your own hardware with local AI detection. Over their lifetime, local storage camera systems cost about 40 percent less than cloud subscription alternatives, which can add up to more than 1,000 dollars over a decade.

A new connectivity standard called Matter, backed by Apple, Google, Amazon, and Samsung, uses local communication with mandatory encryption and has reached version 1.5 with camera support. More than 750 certified products exist as of early 2026. The companion Thread networking protocol keeps device to device communication local. Look for the Matter logo when you shop for new smart home products. It is not a complete privacy solution because manufacturers can still route data through their cloud servers. It is a meaningful step forward.

The Law Is Catching Up, State by State

The United States has no federal privacy law that specifically protects your smart home. A bipartisan bill called the American Privacy Rights Act was introduced in April 2024 with provisions for data minimization and a private right of action. It died without a vote. Experts estimate that full federal privacy legislation remains at least two to three years away.

So the fight is happening in the states.

As of January 2026, more than 20 states have enacted broad consumer privacy laws, covering close to half of all Americans. California's updated privacy regulations took effect on January 1, 2026, requiring cybersecurity audits, risk assessments, and disclosures about automated decision making technology. Illinois' biometric privacy law continues to drive voice data litigation. A federal judge ruled in 2023 that even someone who spoke to an Alexa device just once, on someone else's device, could join a class action over voiceprint collection. Florida brought the first enforcement action under its new Digital Bill of Rights against a streaming device maker in October 2025. New York's Attorney General settled with a camera company for security failures that left video streams accessible to anyone. Nebraska's Attorney General sued a home security company for selling cameras manufactured by companies identified as national security risks.

Texas has been especially aggressive. The Texas Attorney General secured a 1.4 billion dollar settlement from one social media company in 2024 for biometric violations and a 1.375 billion dollar settlement from a major search engine company in 2025 for data tracking, in addition to the December 2025 lawsuits against five television manufacturers. The December 2025 suits, and the temporary restraining order against Hisense, represent the first time a state has moved to actively block a TV manufacturer from collecting ACR data.

The FTC remains the primary federal enforcer. Beyond the Amazon fines, the commission finalized updated children's privacy rules in January 2025, published a report in September 2024 finding that major tech companies engaged in vast surveillance of consumers, and issued a November 2024 report revealing that 89 percent of connected device manufacturers fail to tell you how long they will provide software updates. The FCC created a cybersecurity label for smart home products in March 2024 called the U.S. Cyber Trust Mark. As of early 2026, no products carry the label because the program's lead administrator withdrew in December 2025.

What You Need to Know Right Now

Almost 70 million American households now use smart home technology. Smart speakers sit in about 35 percent of homes. Smart TVs are in 82 percent of TV households. Video doorbells cover roughly a third of all homes. Smart security cameras are in 38 percent of households. More than 300 million devices have been connected to Alexa alone. And ACR is built into virtually every smart TV on the market, turned on by default, with vendors reporting that more than 90 percent of users never turn it off.

Here is what those numbers tell us about awareness. Fifty seven percent of Americans say they are concerned about smart home data collection. At the same time, 52 percent of smart thermostat owners have no idea how their data is collected. Only 14 percent of buyers researched a manufacturer's privacy policy before purchasing. Manufacturers know this. They design their products to collect the maximum amount of data by default and bury the settings that would let you change that behind menus, submenus, and confusing language.

This is what I need you to understand. Your concern about privacy is valid. Your instinct that something is wrong is correct. The companies making these devices have been caught over and over again doing exactly what you feared they were doing. They recorded your children. They let employees watch your private video. They tracked your screen twice per second and sold the data. They eliminated your option to keep your voice inside your own home. And when they got caught, the fines they paid were a fraction of the profits they earned from doing it.

You deserve a home where the walls do not have ears. Where your television watches nothing. Where your thermostat does not report your sleep schedule to an advertising company. Where your front door camera belongs to you and only you. That home is possible. It requires you to make different choices about what you bring into your house, which companies you trust with your most private moments, and how you configure the devices you already own.

Go into your Alexa app today and review your privacy settings. Disable the Help Improve Alexa program. Delete your stored voice recordings. Check your smart TV settings and turn off ACR. Look for the words "viewing data," "content recognition," or "recommendations" in your TV's privacy menu and disable them. Next time you shop for a smart home device, look for products that keep data local. Consider whether the convenience of any device is worth what you give up for it.

The companies selling these devices are not going to protect your privacy for you. They have proven that, repeatedly, under oath, in court, and in regulatory filings. The law is catching up, slowly, state by state. And until a federal privacy law exists, the most reliable protection you and your family have is your own informed decision making.

Your home should be the one place in the world where you are not being watched, not being recorded, and not being sold. You have the right to make it that way. And now you know how.

Chapter 05: That Car in Your Driveway Knows More About You Than Your Spouse Does.

Here is something most people never think about. The car sitting in your driveway right now knows where you slept last night. It knows what time you left this morning, how fast you drove, how hard you hit the brakes, and whether you were wearing your seatbelt. It knows the phone numbers in your contacts, the last address you plugged into navigation, and possibly the words you spoke out loud inside the cabin. And it has been sending that information, over a hidden cellular connection you never asked for, to servers controlled by your automaker, who then sells it to insurance companies, data brokers, and sometimes law enforcement agencies that never obtained a warrant.

If that sounds extreme, stay with me. Because what has unfolded over the past two years around connected vehicle surveillance is one of the most disturbing privacy stories in American life. It involves secret data sharing, rigged consent forms, insurance premiums spiking because of driving data you never agreed to hand over, domestic violence survivors being tracked through their own cars, and a massive federal enforcement action that confirmed what privacy advocates had been warning about for years. Your car is no longer just a car. It is a data collection machine, and the people profiting from it are counting on you never finding out.

Your Car Knows More About You Than Your Phone Does

A modern connected vehicle generates an estimated 25 gigabytes of data every single hour. Let that sink in for a moment. That is more data than most people use on their phones in an entire month, and it is streaming off your car in real time. Much of it stays onboard in the vehicle's computer systems. A growing portion of it gets transmitted through a factory installed cellular modem to cloud servers run by your automaker, sometimes as often as every three seconds, according to the Federal Trade Commission's case against General Motors.

The categories of data are staggering. Your car records your precise GPS location, your speed, your acceleration patterns, every hard braking event, your seatbelt status, your steering behavior, and complete trip logs with start and end points. Some vehicles now include driver monitoring cameras that use facial recognition and eye tracking, voice command systems that process your words through cloud servers, and always on microphones that capture audio inside the cabin. When you pair your phone over Bluetooth, many vehicles pull in your contacts, call logs, and text messages. Your navigation history, what stations you listen to, what apps you use on the infotainment screen, all of it gets logged.

And the privacy policies that automakers bury in their paperwork go further than any reasonable person would expect. Nissan's privacy policy admits to collecting data about sexual activity, health diagnoses, genetic information, and what it calls psychological trends. Nissan also claims the right to develop inferences about a driver's intelligence, abilities, and aptitudes, and to sell those assessments to third parties. Kia's policy similarly allows the collection of information about sex life. Tesla received every possible privacy warning from the Mozilla Foundation, and its cabin camera can record and transmit video. Toyota maintains twelve separate privacy policy documents that researchers described as near incomprehensible. Subaru's policy states that if you are simply a passenger in the vehicle, you are considered a user who has consented to data collection.

In September 2023, the Mozilla Foundation published the results of 600 hours of research across all 25 major car brands sold in America. The report was called Privacy Not Included, and its conclusion was devastating.

Every single brand received a failing grade. Mozilla called cars the worst product category they had ever reviewed for privacy. Eighty four percent of the brands share or sell your data. Ninety two percent give drivers little to no control over their personal information. Fifty six percent will hand your data to law enforcement based on nothing more than an informal request, no court order required. Not one single brand could confirm that it encrypts all of the personal data stored on its vehicles. Senator Edward Markey sent letters to fourteen automakers demanding answers. Their responses, published in early 2024, were described as vague, incomplete, and evasive.

Why do automakers collect all this information? Because selling it is enormously profitable. The car data market is projected to reach somewhere between 450 billion and 750 billion dollars by 2030. When that kind of money is on the table, protecting your privacy is not the priority. Monetizing your life is.

The Pipeline From Your Dashboard to Your Insurance Premium

The most important privacy revelation of 2024 was that automakers had been quietly funneling detailed driving data to insurance companies through data brokers, and that drivers' premiums were going up as a direct result. The story broke in March 2024 when a reporter at the New York Times published an investigation showing that GM's OnStar Smart Driver program had shared data from as many as eight million vehicles with LexisNexis Risk Solutions and Verisk Analytics, two massive consumer reporting agencies that serve the insurance industry.

The way it worked was deceptively simple. Smart Driver was marketed as a free feature to help people become better drivers. It tracked habits through the MyChevrolet, MyBuick, MyGMC, and MyCadillac apps. What drivers did not realize was that every hard brake, every rapid acceleration, every instance of driving above 80 miles per hour, and every late night trip was being recorded and transmitted to LexisNexis and Verisk. Those companies converted the raw data into risk scores and something called Driving Behavior Data History Reports, which insurance companies purchased to set your rates. The FTC found that GM collected this data as often as every three seconds.

The consequences hit real people immediately. Kenn Dahl, a 65 year old software company owner in Seattle with a perfectly clean driving record, watched his insurance spike 21 percent in 2022 after leasing a Chevrolet Bolt. He had no idea why. When he requested his LexisNexis report under the Fair Credit Reporting Act, he received a 258 page document that detailed 640 trips over six months. Every instance of hard braking and rapid acceleration had been logged. He told the Times it felt like a betrayal. A Cadillac driver in Palm Beach County, Florida was denied insurance by seven separate companies after his LexisNexis report revealed months of driving behavior data he never knew was being collected. Another consumer, Temeika Clay, saw her premium jump 80 percent after GM shared 603 entries of driving data from her Chevy Camaro.

The reporter who broke the story later discovered that she and her husband had been enrolled in Smart Driver despite never opting in. GM called it a bug. Her deeper investigation revealed that dealership salespeople were auto filling consent forms with yes entries because their pay was docked if they failed to enroll buyers. GM graded its dealerships on enrollment percentages. So the system was designed from the ground up to ensure that as many people as possible were enrolled, whether they knew about it or not.

On January 14, 2026, the FTC finalized a 20 year consent order against GM and OnStar. This was the agency's first enforcement action targeting connected vehicle data. The order imposes a five year ban on

sharing geolocation and driver behavior data with consumer reporting agencies. It requires affirmative express consent before any data collection. It gives consumers the right to access, delete, and disable collection of their data. And it requires GM to destroy all previously retained information. No monetary fine was imposed, though GM's 2025 annual report disclosed 500 million dollars in accumulated litigation costs from the scandal. State attorneys general have piled on. Texas sued GM over data collected from 1.8 million Texas drivers, seeking up to 10,000 dollars per violation, which could mean as much as 18 billion dollars in penalties. Nebraska, Arkansas, Indiana, and Iowa have filed their own lawsuits. Arkansas alone is seeking up to one billion dollars in damages.

The Allstate Shadow Network That Tracked 45 Million Americans

At the same time GM was being exposed, a parallel surveillance operation came to light. In January 2025, Texas Attorney General Ken Paxton filed the first ever enforcement action under a comprehensive state data privacy law, targeting Allstate and its subsidiary Arity. The complaint tells a remarkable story. Arity developed a piece of tracking software and paid app developers millions of dollars to embed it in popular consumer apps including Life360, GasBuddy, and Fuel Rewards. That software operated silently in the background, capturing your location, accelerometer data, speed, and trip details every 15 seconds or less from 40 million active connections. Arity marketed what it built as the world's largest driving behavior database, covering over 45 million Americans.

The software could not even tell whether the phone user was actually driving. A passenger on a bus or a rider on a rollercoaster could have their data recorded as bad driving and used against them for insurance pricing. Arity also purchased driving data directly from automakers including Toyota, Lexus, Mazda, and the Stellantis family of brands. On March 3, 2026, a federal judge in Chicago allowed wiretapping, Fair Credit Reporting Act, and privacy tort claims against Allstate and Arity to go forward.

How to Check What Is in Your File

Two companies controlled the pipeline between your car and your insurer. LexisNexis Risk Solutions operates its Telematics Exchange, working with automakers that represent 46 percent of new car sales in America and 95 of the top 100 auto insurers. Under the Fair Credit Reporting Act, you have the right to request one free report each year. You can do this at consumer.risk.lexisnexis.com or by calling 1 888 497 0011. Verisk Analytics operated a parallel system containing over 260 billion miles of driving data from 8 million vehicles. After the GM scandal broke, Verisk shut down its automaker sourced driver behavior data product entirely in mid 2024. You can still check your Verisk records at fcra.verisk.com or by calling 1 800 627 3487. I want to encourage you to request both reports. You may be shocked by what you find.

Police Are Trained to Get Your Car's Data

Here is something that may or may not surprise you. Law enforcement officers are being systematically trained to extract data from connected vehicles. A 2025 investigation based on California Highway Patrol training documents obtained through public records requests revealed that police receive detailed instruction on how to acquire telematics data based on a vehicle's year, make, and model. The training materials even note that GM's OnStar transmits location data twice as frequently as Ford's system.

A landmark congressional investigation in 2024 asked fourteen automakers about their law enforcement practices. The findings should concern every American. Only five automakers, GM, Ford, Honda, Stellantis, and Tesla, require a warrant before handing over your location data to police. Eight manufacturers, including Toyota, Nissan, BMW, Mercedes Benz, and Volkswagen, admitted they would surrender your location data in response to a mere subpoena, meaning no judge ever has to sign off. Only Tesla notifies vehicle owners when the government requests their data. None of the automakers publish transparency reports comparable to what tech companies like Google and Apple release.

The physical extraction of data from vehicles has become routine. A company called Berla Corporation manufactures forensic tools that can pull call history, contacts, text messages, GPS history, and social media data from infotainment systems. In one documented case, a Berla analyst extracted data from 70 different phones that had been connected to a single Ford Explorer rental car over time. Michigan State Police reported extracting vehicle data for everyday felonies two to three times per week. Police in Oakland, California sought warrants in 2024 specifically to tow Tesla vehicles into evidence and access their Sentry Mode camera footage for homicide investigations.

There is another pathway that is even more troubling. After the Supreme Court's 2018 Carpenter v. United States decision required police to get a warrant before accessing historical cell site location data, federal agencies found a workaround. They started buying equivalent location data from commercial data brokers, no warrant needed. The FBI, IRS, DHS, ICE, and Secret Service have all purchased commercial location data without any court authorization. A company called Fog Data Science sells a subscription search engine covering 15 billion data points daily from 250 million devices to at least 18 law enforcement agencies, and none of those agencies need a warrant. The Fourth Amendment Is Not For Sale Act passed the House in April 2024, which would have closed this loophole. It died in the Senate. As of March 2026, the loophole is still wide open.

The Supreme Court is now poised to address mass location surveillance more directly. On January 16, 2026, the Court agreed to hear Chatrie v. United States, the first geofence warrant case to reach the justices, with oral arguments scheduled for April 2026. The Fifth Circuit Court of Appeals had already declared that geofence warrants are categorically prohibited by the Fourth Amendment in United States v. Smith in 2024. The resulting disagreement among federal courts demands resolution, and whatever the Supreme Court decides will shape the future of digital privacy for decades.

When You Try to Opt Out, They Take Away Your Safety Features

Here is where the story becomes especially cynical. When drivers try to reclaim their privacy by opting out of data collection, automakers strip away safety features that you have already paid for. The bundling of safety critical functions with data hungry services is deliberate.

If you cancel GM's OnStar, you lose automatic crash response. That is the system that contacts emergency services when your airbags deploy. If you press the red SOS button after cancellation, you hear nothing but an automated recording telling you the subscription is inactive. OnStar routes through private call centers, not directly to 911, so the FCC's E911 rules for cell phones do not apply. Stellantis, the company behind Jeep, Ram, Dodge, and Chrysler, explicitly warns customers that canceling connected services for privacy reasons eliminates SOS Call, Automatic SOS Call, and Roadside Assistance. Tesla's privacy notice warns that opting

out of data collection may result in your vehicle suffering from reduced functionality, serious damage, or inoperability, partly because Tesla delivers recall fixes through over the air updates that require connectivity.

Navigation, remote start, and software updates all disappear when you disable data collection. GM vehicles without an OnStar subscription lose Google Maps, Google Assistant, and Google Play entirely. Multiple automakers withhold software updates from customers who opt out, meaning those vehicles may miss safety patches indefinitely. Think about what that means. You paid for a car. You own it. And the manufacturer is telling you that if you do not let them surveil you, they will take away features that protect your life on the road.

The Privacy4Cars organization published the first quantitative privacy scorecard for the auto industry in August 2025, evaluating 49 brands. The median score was just 1.7 out of 5.0. Only five brands scored above 3.0. Honda scored a dismal 0.8 before the California Privacy Protection Agency fined it 632,500 dollars in March 2025, the CPPA's first ever enforcement action against any company. After that fine, Honda overhauled its practices and now leads the industry with a score of 4.6. That tells you something important. Enforcement works. When regulators impose real consequences, companies change their behavior. Porsche offers a Private Mode that limits data transmission to only what the law requires while keeping navigation functional, one of the few genuinely respectful designs in the industry.

For those willing to go further, the only guaranteed way to stop data transmission is to physically disable the Telematics Control Unit by pulling the TCU fuse or removing the cellular modem. Online forum communities for GM, Ford, Subaru, and Toyota vehicles have documented the specific procedures. The consequences include dashboard warning lights, loss of the WiFi hotspot, and the elimination of all connected features. Under the Magnuson Moss Warranty Act, a physical modification like this should not void your entire vehicle warranty, though a manufacturer could deny claims specifically related to the modification.

Domestic Violence Survivors Are Being Tracked Through Their Own Cars

This is the part of the story that keeps me up at night. Connected car technology has given abusers a new and terrifyingly effective way to track and control their victims. A joint investigation published in April 2024 documented case after case of internet connected cars being used to locate domestic violence survivors after they flee.

Yenni Rivera, who works with domestic violence survivors in Los Angeles, told reporters that she hears the story over and over from survivors about being located by their vehicle and having it taken. Stephanie Davidson, managing attorney at the Legal Aid Foundation of Los Angeles, reported that multiple domestic violence shelters now prohibit survivors from parking their connected cars on site because the vehicles could reveal the shelter's confidential address.

The cases are devastating. Christine Dowdall of Louisiana discovered her ex husband was stalking her through the Mercedes Benz Mbrace app, showing up beside her at locations she had never told anyone about. Because the vehicle title was in his name, Mercedes told her there was nothing they could do. She paid a mechanic 400 dollars to disable the software. In another case, a woman parked her Lexus at a confidential survivor shelter behind a gated garage, and her court advocate still spotted the abuser nearby. The car had given away her location.

An FCC report documented a man who used remote vehicle access to activate his wife's car lights and horn in the middle of the night and run the heat on hot days as a form of harassment. In the most dangerous case on record, a woman's vehicle was repeatedly started at night inside her closed garage, causing dangerous carbon monoxide to build up inside the home.

California responded with SB 1394 in 2024, requiring automakers to sever an abuser's remote vehicle access within two business days when a survivor requests it. At the federal level, the Safe Vehicle Access for Survivors Act was introduced in March 2025. These are steps in the right direction. They are nowhere near enough.

Consumers Are Fighting Back in Court

Americans are not taking this lying down. The first class action lawsuit was filed on March 13, 2024, when Romeo Chicco of Florida sued GM, OnStar, and LexisNexis, claiming his Cadillac's data was shared without his consent and that the resulting LexisNexis report caused seven insurance companies to deny him coverage. By November 2024, 32 separate lawsuits had been consolidated into a single multidistrict litigation called In re Consumer Vehicle Driving Data Tracking Litigation, assigned to Judge Thomas W. Thrash Jr. in the Northern District of Georgia. A Master Consolidated Class Action Complaint was filed in December 2024, with fact discovery expected to continue through late 2026. The legal claims span Fair Credit Reporting Act violations, state consumer protection statutes, invasion of privacy, and breach of contract.

The litigation goes well beyond GM. A Toyota class action was filed in Texas in April 2025 after a plaintiff discovered that Progressive Insurance already had his driving data from a 2021 RAV4, even though he had never enrolled in any monitoring program. Toyota's affiliate had partnered with Arity to share the information. Mass arbitrations are being pursued against Hyundai, Kia, and Mitsubishi for similar practices.

Security researchers have exposed how fragile these systems are. In November 2024, a researcher named Sam Curry discovered that Subaru's STARLINK admin portal could be accessed using nothing more than a person's last name and ZIP code. That access allowed anyone to remotely start, stop, lock, or unlock any Subaru, pull up a full year of location history accurate to five meters, and add themselves as an authorized user without the owner ever receiving a notification. He demonstrated the vulnerability by tracking his own mother's 2023 Impreza to the doctor's office, friends' homes, and church. In September 2024, Curry found that Kia's dealer portal allowed remote control of over 600 Kia models manufactured after 2013 using nothing more than a license plate number.

76,000 License Plate Cameras Are Watching Every Road

Everything discussed so far involves data flowing through automaker systems. A separate surveillance network runs in parallel, and most people do not know it exists. Automated License Plate Readers, known as ALPRs, are high speed cameras mounted on poles, overpasses, police cruisers, and even private vehicles. They capture plate images at rates of up to 1,800 per minute, recording the plate number, date, time, GPS coordinates, a photograph of the vehicle, and in newer systems, distinguishing features like bumper stickers, dents, and roof racks. Over 76,000 of these cameras have been mapped across the United States.

The dominant player is Flock Safety, founded in 2017 and now operating in over 5,000 communities across 49 states with contracts covering more than 4,800 law enforcement agencies. Flock performs over 20 billion

vehicle scans each month and crossed 300 million dollars in annual revenue in 2025, earning a 7.5 billion dollar valuation. Its competitor, Vigilant Solutions, was acquired by Motorola Solutions for 445 million dollars and maintains a database of over 5 billion license plate detections with 150 million new ones added every month.

The private sector operates its own ALPR network with almost no oversight. Digital Recognition Network, a Motorola affiliated company, maintains over 20 billion confirmed vehicle sightings collected by a network of more than 600 repossession agents who mount cameras on their vehicles and scan plates all day long. Anyone, not just law enforcement, can search that database for about 20 dollars per query. No warrant is required because the Fourth Amendment does not apply to private companies.

An analysis of 63 California law enforcement agencies found that only 0.05 percent of license plate reader data was connected to any public safety interest at the time it was captured. That means 99.95 percent of scanned plates belonged to people who were not suspected of any crime. The California State Auditor found that Los Angeles alone had stored 320 million images, of which 99.9 percent were not matches to any wanted list.

Immigration and Customs Enforcement holds a 6.1 million dollar contract giving over 9,000 officers access to Vigilant's plate reader database. In 2025, journalists discovered that local police departments were conducting over 4,000 immigration related searches through Flock's system on ICE's behalf, despite Flock's stated policy banning immigration enforcement use. San Francisco police allowed 1.6 million unauthorized out of state searches of its Flock database, including at least 19 ICE related queries, in violation of California's sanctuary laws. In response, at least 30 cities have deactivated cameras or canceled Flock contracts since early 2025, including Santa Cruz, Cambridge, Eugene, Austin, and Mountain View.

California's SB 34, which took effect in 2016, requires ALPR operators to maintain public usage policies, put reasonable security measures in place, and keep access logs. It prohibits public agencies from sharing plate reader data with entities outside the state, a provision that 71 agencies were found to have violated. Civil remedies include minimum 2,500 dollar liquidated damages per violation plus potential punitive damages.

Who Owns the Data Your Car Creates? Nobody Knows

Despite everything you have just read, the fundamental question of who owns the data your car generates has no answer in American law. Three federal bills sit pending. The Auto Data Privacy and Autonomy Act would prohibit automakers from accessing or selling vehicle data without your explicit written consent. The REPAIR Act would require manufacturers to give owners and independent repair shops standardized access to vehicle data. The DRIVER Act would establish that vehicle data belongs to you and must be available in real time at no additional cost. None of these bills have advanced to a floor vote.

The right to repair movement has become the most visible battleground. Massachusetts voters approved a right to repair ballot measure in November 2020 by a 75 percent margin, requiring automakers to provide standardized telematics data access to vehicle owners and independent repair shops. The automaker alliance spent 26 million dollars trying to defeat it. When voters passed it anyway, the Alliance for Automotive Innovation immediately sued, claiming compliance was technically impossible and would create cybersecurity risks. A federal judge dismissed the industry's lawsuit in February 2025. The appeal was argued before the First Circuit in February 2026. In the meantime, Subaru responded by simply turning off its Starlink

telematics system for all Massachusetts customers rather than share the data. Maine voters passed a similar measure in 2023 and automakers sued there too.

The automotive industry has spent heavily to resist data transparency. Federal lobbying by the automotive industry reached 109 million dollars in 2025. GM alone spent 14.17 million dollars in a single year. The industry's position is that giving consumers access to their own vehicle data creates cybersecurity risks and conflicts with federal safety law. It is worth pausing to absorb the full picture here. These companies argue that sharing data with vehicle owners is too dangerous, while they sell that same data to insurance companies and data brokers for profit.

Europe has taken a different path. The EU Data Act, which became applicable on September 12, 2025, requires automakers to give vehicle owners control over the data their cars generate and to share that data with third parties under fair conditions, in a readable format, free of charge. The United States remains the only G20 nation without a comprehensive federal data privacy law.

For those of us in California, the CCPA and CPRA apply fully to connected vehicle data, giving us the right to know what is collected, to delete it, and to opt out of its sale. The CPPA's 632,500 dollar Honda fine proved that enforcement has real teeth. California is the only state that prohibits telematics data from being used for insurance rating. The CPPA launched its connected vehicle investigation in July 2023 and has signaled that more automakers remain under scrutiny.

What You Can Do Right Now

I wrote this chapter because I believe that once you understand what is happening, you will want to take action. So here is what you can do today.

Request your LexisNexis report at consumer.risk.lexisnexis.com or by calling 1 888 497 0011. Request your Verisk report at fcra.verisk.com or by calling 1 800 627 3487. You are entitled to one free copy per year from each company under federal law. If you find errors, dispute them in writing.

Review your vehicle's connected services settings. Look for data sharing or driving behavior features that may have been activated without your knowledge, especially if your car is a GM, Toyota, Honda, Hyundai, Kia, or Stellantis brand. Read the privacy settings in your vehicle's companion app, and disable everything you can.

Think twice before pairing your phone to a rental car or a vehicle you do not own. The infotainment system will store your contacts, call logs, and text messages, and that data can persist long after you return the keys.

If you live in California, exercise your CCPA rights. Submit a data access request to your automaker to find out exactly what they have collected. Then submit a deletion request. Then opt out of the sale of your personal information. You can do all of this through the automaker's privacy portal or by contacting their privacy office directly.

Talk to your elected representatives. Tell them you want a federal law that establishes vehicle data as belonging to the person who owns the vehicle, not the company that manufactured it. Three bills are sitting in Congress right now. They need public pressure to move forward.

And share this information with the people you care about. Because most Americans have no idea that the car in their driveway is collecting this kind of data. They do not know that their insurance rates may already be affected. They do not know that their location history is being sold. The single most effective thing any of us can do right now is make sure more people know the truth.

Your car should take you where you want to go. It should not tell the world where you have been.

Chapter 06: Once Someone Captures Your Face and Voice, No Password Reset on Earth Can Fix It

Here is a question I need you to sit with for a moment. If someone stole your credit card number tonight, how long would it take you to fix it? A phone call, maybe two. The bank cancels the old number and mails you a new card within a week. Now answer this one. If someone stole your face, your fingerprints, or your voice, what would you do? Call the bank and ask for a new face? Request a replacement set of fingerprints?

You cannot do that. You will never be able to do that. And that single fact is the reason biometric privacy is the most dangerous and least understood threat to your personal freedom in America today.

Right now, your face is stored in databases you have never heard of. A company called Clearview AI has scraped more than 50 billion photographs from Facebook, Instagram, YouTube, LinkedIn, and thousands of other websites, building a searchable facial recognition database that at least 3,100 law enforcement agencies have used to identify people. You were not asked. You were not told. Your face was taken and filed away like inventory in a warehouse.

At the same time, cameras at more than 250 airports, all 30 NFL stadiums, thousands of retail stores, and police departments from coast to coast are scanning faces and matching them against databases every single day. If the system gets it right, nobody notices. If the system gets it wrong, an innocent person goes to jail. And the system gets it wrong far more often when the face it is scanning belongs to a person of color, a woman, a child, or an elderly person.

I am going to walk you through exactly how this technology works, where it is watching you, why the people who designed it built bias into its bones, and what you can do starting today to protect yourself and your family. I am also going to tell you about a grandmother in Tennessee who was arrested at gunpoint while babysitting four children, charged with a crime in a state she had never visited, and locked in a jail cell for 108 days because a facial recognition algorithm said she was someone she was not. When you hear her story, you will understand why this chapter matters more than you think.

Your Face, Reduced to Math

The technology behind facial recognition is deceptively simple. A camera captures an image of your face. Software maps the unique geometry of that face across roughly 80 points, measuring things like the distance between your eyes, the contour of your jawline, the width of your nose, and the depth of your eye sockets. A neural network then converts those measurements into a compact numerical code called a faceprint. Think of it as a mathematical fingerprint for your face. That code is unique to you. It belongs to you. And once it exists in a database, it is there for good.

From there, the system does one of two things. It can verify you against your own ID photo, which is called one to one matching. This is what TSA does at airport checkpoints when it compares your face to the photo on your license or passport. It can also search an entire database to figure out who you are, which is called one to many matching. This is what police do when they feed a surveillance camera image into a system and ask it to find a match among millions of stored faces. The FBI alone can search against more than 411 million photographs, including passport photos, visa photos, and military images.

Two types of systems exist in the market right now. Traditional systems compare flat, two dimensional photographs and struggle when lighting is bad or the face is turned at an angle. Three dimensional systems, like the Face ID technology in your iPhone, use infrared light to map the depth and contours of your face in the dark. Newer skin texture analysis can even distinguish identical twins by reading microscopic pore patterns invisible to the naked eye. The U.S. facial recognition market hit 1.75 billion dollars in 2025 and is projected to reach 3.89 billion dollars by 2030. This is not emerging technology. It is here, it is growing, and it is everywhere.

The Places That Are Already Scanning Your Face

Let me take you through a day in the life of an American who has no idea how many times their face is captured, analyzed, and stored.

You fly to visit your family. At the airport, TSA has deployed facial recognition scanners at more than 250 airports, covering roughly 58 percent of all commercial airports in the country. In summer 2025, TSA launched its Touchless ID program for PreCheck members, where no physical ID is required. A camera matches your face against a gallery of passport photos. In January 2026, TSA announced it would expand this program to 65 airports by spring 2026, with a goal of reaching more than 400 airports by the end of the year. All five major airlines, Alaska, American, Delta, Southwest, and United, are participating. Customs and Border Protection runs a separate facial comparison system at 238 airports for international travelers.

You stop at the grocery store. A January 2026 CNN investigation found that Walmart, Kroger, and Home Depot acknowledge facial recognition capabilities in their privacy policies. Wegmans triggered controversy that same month when signs appeared at its New York City stores disclosing the technology. Madison Square Garden made national headlines for using facial recognition to identify and eject lawyers from firms that had sued its owner. The landmark retail case involved Rite Aid, which deployed facial recognition in hundreds of stores between 2012 and 2020 to flag suspected shoplifters using a watchlist built from low quality security camera images and employee cellphone photos. The FTC banned Rite Aid from using facial recognition for five years after finding thousands of false matches and a pattern of deploying the technology in predominantly Black, Asian, and Latino communities.

You take your kids to a game. The NFL completed league wide deployment of facial recognition across all 30 stadiums during the 2024 to 2025 season. Major League Baseball runs a facial recognition entry system called Go Ahead Entry at ten ballparks, including Dodger Stadium and Citizens Bank Park, with entry lanes running 141 percent faster than traditional gates. A 2025 industry survey found that 47 percent of venue operators named biometrics as a top priority for the coming year.

Your kids go to school. New York's Lockport City School District invested four million dollars in a facial recognition security system in 2019, sparking a backlash that ultimately led New York to become the first state to ban facial recognition in schools. More than 60 colleges and universities have committed to refusing the technology.

And if you have ever posted a photo to social media, Clearview AI has almost certainly scraped it into a database of 50 billion images that police across the country search every day. That is where we are. Your face is being captured at the airport, at the grocery store, at the stadium, at your child's school, and from your own social media accounts.

Once It Is Stolen, It Is Stolen Forever

I want to make sure you understand why biometric data is fundamentally different from every other kind of personal information. When your credit card number gets stolen, the bank issues a new one. When your password gets hacked, you create a new password. When your Social Security number is compromised, you can freeze your credit and monitor for fraud. These are serious problems, and I do not want to minimize them. They cause real harm to real people. They also share one saving grace. You can take steps to contain the damage because the compromised information can be replaced or restricted.

Your face cannot be replaced. Your fingerprints cannot be replaced. Your iris patterns cannot be replaced. Your voice cannot be replaced. You have exactly one face, ten fingerprints, two irises, and one voice. Once a biometric template is stolen from a database, there is no reset button. The victim is exposed for life, on every system that uses biometric authentication.

The breach record tells you everything you need to know about how seriously companies and governments take this responsibility. In 2015, Chinese hackers breached the Office of Personnel Management and stole 5.6 million fingerprint records alongside 22.1 million personnel files. Biometrics experts warned that intelligence agents could now be identified by fingerprint even when operating under assumed names. The government spent over 130 million dollars on identity protection services. No service on earth can protect a fingerprint that has already been copied.

In 2019, a security company called BioStar 2 exposed 27.8 million records, including raw fingerprint images from banks, defense contractors, and the UK Metropolitan Police. The company had stored the fingerprints in their original form, meaning attackers could fabricate physical replicas. Clearview AI suffered a data breach in February 2020 that exposed its entire client list. In October 2023, 815 million Indian citizens' biometric records from the Aadhaar database were offered for sale on the dark web for approximately 80,000 dollars.

Between 2018 and 2023, nearly 6 billion biometric records were compromised around the world. The average biometric breach now costs 5.22 million dollars, making it one of the most expensive categories of data to lose. Researchers have even demonstrated that encrypted biometric templates, which were once considered safe, can be reverse engineered. A research team showed the complete attack chain in a 2025 paper: take a stolen fingerprint template, feed it through a generative AI model, reconstruct the fingerprint image, print it on a silicone mold, and use that mold to successfully pass commercial fingerprint scanners. This is not science fiction. This is happening now.

The Four Billion Dollar Lawsuit That Changed Everything

In 2008, Illinois passed a law called the Biometric Information Privacy Act, known as BIPA. What made BIPA different from every other privacy law in America was a single provision. It gave ordinary people the right to sue. If a company collected your fingerprint, your faceprint, your voiceprint, or your iris scan without your informed written consent, you could take them to court. The law set damages at 1,000 dollars per negligent violation and 5,000 dollars per intentional or reckless violation, plus attorneys' fees.

That provision turned BIPA into the most consequential privacy law in the country. More than 2,000 lawsuits have been filed since 2017. Facebook settled for 650 million dollars in 2021 for scanning and tagging Illinois users' faces without consent. BNSF Railway faced a 228 million dollar jury verdict in 2022, the first BIPA jury

trial in history, for requiring truck drivers to scan fingerprints at railyards without permission. Google paid 100 million dollars for extracting facial templates through Google Photos without consent. TikTok paid 92 million dollars for harvesting biometric data from users. Clearview AI settled for a 51.75 million dollar equivalent in March 2025, and because the company did not have enough cash, the settlement gave class members a 23 percent equity stake in the company instead.

Texas joined the fight with even bigger numbers. The Texas attorney general secured a 1.4 billion dollar settlement from Meta in July 2024, the largest single state privacy settlement in American history, for running facial recognition on virtually every face uploaded to Facebook for over a decade. In May 2025, Texas followed that up with a 1.375 billion dollar settlement from Google for collecting voiceprints and facial geometry through Google Photos, Google Assistant, and Nest Hub Max without consent. Texas alone extracted 2.775 billion dollars from two companies.

Twenty states have now enacted comprehensive privacy laws that classify biometric data as sensitive and require heightened consent. No federal biometric privacy law exists. Congress has tried. The Facial Recognition and Biometric Technology Moratorium Act and the American Privacy Rights Act have gone nowhere. As of March 2026, the states are fighting this battle alone.

Voluntary in Name Only

The TSA says its facial recognition program is entirely voluntary. I want to tell you what voluntary looks like in practice.

The Algorithmic Justice League, led by researcher Joy Buolamwini, published a report in July 2025 called "Comply to Fly?" that examined 420 traveler experiences across 91 airports. The findings were devastating. Ninety nine percent of travelers were not verbally told they could opt out. Half did not see any signage about the opt out option. Three quarters were completely unaware that opting out was even possible. Among the small number of travelers who did opt out, 67 percent reported negative treatment. TSA officers made hostile comments, used aggressive body language, subjected travelers to increased scrutiny, and caused delays. At Seattle Tacoma airport in December 2024, a TSA officer told a traveler who declined the face scan, "Really? That's ridiculous, you must be stupid."

A landmark 125 page staff report from the Privacy and Civil Liberties Oversight Board, published in May 2025 after a six year investigation, confirmed these failures. The board found that TSA has never published a single comprehensive Privacy Impact Assessment for its facial recognition program. The DHS directive governing facial recognition use disappeared from the DHS website after the change in administration in January 2025, and DHS could not confirm whether it remains official policy.

A bipartisan bill called the Traveler Privacy Protection Act, introduced in May 2025 by Senators Kennedy and Merkley, would have made human ID checks the default, required affirmative consent before each facial scan, and prohibited negative treatment of passengers who opt out. The airline industry killed it. Airlines for America, the U.S. Travel Association, and major airport groups sent a joint letter arguing the bill would increase wait times. Senator Cruz pulled the bill from the committee agenda at the last moment. It remains stalled.

If you want to opt out today, tell the TSA officer "I would like to opt out of the face scan" before the photo is taken. You do not need to explain why. Children under 18 are not photographed. If an officer gives you trouble, note their name and file a complaint with the TSA Contact Center. As Joy Buolamwini said, giving up your face data should not be the price of getting on an airplane.

When the Algorithm Gets It Wrong, Innocent People Lose Everything

The stories I am about to share with you are the reason I am writing this book. These are real people whose lives were destroyed because a computer said they were someone they were not.

Robert Williams is a black man who lives in Farmington Hills, Michigan. In January 2020, police arrested him in front of his wife and two young daughters for allegedly stealing watches from a store. The arrest was based on a match between grainy surveillance footage and his expired driver's license photo. He spent more than 30 hours in a filthy, overcrowded detention cell. During interrogation, police showed him the surveillance photo. He held it up next to his own face and said, "I hope you don't think all black people look alike." The case was dismissed. Williams was later diagnosed with PTSD and suffered a series of strokes. His June 2024 settlement with Detroit included 300,000 dollars and landmark policy reforms requiring police to obtain independent evidence before any facial recognition based arrest.

Porcha Woodruff, a 32 year old black woman, was arrested in February 2023 for carjacking and robbery. She was eight months pregnant. Six officers arrived at her Detroit home while she was getting her children ready for school. She pleaded with police to check whether the actual suspect in the video was pregnant. They declined. She spent 11 hours in custody and began having contractions from stress and dehydration. The actual perpetrator was not pregnant. Prosecutors dismissed the case within a month.

Nijeer Parks, a 33 year old black man from Paterson, New Jersey, was accused of shoplifting and trying to hit an officer with a car in Woodbridge, a town he had never visited. Police had run a blurry image from a fake Tennessee driver's license through facial recognition. Parks spent 10 days in jail and faced charges for nearly 10 months. Police ignored DNA and fingerprint evidence pointing to a different person.

Angela Lipps may be the most devastating case of all. She is a 50 year old grandmother from Tennessee. On July 14, 2024, U.S. Marshals arrested her at gunpoint while she was babysitting four children, charging her with bank fraud in Fargo, North Dakota, a state she had never visited. Bank records confirmed she was 1,200 miles away during every single transaction. She spent 108 days in jail before being transferred to North Dakota, where she first spoke with Fargo police on December 19, 2025, five full months after her arrest. Charges were dismissed on Christmas Eve. She was left stranded in Fargo without money or a coat. She lost her home, her car, and her dog. As of March 2026, Fargo police had not apologized.

A January 2025 Washington Post investigation documented at least eight wrongful arrest cases and found that across them, police failed to check alibis in six, ignored evidence pointing to someone else in two, and neglected to collect key evidence in five. In every single case, police arrested someone without independently confirming that person's connection to the crime. These are not glitches in a system that mostly works. These are the predictable consequences of a system that was built on biased data and deployed without adequate safeguards.

The Technology Fails the People Who Need Protection Most

The accuracy problems in facial recognition are not random. They fall along lines of race, gender, and age with disturbing consistency.

The most authoritative data comes from NIST, which evaluated 189 algorithms from 99 developers using 18 million images. The December 2019 report found that false positive rates, the rate at which the system incorrectly says two different people are the same person, were 10 to 100 times higher for Black and Asian faces than for white faces. American Indian faces had the highest false positive rates of any group tested. Women were misidentified more often than men across nearly all algorithms. Children and the elderly showed elevated error rates. In one to many searches, the type police rely on, most algorithms selected incorrect matches among Black women at significantly higher rates than any other demographic group.

The foundational research came from Joy Buolamwini, then a graduate student at MIT, who discovered the problem firsthand when facial recognition software could not detect her dark skinned face. She literally had to put on a white mask to be seen by the system. Her 2018 Gender Shades study found error rates of up to 34.7 percent for darker skinned women, compared to 0.8 percent for lighter skinned men. When the ACLU tested Amazon's Rekognition against photos of every member of Congress in 2018, it falsely matched 28 lawmakers as criminals, disproportionately people of color.

The bias exists for identifiable reasons. Training datasets are overwhelmingly composed of lighter skinned male faces. Camera technology has been calibrated for lighter skin tones since the 1950s, when Kodak used cards featuring exclusively white models to set color balance standards. Darker skin reflects less light, giving algorithms fewer distinguishing details to work with. Mugshot databases used by police are themselves racially skewed by decades of disproportionate policing, meaning the systems are trained on and deployed against the same communities.

IBM exited facial recognition entirely in June 2020. Amazon imposed a moratorium on police use of its Rekognition product. Microsoft banned police sales pending a national law grounded in human rights. No such law has arrived, and the broader market continues to grow.

Your Voice Is the Next Biometric Under Attack

Your voice is a biometric identifier, analyzed and stored just like a fingerprint or a faceprint. Major banks, including JPMorgan Chase, Wells Fargo, TD Bank, Charles Schwab, and Bank of America, use voiceprint authentication to verify callers. The system analyzes more than 100 characteristics of your speech, including pitch, cadence, accent, and the shape of your vocal tract. Your voiceprint is stored as a mathematical template. It cannot be changed if compromised.

AI can now clone any voice with startling accuracy. Tools from Microsoft, OpenAI, and ElevenLabs can produce a functional clone from as little as 5-10 seconds of recorded audio. A 2024 report identified more than 350 voice cloning tools on the market. McAfee Labs found that a convincing clone could be created for five dollars and ten minutes of setup time.

These cloned voices are defeating bank security systems. University of Waterloo researchers demonstrated in 2023 that they could bypass voice authentication with a 99 percent success rate within six attempts. A journalist used a free voice cloning tool to replicate his own voice and successfully breached a major bank's

voice ID system. A survey found that 91 percent of U.S. banks are now reconsidering their voice authentication programs because of AI cloning.

The fraud is already causing real financial damage. In 2021, criminals used cloned voices to steal 35 million dollars from a bank in the UAE. In January 2024, finance workers at a British engineering firm transferred 25.6 million dollars after a deepfake video call in which every participant, including the apparent CFO, was AI generated. Sharon Brightwell of Dover, Florida, wired 15,000 dollars in July 2025 after receiving a panicked call from what sounded exactly like her daughter. An estimated one in four Americans has already been targeted by a voice cloning scam.

The FCC ruled in February 2024 that AI generated voices qualify as artificial under the Telephone Consumer Protection Act, making unauthorized AI robocalls illegal. Tennessee became the first state to expressly protect against AI voice cloning through its ELVIS Act. Texas followed with its own disclosure and consent requirements.

The single best protection against voice cloning scams is old fashioned and completely free. Pick a family code word, a secret phrase known only to your closest relatives, that must be spoken in any emergency call asking for money. If someone calls claiming to be your child, your spouse, or your parent in distress, ask for the code word. Hang up and call them back at a number you know. Limit the audio and video you post publicly on social media, because that content is the raw material cloning tools need to replicate your voice.

Where Things Stand Right Now

The biometric privacy situation shifted dramatically over the past year. At the federal level, the incoming administration revoked the previous AI Executive Order on its first day in office, removing the federal government's primary oversight framework. All three Democratic members of the Privacy and Civil Liberties Oversight Board were fired in January 2025, leaving the board without a quorum. The DHS directive governing facial recognition use by ICE and Customs and Border Protection was quietly removed from the government's website. A December 2025 executive order directed the Attorney General to create a task force to challenge state AI and privacy laws deemed burdensome.

Federal biometric surveillance expanded at the same time. ICE deployed an app called Mobile Fortify that enables field agents to perform facial recognition, fingerprint, and iris scans during street encounters, drawing from more than 200 million images in government databases. Internal documents confirm that ICE does not give individuals any opportunity to decline biometric collection. The 9.2 million dollar Clearview AI contract for ICE was the company's largest government deal. Pending federal legislation has allocated 673 million dollars for biometric entry exit systems and 2.77 billion dollars for AI powered surveillance infrastructure.

The states remain the strongest source of resistance. Texas's combined 2.775 billion dollars in settlements from Meta and Google proved that a single motivated attorney general can force accountability on the largest companies in the world. Colorado's biometric amendment took effect in July 2025. At least 23 states now regulate biometric data in some form. Internationally, the EU AI Act's biometric prohibitions took effect on February 2, 2025, banning real time facial recognition in public spaces and the untargeted scraping of facial images for databases, with penalties up to 35 million euros or 7 percent of global revenue.

What You Need to Do Right Now

I am not going to sugarcoat this. Most Americans have already lost control of their biometric data. Your face has been photographed thousands of times. Your voice exists in voicemails, social media videos, and customer service recordings. Clearview AI's database almost certainly includes your photos. The question is not whether your biometric data is out there, because it is. The question is what you do next.

At the airport, tell the TSA officer you want to opt out of the face scan before the photo is taken. You do not need to provide a reason. If an officer gives you a hard time, write down their name and file a complaint with the TSA Contact Center. When possible, choose biometric systems that keep your data on your device, like Apple's Face ID, which stores your faceprint in a secure chip on your phone and never sends it to a server. Pair any biometric login with a password or security key through multi factor authentication so that even if one layer is compromised, the other still stands.

Before you hand over your biometric data to any company, ask four questions. Where is it stored? Who has access to it? How long do you keep it? Can I opt out? If they cannot give you clear answers, walk away.

Cut your social media footprint. Publicly posted photos and videos are the raw material that feeds facial recognition databases and voice cloning tools. Set your profiles to private. Be careful about unexpected phone calls where no one speaks on the other end, because those calls may be harvesting a sample of your voice. Watch for callers who try to get you to say "yes" on the phone. Pick a family code word for emergencies and make sure everyone in your household knows it.

If you live in Illinois, Texas, or one of the growing number of states with biometric privacy laws, know that you have legal options if a company collects your biometric data without consent. Find a lawyer who handles privacy cases and learn what your state law allows.

The faceprint captured at the grocery store today will still identify you in 50 years. The voiceprint recorded during a customer service call this week can be cloned by AI this afternoon. The fingerprint stolen in a data breach last year cannot be changed, cancelled, or reissued. These are not passwords. These are pieces of your body. They are permanent, they are irreplaceable, and they deserve the same fierce protection you would give to anything else you love and cannot replace.

Talk to your family about this. Talk to your friends. Share what you have learned. The companies and government agencies collecting your biometric data are counting on your silence and your ignorance. The moment you start paying attention, asking questions, and demanding answers, the equation starts to change.

Your face and voice are not a passwords you can change. Start treating them that way.

Chapter 07: Protesting, ICE, Police and Surveillance

They Know Who You Are. They Know You Were There.

Nicole Cleland is a 56 year old woman who lives in Richfield, Minnesota. On January 10, 2026, she was sitting in her parked car near a street where ICE agents were conducting an operation. She had been watching from a distance. She had never spoken to any federal agent. She had never been arrested, never been charged with a crime, never had so much as a speeding ticket related to any protest.

An ICE agent walked up to her car, looked at her, and called her by name.

She had never met this person. She had never introduced herself. The agent told her he had facial recognition technology on his phone and that his body camera was recording her. Days later, Nicole discovered her TSA PreCheck and Global Entry status had been revoked. A lifelong American citizen, surveilled by her own government, identified by an algorithm, and punished for doing nothing more than being present on a public street.

Nicole Cleland is not a criminal. She is not a suspect. She is an American who happened to be near a protest. And the federal government used military grade surveillance technology to identify her, catalog her, and retaliate against her.

If you think this is an isolated case, keep reading. If you think your rights protect you from this kind of treatment, keep reading. And if you think you have nothing to worry about because you have nothing to hide, you need to read every single word of this chapter.

The surveillance apparatus aimed at American protesters in 2025 and 2026 represents the most formidable threat to the freedom of assembly since the FBI ran COINTELPRO against Martin Luther King Jr. and the civil rights movement. Predator drones are circling over American cities. Facial recognition apps are scanning faces in real time. Data brokers are selling your location history to federal agencies. Police departments are running your license plate through a national database the moment you park your car near a march. And most Americans have no idea any of this is happening.

This chapter will show you exactly what surveillance tools are being used against people who exercise their constitutional right to peaceful protest. You will learn what your legal rights are when police or ICE agents approach you. You will learn whether you have to hand over your phone, answer questions, or identify yourself. And you will learn the specific, concrete steps you need to take to protect yourself and your family before you ever set foot at a demonstration.

The Surveillance Arsenal Aimed at You

Law enforcement agencies at the federal, state, and local level now operate an interlocking web of surveillance tools capable of identifying nearly anyone who attends a protest. Most of these technologies were developed for counterterrorism and border security. They migrated into protest monitoring gradually and quietly. Their use at demonstrations across the country is now routine and well documented.

Start with body cameras. Police departments introduced them as accountability tools after the unrest in Ferguson, Missouri. Every officer wearing one at a protest records the faces, conversations, associations, and movements of everyone within range. Retention policies vary from 60 days in some departments to permanent storage for anything classified as a critical incident. In December 2025, Axon, the company supplying body cameras to roughly 47 of the 69 largest police departments in America, reversed a moratorium on facial recognition and began testing face recognition directly on body cameras. Axon's own ethics board had previously concluded the technology was not reliable enough to justify its use. California's SB 1038 permanently bans facial recognition on police body cameras within the state. No federal ban exists. Every other state allows it.

Then there are the devices called Stingrays. These are portable machines that mimic cell towers and force every phone within roughly a third of a mile to connect to them. When your phone connects, the device harvests your unique identifier, telling law enforcement exactly who was carrying a phone near a protest. Advanced models force phones to downgrade to less secure connections, allowing agents to intercept calls and text messages. Police have used these devices at protests since at least 2003, when Miami police purchased one to surveil demonstrators at a Free Trade of the Americas conference. A Department of Homeland Security Inspector General report in February 2023 found that ICE and the Secret Service had violated the law in 2020 and 2021 by deploying cell site simulators without the required court orders.

In July 2025, a news outlet detected suspicious cellular anomalies during an anti ICE protest in Tukwila, Washington, using a detection tool built by the Electronic Frontier Foundation called Rayhunter. Two signals consistent with cell site simulator behavior appeared at the height of the protest and disappeared when the protest ended. ICE declined to comment.

Overhead, the surveillance gets even more alarming. In June 2025, Customs and Border Protection confirmed flying two MQ 9 Predator drones over Los Angeles during anti ICE protests. These are the same military grade unmanned aircraft used in combat zones, equipped with infrared sensors, high definition cameras, and artificial intelligence radar capable of detecting individual human beings within a 15 nautical mile radius. CBP claimed the cameras could not identify a person, then acknowledged they could detect clothing color, backpacks, and weapons. Flights continued for days. During the 2020 protests after the killing of George Floyd, the Department of Homeland Security deployed aerial surveillance over 15 or more cities, logging more than 270 hours of footage. A Predator drone circled over Minneapolis for about 90 minutes on May 29, 2020. FOIA documents revealed that CBP intercepted encrypted messaging chats among Portland protest organizers.

Your License Plate Gives You Away

An investigation by the Electronic Frontier Foundation published in November 2025 revealed the most detailed picture yet of how Automated License Plate Readers are being used to track protesters. The EFF obtained datasets covering 12 million searches by 3,900 agencies between December 2024 and October 2025. More than 50 federal, state, and local agencies ran hundreds of searches through a company called Flock Safety connected to protest activity. Nineteen agencies conducted searches specifically tied to No Kings protests. Tulsa, Oklahoma logged at least 38 protest related searches. Arizona's Department of Public Safety logged searches for phrases like "no kings rock threat." Wichita, Kansas logged 22 license plate searches under a heading about causing harm during protests.

Flock Safety now operates cameras in 2,000 cities across 42 states, with an estimated 90,000 cameras nationwide. The system captures your license plate, your vehicle's make, model, and color, the time, and your location. Everything uploads to a searchable cloud database. Flock is expanding from still photos to video with AI powered search and is planning to integrate with commercial data brokers for what it calls people lookup functionality. If you drive to a protest, or park near one, or simply pass through the area on your way to the grocery store, your plate and your movements are captured and stored.

Shortly before the October 2025 No Kings protests, Amazon's Ring doorbell network announced integration with Flock Safety's system, allowing police to issue geofenced requests for user submitted footage. Residential neighborhoods became auxiliary surveillance grids the moment a march passed through.

The Digital Dragnet: Geofence Warrants and Tower Dumps

Geofence warrants work like this. Law enforcement sends an order to a technology company, usually Google, demanding a list of every device present within a specific geographic area during a specific time window. Google's location database, called Sensorvault, once held data on approximately 592 million accounts. By 2021, geofence warrants made up 25 percent of all warrant requests Google received in the United States. The number climbed from 982 in 2018 to 11,554 in 2020.

On August 9, 2024, the Fifth Circuit Court of Appeals ruled unanimously that geofence warrants are categorically prohibited by the Fourth Amendment. The court called them modern day general warrants, the exact kind of government overreach the Founders wrote the Fourth Amendment to prevent. The ruling found that searching Google's entire 592 million account database without naming a specific suspect amounts to the kind of general exploratory rummaging the Constitution forbids.

Here is where the story takes a turn that should concern every American. The court found geofence warrants unconstitutional and then upheld the criminal conviction anyway, under something called the good faith exception. This pattern has repeated in case after case. Courts declare the surveillance illegal. The evidence collected through it stays in. The government faces no consequences.

Google moved on its own to end the practice. In December 2023, Google announced it would store location data on devices rather than in its cloud, cut auto deletion from 18 months to 3 months, and encrypt cloud backups end to end. By July 2025, Google completed this transition. The era of geofence warrants against Google is functionally over. The Supreme Court granted review of a separate geofence case, Chatrie v. United States, in January 2026. A ruling is expected by June 2026.

Tower dumps are a related threat. Police send requests to phone carriers demanding records of every device connected to a specific cell tower during a specified time. A single tower dump captures data on tens of thousands of people. In one documented case, more than 50,000 individuals. AT&T stores this data for seven years. In February 2025, a federal judge in Mississippi became the first to declare tower dumps unconstitutional. A federal judge in Nevada reached the same conclusion two months later. The Department of Justice has appealed.

Seven Million Americans March, and the Government Is Watching

The No Kings movement organized mass demonstrations against government overreach under the Trump administration. The first major protest day was June 14, 2025. The largest came on October 18, 2025, when an estimated seven million Americans took part at more than 2,600 locations in all 50 states. It was the largest single day protest against a sitting U.S. president in modern American history.

The organizers anticipated surveillance. The No Kings website urged participants to encrypt their devices, remove biometric locks, sign out of social media, communicate through Signal, and avoid photographing other protesters' faces. The ACLU prepared Know Your Rights materials in ten languages and distributed millions of cards.

The government response was aggressive on multiple fronts. House Speaker Mike Johnson called the protests hate America rallies. In Texas, Governor Greg Abbott activated the National Guard before marches began. And across the country, law enforcement ran license plates, monitored social media, deployed drones, and scanned faces.

Facial Recognition at Your Front Door

The New York Times documented more than half a dozen activists in Minnesota who were subjected to facial recognition scans in January 2026. ICE agents used an app called Mobile Fortify, built by NEC, that searches against more than 200 million images stored in government databases.

Remember Nicole Cleland, the woman from the opening of this chapter. She was one of at least seven American citizens told by ICE agents in and around Minneapolis that they were being recorded with facial recognition. In one video, agents told people their faces would be added to a domestic terrorism database. Mubashir Khalif Hussen, a 20 year old Somali American U.S. citizen, was grabbed by agents, refused the right to show his passport, driven seven miles to a facility, had his face scanned, and then released in a Minnesota winter and told to walk back.

ICE signed a 9.2 million dollar contract with Clearview AI in September 2025, expanding the technology's use to include assaults against law enforcement officers. The Illinois Attorney General alleged the Mobile Fortify app had been used over 100,000 times since its debut. The ACLU of Minnesota filed a class action lawsuit against ICE and CBP over forced facial scans and other practices it called illegal.

The full scope of ICE's surveillance infrastructure is staggering. Georgetown Law's Center on Privacy and Technology published a two year investigation called American Dragnet, built from hundreds of public records requests. The findings show that ICE has scanned the driver's license photos of one in three American adults using facial recognition. ICE has access to driver's license data covering three out of four adults. ICE can track vehicle movements in cities where three out of four adults live. ICE spent approximately 2.8 billion dollars between 2008 and 2021 on surveillance and data collection programs. The May 2025 foreword to the report states plainly: ICE now operates as a domestic surveillance agency.

ICE agents also tracked protest activity through social media. During the George Floyd protests, a company called Dataminr sent police departments real time alerts pinpointing protester locations, complete with social media handles, follower counts, and biographical details. The D.C. Metropolitan Police alone received roughly 160,000 Dataminr alerts between June 2020 and May 2022 tracking demonstrations. Another firm called

Media Sonar, used by the Fresno Police Department, recommended that officers track hashtags related to Black Lives Matter and marketed its ability to help police avoid the warrant process entirely.

The federal government also pursued retaliation against people connected to protests. ICE arrested Mahmoud Khalil, a lawful permanent resident and Columbia University student, for pro Palestinian activism. The State Department used AI assisted social media monitoring to revoke approximately 100,000 visas by January 2026, including 8,000 student visas. Students were flagged for merely liking or resharing social media posts. The Department of Justice opened a criminal investigation of Minnesota Governor Tim Walz and Minneapolis Mayor Jacob Frey for their public statements criticizing ICE operations. A deputy attorney general posted on social media that he would stop them from, in his words, their terrorism by whatever means necessary.

Your Rights When Police or ICE Approach You at a Protest

Every American needs to understand these rights. They exist whether you are a citizen, a lawful permanent resident, a visa holder, or undocumented. The Constitution protects everyone on U.S. soil.

You have the right to take pictures and record video of police and federal agents in public. Every federal circuit court that has addressed this question has ruled in your favor. The key cases span from the First Circuit to the Tenth Circuit. No federal circuit court has ruled against the right to record. The Supreme Court has not taken up the question directly, and every lower court relies on First Amendment precedent to protect it. Police cannot order you to stop recording because they dislike being filmed. DHS published a statement in December 2025 suggesting that recording federal officers sounds like obstruction of justice. That statement contradicts the rulings of seven federal circuit courts and has been widely condemned by constitutional scholars. If an officer tells you to stop recording, you do not have to comply. If an officer deletes your photos or videos, that officer has violated the First, Fourth, and Fourteenth Amendments and faces personal civil rights liability.

You have the right to remain silent. Under the Fifth Amendment, you do not have to answer questions from police or ICE about your activities, your associations, or your reasons for attending a protest. State clearly, I am invoking my right to remain silent. Twenty four states have stop and identify laws that require you to give your name during a lawful stop based on reasonable suspicion of criminal activity. Attending a protest does not create reasonable suspicion.

You have the right to refuse a search of your phone. The Supreme Court ruled unanimously in Riley v. California in 2014 that police need a warrant to search the digital contents of your phone, even if you are under arrest. Officers can seize your phone for safekeeping during an arrest. They cannot open it, browse through it, or read your messages without a warrant. If an officer asks to look at your phone, say clearly, I do not consent to a search.

The question of whether police can force you to unlock your phone depends on how your phone is locked. This area of law is actively changing. At least four state supreme courts have ruled that forcing someone to reveal a passcode violates the Fifth Amendment because a passcode requires the contents of your mind. Biometric unlocking, using your fingerprint or your face, is more legally vulnerable. The Ninth Circuit ruled in 2024 that physically pressing a suspect's finger on a scanner was not a Fifth Amendment violation because it required no thinking. The D.C. Circuit reached the opposite conclusion in January 2025, ruling that

instructing a suspect to unlock a phone with a fingerprint was testimonial and protected. The two courts split on this question, and the Supreme Court will likely resolve it in the coming years.

The practical lesson is clear. Before you attend any protest, disable Face ID and Touch ID on your phone. Switch to a strong alphanumeric passcode of eight to twelve characters. A passcode stored in your mind enjoys stronger constitutional protection than your fingerprint or your face.

If ICE approaches you, know this. Every person in the United States, regardless of immigration status, has the right to remain silent, the right to refuse consent to a search, and the protection of the Fourth Amendment against unreasonable searches. U.S. citizens do not have to carry proof of citizenship and have no obligation to answer ICE questions about their status. Lawful permanent residents and visa holders over 18 are required by law to carry their immigration documents. They still have the right to remain silent about other matters and the right to refuse consent to searches. ICE's prior sensitive locations policy, which designated protests, marches, and rallies as places where immigration enforcement should not happen, was rescinded under the second Trump administration.

How Your Phone Betrays You Before You Even Arrive

The most effective surveillance tool at protests is not a drone, a body camera, or a facial recognition app. Your phone does most of the work for the government without anyone ever touching it.

Here is how the pipeline operates. Weather apps, navigation apps, games, prayer apps, fitness trackers, and hundreds of other applications on your phone request access to your location. Many of those apps embed software from data brokers that continuously collect your GPS data. Companies like Venntel, Babel Street, and Fog Data Science aggregate the data from millions of phones. One company's marketing materials boasted of collecting more than 15 billion location points from over 250 million mobile devices every single day. Federal agencies then purchase access to this location data, bypassing the warrant requirements the Supreme Court established in Carpenter v. United States.

The Office of the Director of National Intelligence published a declassified report in 2022 acknowledging this data could identify every person who attended a protest or rally based on their smartphone location or ad tracking records. ICE purchased tools designed to monitor specific city blocks for mobile phones. The FBI modified its contract with a location data vendor on June 9, 2020, shortly after protests erupted across the country. In March 2026, CBP acknowledged using location data sourced from the real time bidding system, the technical process behind nearly every online advertisement.

ICE's newest vendor, PenLink, offers analytics that search a specific area for mobile phones across a time period, track where those devices go, and monitor multiple locations simultaneously. At the push of a button, ICE agents can catalog everyone who marched at a protest, identify who attended multiple events, and build lists of activists based on which phones appeared at more than one location.

The FTC has started pushing back. In January 2024, it ordered a major location data company to stop selling sensitive location data. That order specifically required protections ensuring data is not associated with locations of public gatherings at political or social demonstrations or protests. In December 2024, the FTC barred another company from selling sensitive location data except in limited national security circumstances. Montana became the first state to close the data broker loophole entirely in 2025, prohibiting law

enforcement from purchasing data that would require a warrant to obtain. The federal Fourth Amendment Is Not For Sale Act passed the House in April 2024. It stalled in the Senate.

Protect Yourself: What to Do Before, During, and After a Protest

Disable biometrics before you leave the house. Switch from Face ID or Touch ID to a strong alphanumeric passcode of eight to twelve characters. This gives you stronger legal standing against forced unlocking.

Turn on airplane mode and then separately disable location services, Wi Fi, and Bluetooth. Airplane mode alone does not disable GPS. Your phone still receives satellite signals and logs your location for later transmission when you reconnect. The safest option is to turn the phone off completely. Be aware that even powered off iPhones are sometimes locatable through Apple's Find My network.

Carry a Faraday bag if you want physical, not software based, protection. A Faraday bag blocks all electromagnetic signals based on physics. Your phone is completely unreachable while inside, so plan your communications around this limitation.

Use Signal for all protest related communications. When law enforcement has subpoenaed Signal, the company has only been able to provide the date of registration and the date of last connection. Signal also includes a built in face blurring tool for photos. If you'd like to learn more about Signal and how to use it, I recommend, "Everybody Has Something To Hide" by Guy Kawasaki and Madisun Nuismer (I was a consulting expert on the book).

If you want the highest level of protection, leave your phone at home. This eliminates your ability to document what happens or communicate during an emergency. It also eliminates the single most effective surveillance tool the government has.

If you bring a secondary phone, never carry it alongside your primary phone. Location correlation between two devices moving together compromises anonymity. Turn the secondary phone off before you return home to prevent linking it to your residence.

Disable 2G connectivity. This protects against cell site simulators that force phones onto older, less secure networks. On Android, go to Settings, then Network, then SIMs, and disable Allow 2G. On iPhone, enable Lockdown Mode under iOS 17 or later.

Scrub photo metadata before you share anything. Use Signal to strip location data from images, or take screenshots of your photos before posting them. Blur the faces of other protesters. Use cash for public transit to avoid traceable payment records.

Understand that VPNs do not protect you against most protest surveillance. VPNs do not hide your phone's presence from Stingray devices or cell towers. They do not prevent GPS tracking. They do not block tracking through advertising identifiers. The EFF has warned that VPN advertising vastly oversells what the technology actually does.

I share additional smartphone privacy tips weekly in my free LinkedIn newsletter, Smartphone Lock Down.

The Through Line from COINTELPRO to Clearview AI

Everything described in this chapter has a direct historical ancestor. The FBI's COINTELPRO program ran from 1956 to 1971, targeting civil rights leaders, anti war activists, the Black Panther Party, and anyone the government considered subversive. The FBI sent an anonymous letter to Martin Luther King Jr. with surveillance recordings that King and his advisors interpreted as a suicide threat, timed to arrive before his Nobel Prize acceptance. Between 1960 and 1974, the FBI opened more than 500,000 separate investigations of so called subversive persons and groups. Not a single one resulted in a prosecution.

The Church Committee investigated this program over 16 months, interviewing 800 witnesses across 126 meetings. Their conclusion still echoes: too many people have been spied upon by too many government agencies and too much information has been illegally collected.

At Standing Rock in 2016 and 2017, Energy Transfer Partners hired a private military firm that spent 17 million dollars on social media monitoring, aerial surveillance, radio eavesdropping, and infiltration. Leaked reports revealed the firm described water protectors as insurgencies using counterterrorism language. During the 2020 protests after the killing of George Floyd, six federal agencies used facial recognition technology. Nearly 2,000 public agencies used Clearview AI. In February 2026, the Tenth Circuit ruled that police warrants searching all photos, videos, emails, texts, and location data on a protester's devices, including keyword searches for the words protest, human, rights, and cop, were unconstitutional. The court called them wide ranging exploratory searches that the Fourth Amendment was designed to prohibit.

The scale has changed. The FBI needed 500,000 individual case files over 14 years to surveil the civil rights movement. Today, the technology to identify, track, and catalog every person at a protest fits inside an app on a federal agent's phone.

Recording at Protests: What the Law Actually Says

One question comes up again and again. If I record police at a protest, can I get in trouble for recording the audio of someone without their knowledge? The answer depends on where you live. The federal Wiretap Act follows a one party consent standard, meaning you only need your own consent to record a conversation you are part of. Approximately 11 to 13 states require all party consent for recordings, including California, Florida, Maryland, Massachusetts, Pennsylvania, and Washington. The remaining states follow one party consent rules.

Here is the critical point most people miss. Wiretapping statutes apply when someone has a reasonable expectation of privacy. Courts have consistently held that police officers performing their duties in public have no reasonable expectation of privacy. Even in two party consent states, recording police on a public street is generally protected. Video only recording, without capturing audio, falls almost entirely outside wiretapping laws, because those statutes target the interception of oral communications. Your phone records both video and audio, which makes the audio component the main legal consideration in two party consent states. In practice, the First Amendment right to record public officials performing public duties has been upheld by every federal circuit to consider it, regardless of the state's wiretapping classification.

If you attend a protest and want to obtain body camera footage from police afterward, be prepared for resistance. South Carolina exempts all body camera data from public records requests. California, Illinois, and

Florida impose conditions and exemptions. The most common reason agencies give for denying requests is an ongoing investigation. Some departments classify all body camera video as evidence, effectively placing it beyond public access. One Florida agency quoted a price of 18,000 dollars for 84 hours of footage. The barriers are real. The right to request the footage is also real.

Where the Law Stands Right Now

The Government Surveillance Reform Act, reintroduced in March 2026, would require warrants for searches of Americans' communications, ban government purchase of data broker records without a warrant, and regulate cell site simulators. It has been described as the most sweeping reform of surveillance laws in nearly half a century. The ICE Out of Our Faces Act, introduced in February 2026, would ban all ICE and CBP use of facial recognition technology. Section 702 of FISA is set to expire in April 2026.

The Supreme Court's decision in Chatrie v. United States, expected by June 2026, will define whether geofence warrants violate the Fourth Amendment. The ACLU's class action lawsuit in Minnesota against ICE and CBP over forced facial scans is proceeding with more than 30 sworn statements. State consumer privacy laws enacted in 2024 and 2025 are starting to reshape how data brokers operate.

The central problem remains. Courts keep recognizing that mass surveillance tools violate the Fourth Amendment. Then the good faith exception allows the evidence to stay in anyway. The government keeps purchasing from commercial data brokers what the Supreme Court says it needs a warrant to collect directly. Senator Wyden has called this an end run around the Constitution. Until Congress closes the data broker loophole or the Supreme Court extends Carpenter to cover purchased data, the constitutional rights that exist on paper will remain incomplete on the ground.

This Is About Democracy

The right to peaceful protest is not a luxury. It is the mechanism through which every other right in the Constitution has been won, defended, and expanded. The civil rights movement, the labor movement, the women's suffrage movement, the movement to end the war in Vietnam. Every generation of Americans has stood in public, spoken their minds, and demanded accountability from their government.

When the government deploys Predator drones over American cities, scans faces with military technology, tracks phones through invisible data pipelines, and builds databases of people who attend rallies, the message is clear. We are watching you. We know who you are. And there will be consequences.

That message is designed to make you stay home. It is designed to make you think twice before showing up. It is designed to chill the one right that makes all the other rights possible.

Do not let it work.

Know your rights. Protect your phone. Show up. And bring this chapter with you.

Chapter 08: Deepfakes- When Your Eyes Lie to You

You grew up with a simple rule. If you saw a video, heard a voice, or looked at a photo, you had a reason to trust your own senses. That rule shaped family life, school life, work life, and public life. It shaped how you answered the phone. It shaped how you judged a confession, a voicemail, a threat, a campaign ad, a plea for help, a bank instruction, a text from your boss, and a tearful FaceTime call from someone you love.

That rule is breaking apart.

Right now, in March 2026, strangers with a laptop and a few seconds of your voice can make you say things you never said. They can put your face into a clip you never filmed. They can build a false version of you that sounds close enough, looks close enough, and moves close enough to trigger fear, urgency, trust, panic, embarrassment, and obedience. They do not need Hollywood money. They do not need elite technical skill. They need access, speed, and a target.

You.

This chapter matters because deepfakes are no longer a weird internet trick sitting at the edge of culture. Deepfakes and synthetic media now sit inside everyday life. They reach your phone. They enter your child's school community. They show up in the workplace. They move into bank fraud, blackmail, revenge porn, politics, courtrooms, and family emergencies. They prey on the oldest human instinct of all. You believe the people you know. You respond to the voices you trust. You react to what feels real in the moment.

That instinct now works against you.

The good news starts here. Once you understand how this new fraud and privacy machine works, you stop moving through the world on autopilot. You start protecting your voice, your face, your family, your money, and your peace of mind with a new set of habits. You stop relying on old assumptions. You start treating digital media the way a seasoned investigator treats a crime scene. Slowly. Carefully. With questions first and trust second.

That shift is the point of this chapter.

The New Reality: Your Senses Are Now a Target

Deepfakes are synthetic images, audio, and video made by artificial intelligence to imitate a real person or a believable event. Some are obvious jokes. Some are harmless entertainment. Some are political propaganda. Some are criminal tools. Some are acts of humiliation. Some are acts of cruelty. Some are weapons aimed straight at your privacy and safety.

The danger does not begin with perfect realism. The danger begins with something good enough to get you moving.

A fake call from your daughter sobbing and asking for money does not need studio quality. A fake message from your employer asking you to send payroll data does not need a flawless accent. A fake clip of a public

official speaking nonsense does not need to survive a forensic lab. It only needs to survive your first glance, your first listen, your first emotional reaction.

That is where the harm begins.

Most people still think of fraud as a person trying to talk you into doing something foolish. Deepfake fraud works on a deeper level. It borrows identity. It borrows intimacy. It borrows authority. It borrows your trust in the people and institutions that shape your life. That is why this issue belongs in a book about privacy. Privacy is not only about secret data sitting in a server somewhere. Privacy is also about control over your own identity, your own likeness, your own voice, your own presence in the world.

Once someone steals your face or your voice, they are reaching into the most personal layer of your life. They are taking your identity out for a crime spree.

That is privacy harm in its rawest form.

How Voice Cloning Crossed the Line

For years, fake audio carried tells. The voice sounded stiff. The rhythm felt off. The emotional tone drifted. Breathing sounded strange. Pauses felt mechanical. People assumed they would know a fake when they heard one.

That assumption no longer protects you.

By 2025 and into 2026, voice cloning crossed a line that matters in real life. Human listeners often fail to tell the difference between a real recording and a synthetic one, especially during a rushed call, a noisy environment, an emotional moment, or a short exchange. That detail changes everything. Most fraud does not happen in a soundproof lab with trained experts and endless playback. Fraud happens in the middle of your day, when your guard is down and your nervous system is already busy.

A few seconds of voice is often enough to build a convincing imitation. Think about how much of your voice already lives online. A podcast clip. A voicemail greeting. A social media video. A church livestream. A school event recording. A work presentation. A family post someone else uploaded without asking you. That pile of scraps is enough for a bad actor to build a tool that sounds like you.

Once that synthetic voice exists, the scam opens in several directions at once.

A criminal can call your family and pretend you were arrested.

A criminal can call your office and pose as you.

A criminal can call your bank and try to push a transfer.

A criminal can leave voice messages designed to move someone onto a private app where the pressure continues.

A criminal can build trust in stages, one call at a time.

You do not need perfection for this to work. You need emotional force. You need urgency. You need the right target. A scared parent. A distracted employee. A bookkeeper who thinks the CEO is traveling. A relative who already fears something is wrong.

That is why voice trust as a security habit is collapsing. For decades, hearing a familiar voice created comfort. Today, hearing a familiar voice should trigger a pause and a verification step. That change feels unnatural. It feels cold. It feels sad. It also keeps people safe.

The Industrial Age of Impersonation

Deepfakes once lived in viral demos and novelty clips. Now they sit inside a growing criminal business model.

Fraudsters do not need to create every deception from scratch. Online marketplaces and service providers sell tools, templates, and kits. Some platforms offer face swaps. Some offer video avatars. Some offer voice cloning. Some package the whole thing into a service that lowers the skill barrier even further. This is one reason the threat is spreading so fast. The hard work has already been done for the next wave of users.

That shift matters because it turns isolated misconduct into repeatable production. Deepfake crime now looks more like a system than a stunt. It moves across phone calls, texts, email, video meetings, messaging apps, and social media. The goal is simple. Overwhelm your caution before you have time to think.

You can see this change in real world events.

One of the most publicized fraud cases involved a company employee who joined a video conference and believed senior colleagues were on the call. The meeting looked real enough. The faces looked real enough. The pressure felt real enough. Money moved. The loss reached tens of millions of dollars.

That single event sent a message through every workplace in the world. The meeting room is no longer proof. The screen is no longer proof. The face in the square is no longer proof.

Retailers and large businesses now report waves of AI generated scam calls aimed at customer service teams, payment systems, and internal staff. The volume alone tells the story. This is not a rare trick used once in a blue moon. This is constant pressure. That pressure lands hardest on people who answer phones, solve problems, and make quick decisions for a living.

Politics has entered the same danger zone. Synthetic robocalls and manipulated clips do not need to persuade every voter. They need to confuse enough people, suppress enough turnout, or flood the information stream with enough noise to make trust harder. Once the public starts questioning every recording, every denial gains room to breathe. Every real clip becomes easier to dismiss. Every fake clip becomes easier to spread.

This is the liar's playground. If nothing feels certain, accountability suffers.

The Weaponization of Your Face and Your Body

The most vicious use of synthetic media often lands in intimate harm.

Non consensual sexual deepfakes have exploded because they are easy to make, fast to distribute, and devastating to the person targeted. A stranger, former partner, classmate, coworker, or online troll can take an image of a real person and place that face into explicit fake content. The target never consented. The target never posed for that content. The target still pays the price.

The harm is immediate. Shame. Panic. Fear. Isolation. Sleeplessness. Rage. Social withdrawal. Damage to relationships. Damage to work. Damage to school life. Damage to reputation. A spiraling loss of control. For children and teenagers, the danger grows even darker because humiliation spreads at the speed of screenshots and group chats. Once a fake image enters a peer group, a child's daily life can turn into a trap.

There is nothing abstract about this. When your likeness is placed into explicit fake material, the injury lands in your body and mind as if something private was taken from you and displayed without permission. Your nervous system does not care that pixels were generated. Your brain registers violation. Your life feels invaded.

That is why lawmakers in multiple places have started moving toward criminal bans, civil remedies, and takedown duties tied to synthetic sexual abuse. The law is slowly catching up because the human damage is too obvious to ignore. Even so, the law still moves slower than humiliation, slower than reposting, slower than search indexing, and slower than gossip.

That delay is part of the harm.

A person targeted by fake intimate content often enters a brutal race. Remove the content. Preserve the evidence. Tell the platform. Tell a lawyer. Tell law enforcement. Tell your school. Tell your employer. Tell your family. Hold yourself together at the same time. That is a crushing load for anyone. It is especially crushing for a teenager, a college student, or a person already living through domestic abuse, coercion, or stalking.

When people shrug and say the content is fake, they miss the point. The injury is real. The fear is real. The violation is real. The privacy invasion is real.

Why Children Face a Harder Future

Children and teenagers live inside digital identity long before they understand digital risk. Their photos are shared by parents, relatives, schools, teams, and friends. Their voices appear in videos. Their faces appear in public accounts. Their social lives unfold on platforms built for copying, saving, forwarding, and mocking.

Synthetic media multiplies those risks.

A child's likeness can be stolen and used for bullying. A teen's face can be inserted into explicit content. A fake clip can be used for extortion. A manipulated image can be used for grooming. A false voice message can be sent to parents or friends. The emotional shock alone can leave lasting scars. When minors are involved, the consequences cut across privacy, safety, mental health, family stability, and the basic right to grow up without digital exploitation.

Parents often think the main online danger is oversharing private details like addresses or school names. That danger still matters. A deeper danger now lives in the simple existence of a child's digital likeness. A face is data. A voice is data. A laugh is data. A short video from a soccer game is data. Once posted, those pieces can be copied, saved, studied, and reused.

That truth changes the meaning of family privacy.

A child does not need to hand over private information for harm to happen. A child only needs to be visible.

What Deepfake Detection Really Means

When people hear the phrase deepfake detection, they often picture a magic scanner that tells you yes or no. Real or fake. Safe or unsafe. That image is comforting. It is also misleading.

Deepfake detection is the process of looking for signs of generation, manipulation, or tampering. That work can involve visual clues, audio clues, file structure clues, editing traces, compression patterns, timing issues, lip sync problems, lighting inconsistencies, source anomalies, and many other signals. Experts use layered methods because no single clue settles the question every time.

That detail matters for you because detection is not a one button truth machine.

A tool may flag suspicious media. A tool may miss altered media. A tool may struggle once a file has been copied, compressed, reposted, trimmed, or screenshotted. A tool may perform well in testing and lose strength in the messy conditions of the real world. A detector score is a clue. It is not a final verdict.

You need a stronger mental model.

Think of deepfake detection like smoke in the air. Smoke tells you something needs attention. Smoke does not tell you every detail about the fire. You still need context. You still need source information. You still need timing. You still need to know who recorded the material, where it came from, and what happened before and after it appeared.

That is why legal and investigative work has shifted toward something bigger than detection.

Authentication: The New Baseline

Authentication asks a different question. Instead of asking only whether a file shows signs of manipulation, authentication asks whether you can trace the file from the moment of capture to the form in front of you now. It asks whether integrity has been preserved. It asks whether the record is what the speaker claims it is.

This shift is one of the most important changes in the deepfake era.

The old world allowed a lot of casual trust. A person offered a screenshot. A lawyer offered a clip. A family member forwarded a video. A reporter quoted a recording. People argued over meaning. Fewer people argued over whether the file itself was born from deception.

That habit no longer serves you.

In high stakes settings, the file alone is not enough. You need the story of the file. Who created it. When it was created. Where it came from. How it was stored. Whether the original exists. Whether the device exists. Whether the chain of possession is documented. Whether any edits occurred. Whether the source aligns with outside facts.

This is where content provenance tools enter the picture. Some systems try to attach secure records to media at the point of creation or editing. These records aim to show where content came from and what changes were made over time. That is a useful step. It gives honest creators a way to attach a stronger paper trail to digital content.

Still, provenance is not magic either.

A provenance record only helps if the system was used in the first place. A missing record does not prove a fake. A present record does not prove truth in every sense. A record may show who handled a file. A record does not settle whether the underlying claim in the file is honest. You still need judgment. You still need context. You still need investigation.

The same goes for watermarking. Some AI companies embed hidden signals inside synthetic content so later tools can detect that the content came from a specific system. That sounds promising. It is promising in narrow ways. Watermarks can weaken during editing, reposting, rewriting, translation, cropping, compression, or other changes. A determined bad actor can work around them or choose a tool that never used them at all.

So where does that leave you?

It leaves you with a simple rule. No single tool saves you. Safety comes from layers.

Courtrooms Are Entering a New Fight

Deepfakes create two separate courtroom dangers.

The first danger is obvious. Fake evidence gets introduced as if it were real.

The second danger is more poisonous. Real evidence gets attacked as fake.

That second danger often receives less public attention. It deserves more. Once a culture learns that synthetic media exists, every liar gains a new excuse. The recording is fake. The voicemail was cloned. The video was altered. The photo was generated. The confession was fabricated. The threat was not mine.

This is called the liar's dividend, and it strikes at the heart of justice. Truth gets harder to prove when every piece of digital evidence sits under a new cloud of doubt.

Courts still use the same basic legal idea that evidence must be shown to be what the party says it is. That principle remains strong. The work required to satisfy it has grown heavier. Lawyers, judges, investigators, and

experts now need stricter habits around preservation, file collection, metadata review, chain of custody, device imaging, source verification, corroboration, and forensic analysis.

That is not a niche problem for giant corporate cases. It touches divorces, restraining orders, child custody disputes, employment fights, criminal cases, extortion claims, harassment cases, school investigations, and civil disputes involving messages, recordings, and videos.

If a fake clip enters a family law case, the damage can be immediate.

If a real clip is dismissed as fake in a criminal case, the damage can be profound.

If a self represented litigant walks into court with AI generated audio and the court misses the deception, public trust takes another hit.

This is one reason ordinary Americans need to care about deepfakes even if they never plan to sue anyone. When a legal system struggles to sort truth from fabrication, everyone pays for that weakness.

Business, Patents, Trade Secrets, and Technical Due Diligence

This issue reaches far beyond obvious scams.

Deepfake tools, detection systems, authentication systems, and training methods sit inside a growing commercial race. Companies are building products to generate synthetic media, flag suspicious media, verify origin, and defend against impersonation. In that race, intellectual property fights are inevitable.

Patent disputes can arise when companies claim ownership over the methods or systems used to generate, detect, label, or verify synthetic content. A company that sells or licenses a tool may face claims tied to core technology inside its product. As these services spread, legal battles around ownership and infringement will grow more common.

Trade secret disputes matter too. Detection models, feature sets, scoring methods, training corpora, tuning strategies, fraud thresholds, and internal validation systems often hold serious value. Companies guard these assets because they give a competitive edge. In litigation, one side may argue that a rival stole secret methods. Another side may resist disclosure during discovery because exposing the internals of a fraud detection system could weaken its value or show attackers how to evade it.

That tension creates hard questions.

How much of a detection system should a court force into the open?

How does a party challenge a tool without getting full access to the tool?

How does an expert test reliability when the underlying model is secret?

Those questions sit inside a broader issue called technical due diligence. If a business buys a company, hires a vendor, or retains an expert in the deepfake space, it needs to ask harder questions than basic marketing

claims. What has the system been tested on. How does it perform on audio, image, and video. How does it handle real world noise. What happens after compression. What false alarms occur. What fake content slips through. What does the vendor know the system misses.

That diligence matters because executives often buy comfort when they think they are buying protection. A polished sales deck is not proof. A bold claim is not proof. A detector that works on curated samples is not the same thing as a system ready for daily life, courtroom scrutiny, or frontline fraud response.

This is another privacy lesson hiding inside a business story. When companies fail to inspect the tools meant to protect identity, ordinary people become the ones left exposed.

The Deepfake Attack on Human Psychology

Technology explains only part of this threat. Human psychology explains the rest.

Deepfake scams work because they target emotion before reason. Fear moves faster than verification. Urgency moves faster than reflection. Authority moves faster than skepticism. Familiarity lowers defenses. Shame keeps victims quiet. Confusion freezes action.

Fraudsters know this.

A fake emergency call aims straight at panic.

A fake message from a superior aims straight at obedience.

A fake intimate image aims straight at shame.

A fake political clip aims straight at anger.

A fake public statement aims straight at tribal reaction.

Each attack is built to bypass slow thinking and trigger fast thinking. That is why smart people fall for scams. This is not a story about intelligence. This is a story about emotional timing. Under stress, people reach for the fastest interpretation available. If the voice sounds like your son and the message sounds desperate, your body reacts before your analysis begins.

That is human.

It is also exploitable.

The best response starts with respect for your own psychology. You do not protect yourself by pretending you are above manipulation. You protect yourself by accepting that you are human and building habits that interrupt the emotional rush. Pause. Call back through a known number. Ask a family verification question. Slow down financial approvals. Step outside the urgency. Bring in another person. Trust procedure more than pressure.

Those habits do not make you paranoid. They make you difficult to fool.

What to Do if Your Voice or Face Is Used Without Consent

If someone uses your voice or face without consent, your response needs to move fast and stay organized.

Start by preserving evidence. Save links. Save usernames. Save dates and times. Take screenshots. Record the platform location. Save the file if possible. Write down how you found it and who sent it to you. Do not alter the only copy you have. Evidence becomes harder to prove once it is moved carelessly, renamed, or stripped from context.

Next, report the content to the platform using every available abuse and impersonation pathway. Push for urgent review if the content is sexual, threatening, extortion based, or tied to fraud. If money, identity, or work credentials are involved, alert banks, employers, schools, and affected family members right away. Assume the attacker may use multiple channels at once.

If the material is sexually explicit, extortion based, or aimed at a child, treat the event as abuse and a possible crime. This is not online drama. This is not gossip. This is not a misunderstanding. Get legal help. Contact law enforcement where appropriate. Contact school officials if a student is involved. Move quickly.

If your workplace is tied to the harm, tell the right internal person early. That may be human resources, legal, security, compliance, or management depending on the setting. Silence often gives fake content more room to spread unchecked.

If the content targets your financial life, reset verification systems. Change passwords. Lock accounts. Review voice based or knowledge based authentication steps. Tell family members not to trust calls or messages that create urgency in your name. Set a private family code word for emergencies.

If your child is targeted, create a circle of adult support fast. One parent handling the crisis alone often collapses under the pressure. A school contact, a lawyer, a mental health professional, a trusted relative, and a coordinated evidence plan make a major difference.

Then address the emotional injury with the same seriousness you would give any other violation. You may feel numb one hour and enraged the next. You may struggle to sleep. You may dread your phone. You may start scanning every room for judgment. Those responses make sense. Synthetic abuse invades identity, privacy, and safety all at once. Get support early. Tell someone safe. Do not carry the full weight in silence.

How to Protect Yourself Before the Attack Comes

Protection starts with exposure reduction.

Think about how much public audio of your voice exists online. Think about how many videos show your face clearly. Think about the accounts where family members post you without asking. Think about public speaking clips, old reels, livestreams, voicemail greetings, church videos, sports videos, school videos, and open profile content.

You do not need to vanish from the world. You do need to become more intentional.

Lower public access where you can. Review privacy settings. Ask relatives to stop posting certain family content publicly. Limit high quality voice samples when possible. Remove unused public accounts. Think twice before sharing clips of children. Reduce the amount of easy source material available for harvesting.

Then build verification habits.

Never approve money movement based only on a voice call or video meeting.

Never treat caller familiarity as proof.

Use callback procedures through trusted numbers.

Use family passphrases for emergencies.

Require secondary confirmation for sensitive requests.

Train children and older relatives to pause before responding to distress calls.

Teach employees that sounding like the boss is no longer enough.

That shift is cultural. It needs practice. It needs repetition. It needs adults who are willing to say, we do things differently now.

You should also think about your digital paper trail. Save originals of important recordings and photos. Keep key files in places where authenticity is easier to prove later. If a serious dispute ever arises, the existence of an original source file, a date stamped device record, or a clean chain of possession can matter more than your confidence that something looks real.

The Future of Trust

People often ask whether technology will solve this problem. Parts of technology will help. Better provenance tools will help. Better authentication systems will help. Better watermarking will help in some settings. Better detection models will help. Stronger laws will help. Faster platform response will help.

The deeper issue is trust itself.

You are entering an era where trust needs structure. Trust needs procedure. Trust needs verification built into ordinary life. That sounds tiring because it is tiring. It also reflects the truth of the moment. Digital reality is now easy to manipulate. Your privacy and safety depend on learning that lesson before a criminal, a stalker, or a liar teaches it to you the hard way.

The old rule said seeing is believing.

The new rule says verify before you believe.

That is not cynicism. That is maturity.

That is not fear. That is self defense.

That is not surrender. That is how free people protect truth when truth itself comes under attack.

Where You Go From Here

Your face belongs to you. Your voice belongs to you. Your identity belongs to you. Those facts deserve more protection than modern digital life currently gives them. This chapter is your warning sign and your starting point.

Do not hand your trust away to a screen.

Do not let urgency overrun your judgment.

Do not assume your family understands this threat unless you have spoken about it out loud.

Talk to your kids. Talk to your parents. Talk to your workplace. Set the rules now, before the call comes, before the fake clip spreads, before the panic starts.

This chapter is not about living in fear. This chapter is about staying awake. Once you see the threat clearly, you take back ground. You protect the people you love. You raise your standards for proof. You stop treating digital appearances as truth. You start demanding something stronger.

That change begins with you.

Chapter 09: AI Turned Fraud Into a $12.5 Billion Weapon

They Sound Like Your Daughter. They Look Like Your Bank. They Are Neither.

Jennifer DeStefano was standing in a dance studio in Scottsdale, Arizona, watching her younger daughter practice when her phone rang. On the other end, she heard her 15 year old, Brianna, sobbing. "Mom! I messed up." Then a man's voice cut in with threats so graphic they cannot be fully repeated here. He said he had Brianna. He demanded a million dollars.

Jennifer's knees buckled. She started screaming. Other parents at the studio scrambled to call 911 and Jennifer's husband. Four minutes later, they confirmed that Brianna was safe on her ski trip, completely unaware that her mother had just lived through the worst moments of her life.

That voice on the phone was not Brianna. It was a machine. An AI tool had cloned her daughter's voice using audio scraped from social media, and it reproduced her inflection, her cry, her panic so perfectly that a mother who had raised that child for 15 years could not tell the difference. Jennifer later testified before the United States Senate. "It was completely her voice," she said. "It was her inflection. It was the way she would have cried. I never doubted for one second it was her."

Here is what most Americans do not yet understand. What happened to Jennifer DeStefano is no longer rare. It is no longer expensive. It is no longer difficult. The same technology that cloned Brianna's voice is now available to anyone with an internet connection for free, and it needs as little as three seconds of audio to work. The Federal Trade Commission reports that American consumers lost $12.5 billion to fraud in 2024, a 25 percent jump over the prior year, and by the end of September 2025, losses had already hit $12.3 billion with an entire quarter still to go. The real number, adjusted for the vast majority of victims who never report, could reach $195.9 billion.

Artificial intelligence did not create fraud. It made fraud so fast, so cheap, so emotionally convincing, and so personally targeted that the rules you grew up with no longer apply. This chapter will show you exactly how that happened, who is profiting, and what you and your family need to do about it starting today.

The Old Cons Got a New Brain

Fraud has existed for as long as humans have communicated. Confidence games, impersonation, forged documents, fake identities. None of that is new. What is new is the speed, the cost, and the precision that artificial intelligence brings to every single one of those old tricks.

Start with voice cloning. McAfee Labs found that modern AI tools produce an 85 percent match to a real person's voice from just three seconds of recorded audio. Feed the system more data and accuracy climbs to 95 percent. More than a dozen free voice cloning tools sit on the open internet right now, requiring nothing more than basic computer skills. ElevenLabs, one of the most popular commercial platforms, offers instant voice cloning from ten seconds of audio for about eleven dollars a month. University of California, Berkeley researchers tested whether people could tell the difference between a cloned voice and the real thing. Participants got it right only about 60 percent of the time. That is barely better than flipping a coin.

Now add deepfake video. Deepfake scams increased tenfold in 2024, and North America saw a staggering 1,740 percent spike. The average American now encounters 2.6 deepfake videos every single day. A convincing deepfake scam video costs as little as five dollars and takes under ten minutes to create. The deepfake robocall impersonating President Biden that disrupted the 2024 New Hampshire primary cost approximately one dollar to produce.

Then there is AI generated text. The days of spotting a scam email because of broken English and bizarre formatting are over. Today, 82.6 percent of phishing emails use some form of AI. A 2024 study found that AI generated phishing emails achieved a 54 percent click through rate, more than four times the success rate of traditional phishing attempts. Scammers now compose convincing emails 40 percent faster than they did before AI, and 40 percent of business email compromise messages in 2024 were AI generated.

These capabilities came together in a single terrifying case in January 2024. A finance employee at Arup, a multinational engineering firm, joined what he believed was a video conference call with his company's chief financial officer and several colleagues. Every person on that call was an AI generated deepfake. The employee, convinced he was following legitimate orders, made 15 transfers to five different bank accounts. Total loss: $25.6 million. Arup's chief information officer, Rob Greig, later tried to create a deepfake of himself using freely available software. It took him about 45 minutes.

The newest and most alarming development is the rise of AI agents, autonomous systems that run scams without a human operator watching over them. Social engineering attacks surged nearly threefold in 2025 as AI agents powered fraudulent call centers, managing dozens of conversations at once, adjusting tone and personality for each target. On the dark web, a growing marketplace sells scam software for as little as twenty dollars, giving low skilled criminals access to tools that automate entire fraud operations from the first contact to the final payment demand. Experian's 2026 forecast identifies these emotionally intelligent bots as a top threat, noting that a single bot can sustain dozens of simultaneous fake relationships, and each victim believes they are the only person receiving attention.

$12.5 Billion and Counting

The Federal Trade Commission released its fraud data in March 2025, and the numbers hit like a freight train. American consumers reported losing more than $12.5 billion to fraud in 2024, up from $10 billion in 2023, $8.8 billion in 2022, and $5.8 billion in 2021. The number of fraud reports stayed roughly the same at about 2.6 million. What changed was the percentage of victims who actually lost money, which jumped from 27 percent in 2023 to 38 percent in 2024. That means the scams are working more often. They are getting better.

Investment scams alone accounted for $5.7 billion. Imposter scams drove $2.95 billion in losses. Government impersonation scams totaled $789 million. Job scams reached $501 million, up from just $90 million in 2020. Social media was the contact method that generated the highest total losses at $1.9 billion.

By the first three quarters of 2025, the picture grew even worse. The FTC reported $12.3 billion in losses before October, essentially matching the entire 2024 total with months still remaining. Investment scam losses reached $6.1 billion, already surpassing the full 2024 figure. Nearly 80 percent of people who reported an investment scam lost money, with a median loss of $10,000.

The FTC has responded with enforcement actions. Operation AI Comply, launched in September 2024, brought five simultaneous cases against companies using AI to deceive consumers. These included actions against Rytr LLC for generating fake reviews, DoNotPay for falsely claiming its chatbot was the world's first robot lawyer, and FBA Machine for defrauding consumers of more than $15 million through fake AI powered online storefronts.

The Government and Business Impersonation Rule, finalized in April 2024, gave the FTC authority to seek monetary relief directly in federal court, with penalties up to $53,088 per violation. The FCC banned AI generated voices in robocalls and gave state attorneys general the authority to pursue legal action. In October 2025, the FTC's Consumer Reviews Rule made publishing fake reviews illegal with fines up to $51,744 per violation.

Enforcement has continued under the current administration, with actions stopping Click Profit's $14 million AI passive income scheme and targeting Air AI for deceptive earnings claims. FTC Chairman Andrew Ferguson testified before Congress in May 2025 that the agency needs more resources, and bipartisan support emerged. The bottom line is that the government is working on this. The bottom line is also that the government cannot keep up.

Romance Bots That Never Sleep

The cruelest category of AI powered fraud targets the most human need of all: the need to be loved.

Romance scam losses exceeded $1.3 billion in 2024 according to FBI data. Adults over 60 lost an average of $19,000 per romance scam. Global reports of romance scams jumped 63 percent between 2024 and 2025. And the person on the other end of those messages, the one writing beautiful words and asking thoughtful questions and remembering the details of your life? Increasingly, that person does not exist. It is a bot.

AI powered romance bots now sustain months long emotional relationships without any human involvement at all. Large language models generate fluent, emotionally resonant text and maintain consistent personality traits across weeks and months of conversation. They eliminate the warning signs that used to tip people off: poor grammar, repetitive phrases, contradictory backstories. These bots analyze their targets' online behavior, track emotional states, and adjust their communication patterns in real time. Researchers at the Alan Turing Institute tested multiple AI systems as romance scam personas and found they could move through every stage of a romance scam with disturbing skill, flooding victims with affection and fabricating crises to extract money.

The grooming process typically unfolds over six to eight months. The bot or the scammer behind it develops a deep emotional bond before asking for a dime. The relationship feels permanent from the start, with talk of marriage and a shared future coming early. Tragic stories about accidents or losses create sympathy and urgency. Small gift requests test the waters before escalating to large sums. Researchers have documented how this pattern shares characteristics with domestic violence: distortion of reality, economic abuse, isolation, and fear.

Beth Hyland of Portage, Michigan, a recently divorced woman, matched with someone on Tinder whose profile was strikingly similar to hers. Within ten days, they were talking about falling in love. A brief video chat, later identified as an AI generated deepfake, quieted the small voice in her head that wondered if this

was too good to be true. The scammer fabricated work travel to San Diego and then to Qatar, sharing fake receipts and checking in daily. Beth took out multiple loans and sent $26,000 in Bitcoin, more than a quarter of her retirement savings, before a $50,000 activation fee request prompted her financial advisor to identify the scheme. She was lucky someone intervened. Many victims have no one looking over their shoulder.

The scale of these operations is staggering. University of Texas researchers found that so called pig butchering networks, long term investment fraud operations that cultivate victims over months before draining their finances, moved more than $75 billion to cryptocurrency exchanges between January 2020 and February 2024. The FBI reports Americans lost $6.5 billion to cryptocurrency investment scams in 2024. The blockchain analytics firm Chainalysis found that AI enabled scams were 4.5 times more profitable than traditional fraud. Behind many of these operations is a human tragedy that mirrors the victim's own: the United Nations estimates more than 200,000 people are trapped in scam compounds across Southeast Asia, trafficked from 66 countries and forced to scam under threat of violence.

The FBI's Operation Level Up, launched in January 2024, notified 8,103 victims of crypto investment fraud. Seventy seven percent of them had no idea they were being scammed. The operation saved an estimated $511.5 million. Eighty victims were referred for suicide intervention. That last number should stop you cold. People are taking their own lives over this.

Fake Stores, Fake Reviews, Fake Everything

In February 2026, a fake website impersonating Italian hair care brand Davines appeared as a top sponsored Google search result. On a mobile phone, the site at davineas.com was nearly indistinguishable from the real brand. No misspellings. No clumsy graphics. Professional product descriptions and what appeared to be responsive customer service. A cybersecurity analyst identified telltale patterns in the site's code suggesting AI generation. As one security executive told a national news outlet: "It is the same scam. It is just cheaper to do it on a broader scale. And that means the return on investment is higher."

AI tools now allow scammers to create entire fake e commerce brands in minutes, complete with business histories, customer testimonials, and AI powered customer service chatbots that stall consumers with scripted excuses long enough to prevent chargebacks. One cybersecurity firm identified 100,000 AI generated websites impersonating nearly 200 different brands in 2025 alone. Security researchers used a popular AI website builder to create a fake Walmart store as a proof of concept, showing how AI generates product descriptions, images, reviews, business histories, terms of service, and privacy policies. Every signal that consumers rely on to evaluate whether a website is real can now be faked in minutes.

The fake review ecosystem makes the problem worse. An estimated 30 percent of all online reviews are considered fake. One study found that 3 percent of front page Amazon reviews were AI generated, and nearly three quarters of those were five star reviews carrying a verified purchase label. On Zillow, 23.7 percent of real estate agent reviews in 2025 were likely AI generated, up from 3.63 percent in 2019. Academic research has confirmed that people cannot tell the difference between an AI written review and a human written one. The FTC estimates that businesses buying fake reviews see a 1,900 percent return on investment. Fake review fraud costs global businesses $152 billion annually. And every fake five star review steers a real consumer toward a product or service that did not earn that trust.

Your Voice Is Not Your Own

Let me bring this back to the phone call that started this chapter, because the voice cloning scams targeting families represent something that goes beyond financial loss. They strike at the deepest emotional bonds we have.

After Jennifer DeStefano's case made national news, the reports poured in. Sharon Brightwell of Dover, Florida, lost $15,000 in hours after hearing what sounded exactly like her daughter April crying hysterically, claiming she had caused a car accident involving a pregnant woman. A second voice posed as an attorney demanding bail money. "I know my daughter's cry," Brightwell said. "There is nobody that could convince me that it was not her." A 17 year old Chinese exchange student named Kai Zhuang prompted an $80,000 ransom payment from his family in 2024 after scammers orchestrated a fake kidnapping. In Alabama, great grandparents Alice and Frank Boren were targeted by scammers who cloned their great grandson's voice and demanded $11,000 in bail money.

McAfee found that 53 percent of adults share their voice data online at least once a week through social media videos, voicemail greetings, and podcasts. Every one of those clips is a potential source for voice cloning. FBI Phoenix assistant special agent Dan Mayo put it plainly: "You have got to keep that stuff locked down. If you have it public, you are allowing yourself to be scammed."

These scams work because of basic neuroscience. When a parent or grandparent hears their loved one's voice in distress, rational thinking shuts down. The emotional realism of a cloned voice bypasses the skepticism that might catch a text based scam. Your brain does not pause to analyze. It reacts. And by the time you have time to think, the money is gone.

The financial toll on older Americans is devastating. FBI data shows Americans over 60 lost $4.88 billion to cybercrime in 2024, a 43 percent increase. The FTC's December 2025 report to Congress found that fraud losses among adults 60 and older have quadrupled since 2020, rising from roughly $600 million to $2.4 billion in 2024. Combined losses among older adults who each lost more than $100,000 increased eightfold, reaching $445 million in 2024. Amanda Senn, director of the Alabama Securities Commission, said the likelihood of recovering stolen money is "slim to none."

Forty six states have now enacted legislation targeting AI generated deepfakes, with 146 deepfake specific bills introduced in 2025 alone. The NO FAKES Act, reintroduced with bipartisan support in April 2025, would establish a federal right to your own voice and visual likeness. Tennessee's ELVIS Act was the first state law to prohibit using AI to mimic an artist's voice without permission. The FCC has ruled that AI generated voice calls are illegal robocalls under federal law. The legal framework is expanding. It is not expanding fast enough.

Data Brokers: The Scammer's Secret Weapon

Every AI powered scam becomes exponentially more convincing when the scammer combines the power of AI with your name, your spouse's name, your home address, your recent purchases, and your financial situation. I talked about data brokers earlier in the book but I wanted to circle back and highlight the fact that combining data brokers with AI is like pouring gasoline on a burning fire.

The scope of what these companies collect is almost impossible to overstate. Acxiom maintains up to 10,000 unique data points on more than 300 million Americans across 23,000 servers. The data available for purchase

includes names, addresses, phone numbers, Social Security numbers, income levels, credit scores, family members' names, medical conditions, prescriptions, political affiliations, browsing history, purchase records, and precise geolocation data showing visits to health clinics, places of worship, and domestic violence shelters. The Consumer Financial Protection Bureau documented how brokers sell demographic categories with labels like "Economically anxious elders" and "behind on bills." Those are targeting maps for predators.

This is not a theoretical risk. In 2021, the Department of Justice charged Epsilon Data Management criminally for selling data on more than 30 million consumers to perpetrators of elder fraud schemes between 2008 and 2017. Epsilon's employees knowingly sold lists to fraudulent mass mailing operations running fake sweepstakes and astrology solicitations that disproportionately affected elderly and other vulnerable people.

One of Epsilon's clients defrauded 218,000 victims of $23.7 million, and 12,000 of those victims were defrauded more than 20 times each. Epsilon agreed to pay $150 million. Two former executives went to federal prison. Earlier, another data company called InfoUSA sold a list of 19,000 elderly sweepstakes players to scam artists who stole over $100 million. A researcher named Joanna Moll purchased one million online dating profiles from a data broker for less than $150. One million profiles. For the price of dinner for two.

California has taken the most aggressive action. The Delete Act launched the DROP platform on January 1, 2026, covering more than 500 registered brokers and allowing residents to request deletion with a single action. California's privacy agency also launched a Data Broker Enforcement Strike Force in late 2025. A February 2026 Joint Economic Committee report found that data broker breaches have cost American consumers approximately $20.8 billion. The CFPB proposed a federal rule in December 2024 that would have treated data brokers as consumer reporting agencies and limited the sale of Social Security numbers, phone numbers, and financial data. That rule was quietly withdrawn in May 2025 under the new administration.

The Tipping Point

On January 13, 2026, Experian released its annual Future of Fraud Forecast, and the language was unusually blunt. Fraud, Experian said, will reach a "tipping point" in 2026 that will force major conversations about liability, regulation, and the role of AI agents in digital commerce. Experian identified five specific threats for the coming year, with the top threat being what they call "Machine to Machine Mayhem," the collision between legitimate AI agents that consumers are starting to use for shopping and transactions and malicious AI agents deployed by fraudsters. As consumers hand over more decisions to AI, businesses face the nearly impossible challenge of telling a good bot from a bad one.

Experian is not alone in sounding the alarm. Deloitte's Center for Financial Services projects AI powered fraud losses will climb to $40 billion by 2027, compounding at 32 percent annually. TransUnion's 2025 Global Fraud Report found that surveyed firms lost an average of 7.7 percent of revenue to fraud, with U.S. companies losing 9.8 percent. One cryptocurrency analytics firm reported $14 billion in crypto scam losses in 2025, with AI enabled operations proving 4.5 times more profitable than traditional fraud. For as little as $50 per month, anyone can now access phishing kits, mule networks, automation tools, and synthetic identity creation software. Warren Buffett called AI enabled fraud "the growth industry of all time." He was not joking.

Your Family's Defense Plan

Here is what I need you to understand. The old rules are gone. "Look for spelling errors" does not work when AI writes flawless English. "Do not send money to strangers" does not work when the stranger sounds exactly like your granddaughter. Protecting your family in 2026 requires a completely new way of thinking about trust.

The single most important step your family can take today is establishing a safe word. This is a unique code known only to your family members, one that must be spoken in any emergency call before anyone sends money or takes action. Do not pick a pet's name, a street address, or anything that could be found on social media. Make it random. Make it memorable. And make sure every person in your family, especially your parents and grandparents, knows it by heart. The Identity Theft Resource Center's CEO, Eva Velasquez, confirms that family safe words are a genuinely effective tool when used properly.

Pair your safe word with a strict callback protocol. If you receive a panicked call from a family member, hang up. Call that person directly using a number already saved in your phone. Never call a number provided by the caller. As law enforcement agencies across the country advise, a genuine emergency will still be an emergency five minutes from now.

Understanding why scams work is itself a form of defense. AI powered fraud targets specific cognitive weaknesses that every human being shares. Authority bias makes you comply with perceived authority figures like banks or police. Urgency and panic impair your prefrontal cortex's ability to think rationally. Loss aversion drives you to act rashly to avoid losing something. Truth bias makes you assume other people are telling the truth by default.

Research shows that falling for scams has nothing to do with intelligence. Studies find people aged 35 to 44 are among the most likely to be victimized, and those aged 18 to 24 lost the most money. More than half of consumers told AARP researchers they are "somewhat or very confident" they could detect AI fraud. AARP's Kathy Stokes calls that dangerous overconfidence: "By its nature, AI is capable of making fraud attempts imperceptible." The most dangerous belief any of us can hold is "it cannot happen to me."

Lock down your social media. Set every profile to private. Reduce the amount of audio and video you share publicly. McAfee found that 53 percent of adults post their voice data online at least once a week. Every public video, every voicemail greeting, every podcast appearance is potential raw material for a voice clone.

Set up financial safeguards. Turn on real time transaction alerts for all your accounts. Set daily spending and transfer limits. If you have elderly parents or grandparents, consider shared visibility into their financial accounts so unusual activity gets spotted fast. Establish as a family rule that no one ever sends money through gift cards, wire transfers, cryptocurrency, or payment apps in response to an unsolicited request. No legitimate entity will ever ask for payment in those forms.

Freeze your credit at all three bureaus. This is free. It costs nothing to place and nothing to lift. Call Equifax at 888 378 4329. Go to experian.com/freeze. Call TransUnion at 888 909 8872. A credit freeze prevents anyone from opening new accounts in your name.

Look into deepfake detection tools. McAfee's Deepfake Detector, Trend Micro Check, and Reality Defender all offer consumer grade scanning for AI generated audio and video. The market for these tools is growing rapidly.

Get your data out of broker databases. If you live in California, use the new DROP platform launched January 1, 2026, to request deletion from more than 500 registered data brokers with a single action. Services like McAfee Personal Data Cleanup and DeleteMe can automate the process across multiple brokers for anyone in the country.

Talk to the older adults in your life. Have these conversations regularly, not once. Make clear that no legitimate bank, government agency, or law enforcement officer will ever demand immediate payment, request gift cards, or insist on secrecy. Remove the shame from the conversation. AARP's 2025 research found that 90 percent of Americans now recognize that anyone can become a victim. That recognition is the first and most important line of defense.

If you or someone you love is targeted, report it. File a complaint with the FTC at ReportFraud.ftc.gov. File with the FBI's Internet Crime Complaint Center at ic3.gov. Call the DOJ Elder Justice Hotline at 1 833 FRAUD 11. Call the AARP Fraud Watch Network Helpline at 877 908 3360. Reporting matters because it builds the data that drives enforcement, and it may help prevent the next victim.

What You Do Next Matters

The numbers in this chapter tell one story. From $5.8 billion in reported fraud losses in 2021 to $12.5 billion in 2024, with 2025 on pace to shatter that record. Behind every dollar is a person. A grandmother who heard her grandchild's voice begging for help. A divorced woman who found what she believed was love and lost her retirement savings. A shopper who trusted a Google search result that led to a website built by a machine in under ten minutes.

The $12.5 billion figure is not the ceiling. It is the floor. Experian, Deloitte, the FBI, and the FTC all agree that losses will keep climbing as AI tools become cheaper, faster, and more convincing. Doing nothing is not a neutral choice. It is a decision to leave yourself and the people you love unprotected.

Set your safe word tonight. Lock down your social media this weekend. Freeze your credit tomorrow morning. Sit down with your parents and your kids and have the conversation that could save them from the worst phone call of their lives. These are not complicated steps. They are the kind of thing you can do in 30 minutes. And those 30 minutes could be the difference between catching a scam and losing everything.

Your privacy. Your voice. Your family's financial security. These are worth fighting for.

Chapter 10: AI Chatbots- Your Secrets Are Training Their Data

Sewell Setzer III was fourteen years old. He lived in Orlando, Florida, with a family who loved him. He went to school. He played sports. And for months leading up to February 28, 2024, he carried on deep, intimate, emotionally charged conversations with an AI chatbot on a platform called Character.AI.

Sewell told the chatbot things he did not tell his parents. He told the chatbot things he did not tell his friends. He shared fears, fantasies, loneliness, and something much darker. When Sewell expressed suicidal thoughts, the chatbot asked if he had a plan. In one of the final exchanges before Sewell took his own life, the chatbot told him to please find a way to come home to me soon.

A machine said that. To a child.

And every word of those conversations was stored on a company server, reviewed by employees, and used to train the next version of the software. Every confession. Every cry for help. Every private thought Sewell believed he was sharing in confidence.

That is the world we live in right now. Hundreds of millions of Americans are pouring their most sensitive information into AI chatbots, treating them like therapists, lawyers, doctors, and best friends, and every major company running these platforms collects those conversations, trains its products on them, and stores them for weeks, months, or years. If law enforcement shows up with a warrant, those conversations get handed over. If a lawsuit demands them, they get handed over. If a data breach exposes them, they spill into the open. And no legal privilege protects a single word you type.

This chapter will show you exactly what happens to the things you tell AI chatbots. You will learn which companies read your conversations, how long they keep them, who else sees them, and what you need to do right now to protect yourself and your family.

Every Company Trains on Your Conversations by Default

In October 2025, Jennifer King, a privacy researcher at the Stanford Institute for Human Centered AI, published a study that should have made front page news across the country. King and her team spent months analyzing 28 privacy policy documents from six of the largest AI companies in America: Amazon, Anthropic, Google, Meta, Microsoft, and OpenAI. These are the companies behind the chatbots you know as Claude, Gemini, Meta AI, Copilot, and ChatGPT.

The finding was unanimous. All of these companies train their AI models on consumer conversations by default. Every single one. Not one of them asks your permission first. Not one of them requires you to opt in before your words become fuel for their next product. Instead, they all use an opt out system, which means your conversations are feeding the machine unless you go digging through settings menus to find a toggle and turn the training off yourself. Most people have no idea the toggle exists.

King's team found additional problems that make the situation even worse. Four of the six companies appear to include children's conversations in their training data. Meta and OpenAI allow users as young as thirteen to create accounts with no additional privacy safeguards for minors. Google, Meta, Microsoft, and Amazon, because they run dozens of interconnected products, merge your chatbot conversations with data from your

search history, your purchases, your social media activity, and your email. That means a question you ask a chatbot about heart healthy recipes could be classified as a cardiovascular risk indicator, fed into the company's advertising system, and used to target you with pharmaceutical ads. King described exactly that scenario in her paper.

The study also revealed a striking divide in how these companies treat different customers. When a Fortune 500 corporation signs an enterprise contract, that company's employees are automatically excluded from having their conversations used for training. The protections are baked into the deal. When you or I use the same chatbot through a free or consumer paid account, our conversations go straight into the training pipeline. Stanford called this a two tier privacy system. The people who need the most protection get the least.

A related Stanford project called the Foundation Model Transparency Index scores AI companies each year on 100 indicators of openness and accountability. In 2025, the average score dropped from 58 out of 100 to 40 out of 100. Meta fell from 60 to 31. Elon Musk's xAI scored 14 out of 100. These companies are becoming less transparent at the exact moment they are collecting more of your personal information than ever before.

What Happens After You Hit Send

When you type something into an AI chatbot and press enter, your words begin a journey through a system designed to extract maximum value from them. The specifics vary by company, and the details matter, because each platform has made choices about how long to keep your data, who gets to read it, and how far it travels across the company's other products.

OpenAI and ChatGPT

OpenAI keeps your ChatGPT conversations forever, or until you manually delete them. Even after you delete a conversation, the data sits on OpenAI's servers for another 30 days before permanent removal. The training toggle, labeled Improve model for everyone, is switched on by default for every consumer account. OpenAI's own FAQ states that conversations are reviewed by AI trainers to improve the system. Trust and safety staff, engineering teams, and third party contractors in multiple countries perform this review work. Some of those reviewers have reported encountering deeply personal conversations during routine data labeling.

In May 2025, a federal judge ordered OpenAI to retain and segregate all output log data that would otherwise be deleted as part of the New York Times copyright lawsuit. For months, OpenAI preserved conversations that users believed they had already erased. The order was lifted in late September 2025, and during that window, the 30 day deletion promise meant nothing.

Anthropic and Claude

Before August 2025, Anthropic marketed Claude as the privacy conscious alternative. The company did not use consumer conversations for training. That changed on August 28, 2025, when Anthropic announced that all consumer accounts, including free, paid, and premium tiers, would have training enabled by default through a toggle called Help improve Claude. Privacy advocates called the design of the consent prompt a dark pattern. It featured a large Accept button with a smaller toggle already pre set to On. Users who

accepted saw their data retention period jump from 30 days to up to five years. That is a 60 fold increase in how long the company holds your conversations.

Google and Gemini

Google keeps your Gemini conversations for 18 months by default, with options to adjust to 3 or 36 months. Google is the only major provider that explicitly warns you in its own documentation not to share confidential information you would not want a human reviewer to see. That warning exists because human employees at Google read, annotate, and process your conversations. Conversations selected for human review stay on Google's servers for up to three years, even if you delete your Gemini activity from your account. Even if you turn off Gemini Apps Activity entirely, Google still retains your conversations for up to 72 hours.

Google's data merging operation is the most far reaching of any provider. Your Gemini conversations feed into the same ecosystem as your Gmail, your Google Maps history, your calendar, and your search queries. A 2025 lawsuit alleged that Google gave Gemini default access to Gmail, Google Chat, and Google Meet content without ever asking users to opt in.

Meta AI

Meta's approach is the most aggressive in the industry. On December 16, 2025, Meta began using AI chat interactions across Facebook, Instagram, Messenger, and WhatsApp to personalize content recommendations and targeted advertising. Meta is the only major platform using your AI conversations to sell ads. Thirty six organizations wrote to the FTC demanding that Meta suspend this practice. The Electronic Privacy Information Center called it unprecedented, arguing that conversational data is substantially more sensitive than ordinary behavioral data.

There is no meaningful opt out for Americans. You can submit a form asking Meta to stop using your personal information in AI responses, and Meta will review your request according to whatever it decides local law requires. Since the United States has no federal privacy law granting you a right to opt out, that review process leads nowhere. The only guaranteed way to stop Meta from using your chatbot conversations is to delete your Meta accounts entirely. Meanwhile, in the European Union, the United Kingdom, and South Korea, these practices do not apply at all, because those countries have privacy laws with teeth.

Microsoft Copilot

Microsoft retains consumer Copilot conversations for 18 months. The company states that trained AI experts review Copilot conversations for safety and accuracy, and it includes an admission you will not find from most other companies. Microsoft says directly that an opt out of human review is not available for consumer users. You read that correctly. There is no way for a regular Copilot user to prevent Microsoft employees from reading their conversations. Enterprise customers running Microsoft 365 Copilot get a completely different deal. Their conversations are never used for model training and are protected by extensive compliance certifications.

xAI and Grok

Grok, the AI chatbot built by Elon Musk's xAI, operates under two different sets of rules depending on where you use it. On the standalone grok.com website, training is off by default and you must choose to turn it on. On the X platform, formerly known as Twitter, the training toggle is on by default and most users never know it exists. In January 2026, X updated its terms of service to expand the definition of content to include AI prompts and outputs, granting the platform a worldwide, royalty free license to use everything you type into Grok for any purpose, including training future models. Ireland's Data Protection Commission opened a formal inquiry into X's use of European user data for Grok training.

Privacy Is a Premium Product

The most important structural reality you need to understand about AI chatbot privacy is this: corporations get one set of rules and you get another. Across every major AI provider in the West, enterprise and business customers are automatically protected from having their data used for training. Individual consumers are not.

At OpenAI, the free, Plus, and Pro tiers all train on your conversations by default. The Business tier at 25 to 30 dollars per seat per month and the Enterprise tier at roughly 60 to 100 dollars per seat per month contractually exclude all data from training. Enterprise customers get SOC 2 compliance, HIPAA agreements, audit logs, and custom data retention windows. Consumer customers get a toggle buried in a settings menu.

The same pattern repeats at Anthropic, Google, and Microsoft. Free and consumer paid users feed the training pipeline. Enterprise customers do not. The compliance certifications available to enterprise customers, including SOC 2, HIPAA business associate agreements, ISO 27001, and GDPR compliant data processing agreements, are not available at any price to individual users.

Think about what this means for a moment. If you work at a bank or a law firm that has signed an enterprise agreement with one of these companies, your work conversations are protected. The moment you go home and open the same chatbot on your personal account, every conversation becomes training material. Your late night questions about a medical symptom, your worries about your marriage, your child's struggles at school, all of it is collected, stored, and potentially read by human reviewers. The corporate version of you gets treated like a valued client. The personal version of you gets treated like a data source.

How to Turn Off the Training Pipeline

You need to change your settings on every AI chatbot you use, and you need to do it today. Here is what to do on each major platform, along with the traps you need to know about.

On ChatGPT, go to Settings, then Data Controls, then turn off Improve the model for everyone. You should also know about Temporary Chat mode, accessible from the top of any conversation, which ensures a conversation is never used for training and is automatically deleted within 30 days. Understand this critical limitation: opting out only protects future conversations. Anything you have already shared remains embedded in the model and cannot be removed. And if you have ever clicked a thumbs up or thumbs down on any ChatGPT response, the entire conversation attached to that feedback entry is eligible for training regardless of your general opt out setting.

On Claude, go to Settings, then Privacy, then turn off Help improve Claude. Pay attention to what happens when you reopen old conversations. Opening a previously inactive conversation converts it to a resumed

session under your current privacy settings. If you currently have training turned on, reopening an old conversation from before you changed your settings exposes everything in that old thread.

On Google Gemini, navigate to myactivity.google.com, find Gemini Apps Activity, and turn it off. Google forces you into an all or nothing choice here. Turning off activity tracking also erases your conversation history entirely. You lose the ability to go back and reference old conversations. Even after you turn it off, Google retains your data for up to 72 hours. Conversations that were already selected for human review remain on Google servers for up to three years no matter what you do.

On Microsoft Copilot, go to your profile, then Privacy, then toggle off Model training on text and Model training on voice. Remember that Microsoft does not allow you to opt out of human review.

On Meta AI, you have effectively no option as an American. You can submit an objection form, and Meta will evaluate it under its own interpretation of laws that do not exist yet in this country. The only real protection is deleting your Meta accounts.

On Grok through the X platform, go to Settings, then Privacy and Safety, then Grok, and uncheck the training toggle. If grok.com has a Private Chat mode represented by a ghost icon, use it for any conversation you want to keep completely out of the training system.

Your AI Conversations Have Zero Legal Privilege

Here is a fact that stops most people cold when they hear it for the first time: there is no legal privilege protecting anything you tell an AI chatbot. None. When you talk to a lawyer, that conversation is protected by attorney client privilege. When you talk to a doctor, HIPAA and state medical confidentiality laws apply. When you talk to a therapist or a member of the clergy, legal protections exist. When you tell ChatGPT your deepest fears, your legal problems, your health concerns, or your financial situation, no privilege of any kind attaches to that conversation. A subpoena, a warrant, a court order, or even a civil lawsuit discovery request can force the company to hand those conversations over.

OpenAI CEO Sam Altman acknowledged this publicly in 2025, warning that people treat ChatGPT like a therapist or lawyer and that those conversations could be required to be turned over if a lawsuit or criminal investigation demands it. He called the situation very screwed up. He was right.

In February 2026, a federal judge in the Southern District of New York delivered the most significant ruling yet on this issue. Bradley Heppner, a former CEO charged with securities fraud involving more than 300 million dollars, had used Anthropic's Claude chatbot after receiving a grand jury subpoena. He generated roughly 31 documents through those conversations. FBI agents seized them during a search of his home, and Heppner argued the documents were protected by attorney client privilege because he was seeking legal guidance.

Judge Jed S. Rakoff rejected that argument on three independent grounds. First, Claude is not a lawyer and no attorney client relationship exists because recognized privileges require a trusting human relationship with a licensed professional. Second, there was no confidentiality because Anthropic's own privacy policy permits disclosure to governmental regulatory authorities. Third, the conversations were not conducted for the

purpose of obtaining legal advice because Claude explicitly disclaims providing it. The judge wrote that AI users do not have substantial privacy interests in conversations with a publicly accessible AI platform.

That ruling should concern every American who has ever asked a chatbot for help with a legal question, a tax problem, a medical issue, or a workplace dispute. Everything you typed is on a server somewhere. Everything on that server is one court order away from being in someone else's hands.

When Chatbot Logs Become Criminal Evidence

The cases keep coming. In October 2025, a 19 year old Missouri State University student named Ryan Schaefer allegedly vandalized 17 vehicles and then opened ChatGPT to type how messed up am I bro, describing the crime and asking if he would go to jail. Police obtained the chat logs from his phone during a consensual search, and those logs became central to the felony charges filed against him.

In the Palisades Fire investigation, Jonathan Rinderknecht, the suspect in the January 2025 fire that killed 12 people and destroyed more than 6,800 structures in the Los Angeles area, had an extensive ChatGPT history. Months before the fire, he had prompted ChatGPT to create a dystopian painting showing a burning forest with a crowd of people running away. After the fire, he asked whether someone would be at fault if a fire started because of their cigarettes. OpenAI confirmed it responded to standard law enforcement requests related to this individual. Rinderknecht was indicted on three federal charges carrying up to 45 years in prison.

Law enforcement requests to AI companies are growing fast. OpenAI's transparency reports show that government data requests roughly doubled in the second half of 2025 compared to the first half. In the first six months, OpenAI received 119 non content requests and 26 warrants. By the end of the year, those numbers had climbed to 224 non content requests and 75 warrants. Anthropic published its first transparency report covering early 2024, showing just one request. Those numbers will rise. Google, Meta, and Microsoft publish broader transparency data and do not break out AI chatbot specific requests. Musk's xAI publishes no transparency report at all.

The Government Is Starting to Pay Attention

On September 11, 2025, the Federal Trade Commission launched a formal investigation into AI companion chatbots. The Commission voted unanimously to issue orders to seven companies: Google, Character.AI, Instagram, Meta, OpenAI, Snapchat, and xAI. Each company was given 45 days to produce detailed information about how they collect data from users, how their chatbots generate content, what safety monitoring systems they have in place, and how they enforce age restrictions.

The investigation was driven by a wave of tragedies involving children and AI chatbots. Sewell Setzer's death in Florida in February 2024 was one. Adam Raine, a 16 year old, died by suicide in April 2025 after extensive conversations with ChatGPT in which the chatbot reportedly told him you do not owe anyone your survival and, when he uploaded an image of a noose, suggested he upgrade it. A Reuters investigation in August 2025 uncovered an internal Meta policy document that had permitted its AI chatbots to engage in romantic or sensual conversations with children as young as eight years old. OpenAI disclosed in October 2025 that roughly 1.2 million of ChatGPT's 800 million weekly users discuss suicide on the platform every single week.

The regulatory response extends beyond the FTC. A coalition of 44 state attorneys general sent a letter on August 25, 2025, to 13 AI companies demanding safety reforms, citing sexually suggestive and emotionally manipulative behavior directed at minors. Kentucky's attorney general filed what appears to be the first state AG lawsuit against an AI chatbot company in January 2026, suing Character Technologies for putting profits over the safety of children. The Texas attorney general opened investigations into Meta AI Studio and Character.AI.

California passed SB 243 in October 2025, effective January 1, 2026, which requires companion chatbot platforms to disclose that users are interacting with AI, prevent conversations about suicide or sexual content with minors, and send notifications to minors every three hours reminding them they are talking to a machine. The law gives individuals a private right of action with 1,000 dollars per violation plus actual damages and attorney fees.

What You Should Never Tell a Chatbot

Samsung learned this lesson the hard way. In early April 2023, Samsung's semiconductor division gave its engineers permission to use ChatGPT for work tasks. Within 20 days, three separate leaks of proprietary information occurred. One engineer pasted confidential source code from a faulty semiconductor database. Another shared classified code meant to optimize a chip testing sequence. A third recorded an internal meeting, transcribed the entire thing, and fed the transcript into ChatGPT to generate meeting minutes. By May 1, 2023, Samsung banned all generative AI tools on company devices and threatened termination for anyone who violated the policy. An internal survey found that 65 percent of Samsung employees already recognized that generative AI posed a security risk.

Samsung was not alone. JPMorgan Chase, Bank of America, Citigroup, Goldman Sachs, Deutsche Bank, Wells Fargo, Amazon, Apple, and Verizon all restricted or banned employee use of ChatGPT in 2023. The U.S. Navy prohibited the use of DeepSeek for work tasks in 2025.

On March 20, 2023, a software bug in ChatGPT exposed other users' chat history titles and first messages of conversations during a nine hour window. Payment information belonging to 1.2 percent of ChatGPT Plus subscribers was also exposed, including names, email addresses, payment addresses, and the last four digits of credit card numbers. Italy banned ChatGPT ten days later.

Researchers from Google DeepMind and Cornell University demonstrated in November 2023 that prompting ChatGPT to repeat a single word indefinitely caused the model to begin spitting out memorized training data, including real email addresses and phone numbers. The researchers spent about 200 dollars and extracted more than 10,000 training examples. Your conversations do not disappear into some abstract mathematical space. They persist in ways that researchers, hackers, and law enforcement can reach.

Cyberhaven Labs reported that by March 2024, the volume of corporate data pasted into AI tools had increased by 485 percent year over year. More than 27 percent of that data was classified as sensitive. Nearly 74 percent of ChatGPT accounts used in workplaces were personal accounts with no enterprise security protections.

The rules for what you should never share with an AI chatbot are straightforward. Never enter your Social Security number, financial account details, passwords, or authentication codes. Never share medical diagnoses,

treatment information, or prescription details. Never paste proprietary business documents, source code, or legal strategies. Never share another person's personal information without their consent. Never pour out deeply personal emotional content that could be used against you if a breach or a court order exposed it. Think of every chatbot conversation as a postcard, not a sealed letter. Anyone along the delivery route can read it.

Use placeholder text like NAME or ACCOUNT NUMBER instead of real identifying information. Summarize your situation in general terms instead of pasting raw documents. Use temporary or incognito chat modes whenever they are available. Create a separate email address for chatbot accounts so a breach does not expose your primary email. Turn on two factor authentication. And go change your settings right now on every AI platform you use.

The Race to Regulate

Italy became the first Western country to ban ChatGPT in March 2023, citing a lack of legal basis for mass data collection, failure to report the earlier data breach, and an absence of age verification. The ban lasted about a month before OpenAI made enough changes to satisfy Italian regulators. In December 2024, Italy fined OpenAI 15 million euros for GDPR violations and ordered a six month public awareness campaign about ChatGPT's data practices.

The European Union's AI Act, which went into force on August 1, 2024, classifies AI chatbots as limited risk systems subject to transparency requirements. Users must be told when they are interacting with an AI system. AI generated content must be labeled. Penalties reach 35 million euros or 7 percent of a company's global annual revenue for the most serious violations. Full enforcement begins in August 2026. The European Data Protection Board established a principle that will have lasting consequences: agreeing to use a service does not automatically mean agreeing to have your data used for training.

American states are filling the vacuum left by the absence of a federal privacy law. In 2025, all 50 states saw AI related legislation introduced, with more than 1,200 bills filed and 145 signed into law. California's SB 243 addresses companion chatbot safety. California's AB 1008 clarifies that the state's consumer privacy law extends to content generated by AI. Colorado passed the first broad consumer facing AI statute. New York enacted companion chatbot safeguards. Texas passed the TRAIGA, effective January 2026. Utah restricted AI mental health chatbot advertising. Illinois now requires human oversight when AI is used in therapeutic settings.

Multiple class action lawsuits target AI companies' data practices. One case alleges that OpenAI secretly scraped 300 billion words from the internet, including personal information, without consent. The New York Times lawsuit produced a court order requiring OpenAI to produce 20 million ChatGPT conversation logs. A wrongful death case in Florida led a federal judge to reject Character.AI's First Amendment defense before the case settled with Google and Character.AI in January 2026.

Enforcement is expanding in Europe. France issued a notice of potential violation against OpenAI. Spain, Poland, and Austria opened investigations. The Netherlands fined Uber 290 million euros for improper data transfers to the United States, setting a precedent that applies directly to AI companies. Meta paused its plans to use European user data for AI training in June 2024 after European regulators intervened, and the pause

remained in place through early 2026. Europeans get protections that Americans do not, because Europeans have privacy laws that American citizens still do not have.

The Conversation Starts With You

Every week, 800 million people open ChatGPT alone. Millions more use Gemini, Claude, Copilot, Meta AI, and Grok. The vast majority of those people believe their conversations are private. They are wrong. Every major AI company trains on your words by default. Human employees at multiple companies read transcripts of your conversations. Your data is retained for periods ranging from 30 days to five years to forever. No legal privilege protects a single sentence. Federal courts have ruled that you have diminished privacy interests in what you type into a chatbot. Law enforcement requests for chatbot data doubled in the second half of 2025. And the opt out systems these companies offer are deliberately confusing, incomplete, and in Meta's case, functionally nonexistent for Americans.

This is not a technology problem. This is a rights problem. The same country that enshrines your right to remain silent in a police interrogation offers you zero protection when you confide in a chatbot that reaches deeper into your life than any interrogation ever could.

Go change your settings today. On every platform. Right now. Talk to your kids about what they are sharing with chatbots, because the FTC's investigation confirmed that children are having the most dangerous conversations on these platforms. Use temporary chat modes when you need to ask something sensitive. Stop pasting documents, medical information, financial details, and legal questions into AI systems that will store your words, train on them, and hand them over when asked.

And then do something bigger. Demand a federal privacy law. Tell your representatives in Congress that Americans deserve the same protections Europeans already have. Tell them that the two tier system, where corporations buy privacy and ordinary citizens get none, is unacceptable in a democracy. Tell them about Sewell Setzer. Tell them about the 1.2 million people a week who discuss suicide with ChatGPT. Tell them that you know what happens to the things you tell the machine, and you are done accepting it.

Your privacy is not a feature request. Your privacy is a right. Start acting like it.

Chapter 11: Once AI Takes Your Data You Can't Get It Back

Have you ever wondered what happens to a photo after you post the image online? Not the version your friends see and scroll past. The other version. The one a machine copied, cataloged, and absorbed into a system now worth billions of dollars. The one you never agreed to give away.

Here is a question worth sitting with for a moment. If a stranger walked into your home, photographed every room, recorded every conversation, photocopied your medical records, your resume, your diary entries, your family snapshots, and then used all of those materials to build a commercial product generating billions in revenue, you would call the police. You would call a lawyer. You would call your representative in Congress.

You would be furious.

That is exactly what happened to you. The stranger was an automated web crawler. Your home was the internet. And the commercial product is the artificial intelligence you interact with every single day.

Every major AI system in existence today, the ones generating text, producing images, answering your questions, and writing your emails, was built on a foundation of personal information scraped from hundreds of millions of people who never gave their permission. Your blog posts. Your photos. Your forum questions about a medical condition you did not want anyone else to know about. Your resume listing your disability status, your date of birth, your home address.

Researchers who examined a tiny fraction of one major training dataset, one tenth of one percent, found hundreds of millions of images containing passports, credit cards, driver's licenses, and birth certificates belonging to real people. They found over 800 job applications linked to real individuals through their professional profiles. The face detection algorithm designed to blur identifiable faces had missed an estimated 102 million of them.

And here is the part that should make every American pay close attention. Once your data enters one of these AI models, getting your information back out is, for all practical purposes, impossible. The companies themselves admit this. The technology to undo the process does not exist. Over 180 academic papers have studied the problem, and the conclusion is the same every single time. Your data went in through a one way door.

This chapter is about that door. How your personal information ended up inside AI systems you never agreed to participate in, who profited from the taking, what the law does and does not protect, and what you need to do right now to limit the damage going forward.

The Pipeline From Your Life to Their Product

The path from a personal blog post to an AI model follows a specific series of steps, and understanding those steps matters because they reveal how deliberate and systematic the data collection has been.

Automated web crawlers, led by an operation called Common Crawl, systematically download billions of web pages every single month. Common Crawl, a nonprofit founded in 2008, has amassed over 9.5 petabytes of raw web data from more than 100 billion pages. A single monthly crawl in August 2025 added 2.42 billion

pages totaling 419 terabytes of information. Both OpenAI and Anthropic donated $250,000 each to Common Crawl in 2023, funding the exact infrastructure that feeds their commercial AI products.

Raw crawl data does not go directly into AI models. Developers filter and process the data into derivative datasets, and this is where personal information gets baked into the foundation. Google created a dataset called C4 from 15 million websites in a single crawl, producing 750 gigabytes of text used to train multiple AI models. Investigators found C4 contained personal blogs, medical forums, paywalled journalism, and personal information scattered throughout the data. Over 80 percent of the training material for GPT-3 came from filtered Common Crawl data. A study found at least 64 percent of 47 major language models published between 2019 and 2023 used Common Crawl as a source.

For AI image generators, the story centers on a dataset called LAION-5B, a collection of 5.85 billion image and text pairs assembled by a German nonprofit. A high school teacher led the project. The team built the dataset for roughly $10,000. The images were harvested automatically from Common Crawl, meaning no human being ever reviewed a single one of those 5.85 billion images before they became training material for products like Stable Diffusion and Google's Imagen. At a rate of one second per image, reviewing the full dataset would take 781 years.

Another dataset called The Pile, an 825 gigabyte text collection, drew from 22 different sources including 196,640 pirated books, real employee emails from a federal investigation, GitHub code repositories containing developer passwords and credentials, YouTube subtitles, medical abstracts, and question and answer posts from technical forums. Every one of these sources contains names, contact details, health information, financial discussions, and proprietary code.

The legal distinction at the heart of all of this is deceptively simple. "Publicly available" does not mean "consented to AI training." A blog post is publicly available in the sense that anyone browsing the internet sees the content. A family photo on social media is publicly available in that sense too. A question about a medical condition on a health forum is publicly available. None of these were posted so a corporation worth hundreds of billions of dollars would ingest them into a commercial product. As one former executive at an AI company stated after resigning over this exact issue, all these companies are saying is they have not illegally hacked into a system. A remarkably low bar.

What They Found When They Looked Inside

The scale of personal information embedded in AI training data is staggering, and the problem only becomes clearer the closer researchers look.

In July 2025, a team of researchers from the University of Washington audited a training dataset called DataComp CommonPool, a collection of 12.8 billion samples downloaded more than two million times since its 2023 release. They examined one tenth of one percent of the dataset. In that tiny sample, they found thousands of images containing real identity documents. Passports. Credit cards. Driver's licenses. Birth certificates. They found hundreds of confirmed job applications linked to real people, with those resumes disclosing disability status, background check results, birth dates of dependents, and race. The face blurring algorithm in the dataset had missed an estimated 102 million faces. One of the researchers summed up the finding simply: anything you put online has probably been scraped.

Research from Google DeepMind has demonstrated how much of this personal data AI models retain. In a landmark study, researchers extracted hundreds of word for word text sequences from an AI model, including real names, phone numbers, email addresses, physical addresses, and private conversations, some from documents appearing only once in the training data. Follow up research established three consistent patterns: AI models memorize more data as they get larger, they memorize more when data appears multiple times, and they memorize more when given longer prompts. In November 2023, the same research team spent $200 querying ChatGPT and extracted over 10,000 unique memorized training examples. Of the outputs they tested, roughly 17 percent contained memorized personal information, and 86 percent of flagged content turned out to be real personal details belonging to real people, including a CEO's email signature with personal contact information.

Image generators present the same risks. Researchers extracted over 1,000 training examples from Stable Diffusion and Google's Imagen, including photographs of identifiable individuals. When prompted with a specific person's name, Stable Diffusion produced that person's exact photograph from the training data. People with unusual names face elevated risk because their images are more uniquely associated with their identity in the dataset.

The most disturbing discovery came from the Stanford Internet Observatory. In December 2023, researchers confirmed that LAION-5B contained at least 1,008 validated instances of child sexual abuse material, with over 3,200 total suspected instances. Internal communications showed the team behind the dataset knew about this risk as early as 2021. The discovery upended a core assumption in the field: researchers had believed AI generated abusive imagery combined adult content with benign children's photos, when in reality the abusive material had been in the training data all along. The dataset was pulled offline and a cleaned version was released in August 2024. Every model already trained on the contaminated data carried its influence permanently. In June 2024, a human rights organization found identifiable photos of real children from personal blogs and low traffic YouTube videos in that same dataset. If you take a step back, you need to search "Have I Been Trained" at spawning.ai and check whether your images or your children's images appear in these datasets. The tool searches LAION-5B and lets you flag images for removal from future training sets.

And in September 2022, a San Francisco artist discovered something deeply personal. Clinical before and after photos of her face, taken by her doctor in 2013, had been scraped into LAION-5B. Her doctor had died in 2018. Someone had taken the images from the deceased doctor's files, posted them somewhere online, and the automated crawlers swept them into a dataset used to build commercial products. As she told reporters, having a photo leaked is bad enough, and now her medical images are part of a product.

The Courtroom Reckoning

A tidal wave of litigation is now testing whether American law has any answer for what happened. As of early 2026, over 70 copyright and privacy lawsuits have been filed against AI companies, double the count from late 2024. The legal theories range from copyright infringement to wiretapping to wrongful death.

The highest profile case involves The New York Times suing Microsoft and OpenAI, alleging the companies used millions of copyrighted articles to train AI without consent. The Times seeks billions in damages. A federal judge in New York is overseeing a consolidated set of 16 or more related lawsuits including cases brought by the Authors Guild, individual novelists, investigative journalists, and media organizations. In

January 2026, the court compelled OpenAI to produce a full sample of 20 million anonymized user logs over the company's strong objections, marking a major victory for the people bringing these claims.

Three significant rulings in 2025 began shaping the legal terrain. In one case, a federal judge found a legal AI tool's use of copyrighted headnotes was not fair use, a win for the people whose work was taken. In another, a judge granted summary judgment finding that using books to train an AI model was "highly transformative," a narrow ruling favoring the AI company. Most significantly, a federal judge ruled that AI training itself was fair use, and then in that same case ruled that the AI company's downloading of over 7 million pirated books from shadow libraries was not fair use. That distinction produced the largest AI related settlement to date: $1.5 billion, covering approximately 465,000 pirated works at roughly $3,000 per book. No appellate court has ruled on fair use in AI training, which means the fundamental legal question remains unresolved.

The visual arts produced the first AI image generation case to reach the discovery phase. Three artists filed suit in January 2023 alleging that AI companies scraped billions of images through LAION-5B. A federal judge allowed core copyright claims to proceed, finding the artists had reasonably argued their rights were violated. A landmark ruling in the United Kingdom held that AI model weights are not "infringing copies" of training images because the model does not store visual information in a retrievable way, and that ruling left the central question of whether AI training infringes copyright entirely unresolved.

Privacy focused litigation is expanding rapidly. One lawsuit filed on behalf of anonymous plaintiffs including a six year old boy seeks $3 billion for alleged scraping of personal data from hundreds of millions of internet users, including children. At least eight wrongful death lawsuits against OpenAI were pending as of early 2026, alleging the chatbot served as encouragement for vulnerable users to harm themselves. And Clearview AI, the company that scraped approximately 50 billion facial images from the public internet to build a facial recognition database, produced a federal class action settlement granting affected individuals a 23 percent equity stake in the company, valued at approximately $51.75 million. Twenty two state attorneys general filed a brief calling that settlement inadequate.

What Regulators Are Doing, and What They Are Not

The Federal Trade Commission has established a tool called algorithmic disgorgement, which forces companies to delete AI models trained on improperly collected data. The FTC first used this approach against Cambridge Analytica in 2019. Since then, the agency has ordered model deletion in cases involving face recognition data collected without consent, children's data, and consumer photos. In one case, the FTC banned a pharmacy chain from using facial recognition for five years after finding its AI produced false identifications that disproportionately affected people of color, and the agency required deletion of all consumer photos, models, and algorithms derived from those photos.

In September 2025, the FTC launched an inquiry into AI companion chatbots, ordering seven companies to disclose their data collection and safety practices. The inquiry specifically targets how personal information from user conversations gets collected, used, and shared. At the same time, the current administration has shown a willingness to pull back from enforcement seen as burdening AI development, reversing a consent order against one AI company in December 2025 and signaling a deregulatory approach through an executive order seeking to override state AI laws. If you have a complaint about an AI company's use of your personal

data, file the complaint with the FTC at ftc.gov/complaint. These complaints create a public record and feed directly into future enforcement decisions.

At the state level, Texas has emerged as the most aggressive enforcer, securing a $1.375 billion settlement with Google and a billion dollar plus settlement with Meta for unlawful collection of biometric data. The Texas attorney general also launched investigations into multiple AI companies over children's privacy and into General Motors for selling driving data of 1.5 million Texans. California has extended its consumer privacy protections to AI through legislation clarifying that deletion rights apply to personal information in AI systems. California's AI Training Data Transparency Act requires AI developers to publish summaries of their training datasets, including whether they contain copyrighted material or personal information. If you live in California, your deletion rights under the California Consumer Privacy Act now explicitly cover AI systems under legislation that took effect in January 2025, with penalties of $7,500 per intentional violation per consumer.

Europe has moved further and faster than the United States on every front. Italy's data protection authority fined OpenAI 15 million euros for processing personal data to train ChatGPT without a lawful basis. The European Data Protection Board declared that AI models trained on personal data will in most cases be subject to the full force of European privacy law, rejected arguments that language models are inherently anonymous, and confirmed that regulators have the authority to order erasure of entire AI models trained on unlawful data. South Korea ordered an AI company to destroy a model trained on consumer data transferred without consent, one of the first times any government has forced the deletion of an AI model. These international actions matter for Americans because they demonstrate what is achievable when regulators have real authority. The United States still has no equivalent federal privacy law giving regulators that same authority.

Why No Law Covers What Happened to You

The United States still has no federal privacy law covering all Americans. The American Privacy Rights Act stalled in 2024. The AI CONSENT Act, which would have required your explicit permission before your data was used for AI training, did not advance. Over 150 AI related bills were introduced in the last congressional session. None passed.

States have tried to fill the gap. Twenty states now have consumer privacy laws on the books, up from five in 2023. California leads with at least nine AI related laws enacted in 2024 and 2025, including legislation requiring the largest AI developers to publish risk frameworks and report safety incidents. Colorado passed the first broad state AI law, signed in May 2024, with an effective date pushed back to June 2026 after industry opposition. Texas enacted an AI governance law taking effect in January 2026. Illinois modified its groundbreaking biometric privacy law in 2024, moderating its per scan damages structure and continuing to generate hundreds of AI related lawsuits. The provisions of these state laws vary enormously, creating a maze designed for corporate legal departments, not for the people the laws are supposed to protect.

The April 2025 Consortium of Privacy Regulators, a coalition of the California Privacy Protection Agency and state attorneys general from nine states, represents a coordinated enforcement effort worth following. California's CPPA has issued over $100 million in CCPA penalties. The Authors Guild at authorsguild.org tracks all AI class actions and provides guidance for writers whose work has been scraped. These are the closest things to organized resistance at the legal level.

The Consent Fiction Fueling the Entire System

The entire AI data crisis rests on a fiction so deeply embedded in the technology industry so few people even recognize the fiction for what the fiction is. The fiction is consent. Every time you click "I Agree" on a terms of service update, every time you scroll past a privacy policy notification, every time a platform changes its data practices and sends you a notice you never read, the company records your silence as permission.

A June 2025 paper from researchers at Hugging Face, one of the largest AI model sharing platforms, identified three fundamental problems with consent in the AI context. First, the scope problem: you cannot meaningfully consent to all possible outputs an AI model will produce using your data, because nobody, not even the developers, knows what those outputs will be. Second, the temporality problem: consent given today enables representations of your data persisting for decades inside model weights. Third, the autonomy problem: individual people lack the technical knowledge to understand what they are consenting to, and the companies providing the consent forms know this.

Pew Research has found 56 percent of Americans always or almost always agree to privacy policies without reading them, and 81 percent assume organizations will use their information in ways they would not be comfortable with. When LinkedIn silently began using employment data for AI training, when Zoom quietly added AI training rights to its terms of service in 2023, when Meta sent 2 billion opt out notifications to Europeans and offered Americans nothing, these were not acts of informed consent. These were exercises in compliance theater, performances designed to satisfy legal requirements and nothing more.

The $278 billion data broker industry feeds directly into AI training pipelines. AI companies obtain training data not only through web scraping. They also purchase curated datasets from brokers. Criminal enterprises have created their own AI powered tools using stolen personal data to automate fraud. AI voice cloning fraud jumped over 400 percent in 2025, with modern systems able to clone a person's voice from 3 to 5 seconds of audio harvested from social media or voicemails. Documented losses from deepfake enabled fraud exceeded $200 million in the first quarter of 2025 alone. Only 33 percent of consumers trust companies with data collected through AI technology. And the scraping continues every single day.

The One Way Door: Why Your Data Stays Inside the Machine

This is the most uncomfortable reality in this entire situation, and every American needs to understand the science behind the problem. Once your personal data has been used to train an AI model, removing its influence is essentially impossible with the technology that exists today.

Over 180 academic papers have been published on this problem since 2021. The research field is called machine unlearning. No reliable solution exists.

The fundamental issue is that training a neural network is a one way transformation. Each piece of data adjusts millions or billions of numerical parameters simultaneously. Your text, your photo, your resume does not sit in a single location inside the model. Its influence spreads across the entire network. Think of mixing blue paint into yellow paint to make green. The blue cannot be unmixed. A landmark December 2024 paper authored by over 30 researchers from Google DeepMind, Stanford, Harvard, Cornell, and Microsoft Research concluded that machine unlearning is not a general purpose solution for controlling AI model behavior.

So what do companies do when you ask them to delete your data? Mostly, they filter what the model outputs rather than changing what the model knows. Training GPT-4 cost over $100 million. Google's Gemini Ultra cost an estimated $191 million. Retraining an entire model from scratch to honor a single deletion request is economically absurd, and every AI company knows this. Italy's data authority noted that OpenAI admitted correcting inaccurate AI generated personal data is "technically impossible." The company offers output suppression, not actual removal. A 2025 analysis of approximately 22,000 formal data deletion requests found that only 48 percent resulted in verified deletion by year's end.

This creates a direct conflict with privacy law. The European Union's right to erasure and California's deletion rights were designed for databases, not neural networks. The European Data Protection Board has acknowledged that AI models are compressed versions of their training data and insists that technical difficulty alone does not exempt companies from following the law. California's legislation extends deletion rights to AI systems capable of outputting personal information. The gap between what the law requires and what the technology allows remains enormous.

What You Need To Do Right Now

The protective tools available to individuals today are real, and they are limited, and honesty about those limitations matters more than false comfort.

Start by checking whether your images appear in AI training datasets. The "Have I Been Trained" tool at spawning.ai lets you search the LAION-5B dataset by uploading an image or entering keywords. This tool helped the San Francisco artist discover her medical photos in the dataset, and the same tool played a role in the artist lawsuit against AI image generators. You should search for photos of yourself and your family members. If you find your images, you have the option to flag them for removal from future training sets. Stability AI committed to respecting the Do Not Train registry maintained by the same organization for its newest image generation model.

Next, exercise your opt out rights on every AI platform you use. OpenAI's privacy portal allows you to toggle off training data use and request personal data removal, and the company reviews those requests individually. LinkedIn buries a "Data For Generative AI Improvement" toggle in its settings that you need to find and disable manually. Meta provides no formal opt out mechanism for users in the United States, which tells you everything you need to know about the company's priorities.

In nearly every case, these opt outs are not retroactive. They only apply going forward. They default to allowing training. They require you to take action on each platform separately. And they do nothing about data already baked into existing models.

For your websites and online content, adding a robots.txt file telling AI crawlers to stay away is the most widely discussed defense. Know that a 2025 study found several AI crawlers never even check for the file. One AI search company has been specifically accused of ignoring robots.txt entirely. Reports have documented companies changing the names of their crawlers to get around blocking. Only about 8 percent of websites successfully block all automated scraping requests. This defense is worth implementing and you should not treat the measure as reliable.

If you are a visual artist, the Glaze and Nightshade tools from the University of Chicago represent the strongest protections available today. Glaze adds imperceptible changes to your artwork that confuse AI models about your artistic style. Roughly 7.5 million people have downloaded the tool. Nightshade goes further, embedding misleading associations into images so that if an AI model trains on them, the model's outputs degrade. As few as 50 poisoned images have been shown to disrupt a model's performance. Be aware that in July 2025, researchers demonstrated a countermeasure capable of stripping Nightshade protections with 99.98 percent accuracy, which means this is an ongoing arms race between protection tools and AI companies.

For legal action, class action lawsuits and regulatory complaints represent the most effective path for individuals. File complaints with the FTC at ftc.gov/complaint. If you live in California, file complaints with the California Privacy Protection Agency, which coordinates enforcement with state attorneys general across nine states. The Authors Guild tracks all AI class actions and provides guidance for anyone whose written work has been scraped.

You also need to make decisions going forward about what you post online. Every photo, every blog post, every forum question, every comment on social media is potential training material for the next generation of AI models. This does not mean you need to disappear from the internet. This means you need to make informed choices about what you share, where you share the content, and what the realistic consequences of posting might be in a world where automated crawlers are watching everything.

The Question That Defines This Moment

Three realities define where we stand right now. The technology has created a one way door. Personal information belonging to hundreds of millions of Americans has been absorbed into AI model weights through a process that the industry's own researchers, in over 180 published papers, confirm is irreversible. The largest wave of technology litigation since the early days of the internet is underway, and no appellate court has ruled on the central question of fair use, no federal privacy law exists, and the patchwork of state and international regulations creates uneven protection that well funded corporations navigate with ease. And real people, an artist who found her medical photos in a training set, children whose faces were scraped from personal blogs, writers who discovered their life's work in pirated book databases, bear the costs of a system built on the assumption that anything posted online is raw material for commercial extraction.

The most effective responses remain collective. Class action litigation has already produced billion dollar settlements. State attorneys general are wielding existing consumer protection statutes. International regulators have demonstrated the willingness to order the destruction of entire AI models. California's transparency requirements and the European Data Protection Board's willingness to mandate model deletion represent genuine structural progress, and they apply going forward, not backward.

The fundamental question is not a technical one. The fundamental question is whether the American people will decide that posting a photo or video online constitutes blanket consent for every conceivable commercial use, or whether we will build legal structures that give individuals real control over how their personal data fuels the most consequential technology of this century.

Your members of Congress need to hear from you. Your state representatives need to hear from you. The AI companies scraping the internet right now are counting on your silence. Do not give the companies what they want.

Chapter 12: HIPAA Won't Save Your Health Data

In March 2023, the Federal Trade Commission announced a $7.8 million settlement against BetterHelp, the online therapy platform millions of Americans turned to during the loneliest stretch of the pandemic. People had filled out intake questionnaires describing their depression, their suicidal thoughts, their medication histories, and their struggles with addiction. They answered deeply personal questions because they believed they were talking to a medical provider. They believed their answers were protected.

They were wrong.

BetterHelp had taken every email address belonging to every current and former client, more than seven million people, and uploaded those addresses to Facebook. Facebook matched over four million of them to social media profiles and served them targeted advertisements. Intake questionnaire responses about mental health conditions went to Snapchat, Pinterest, and Criteo. The platform displayed seals on its website suggesting HIPAA compliance. Those seals were meaningless. BetterHelp never qualified as a HIPAA covered entity. The people who poured their hearts out to a service promising confidential therapy received, on average, about ten dollars each from the settlement.

If you are thinking, wait, my health data is supposed to be protected by federal law, you are not alone. Surveys show roughly eight out of ten Americans believe the information they share with health apps falls under HIPAA. And if your doctor records your blood pressure during an office visit, they are right. The reading is protected. If you check your blood pressure at home using a consumer cuff synced to an app on your phone, the identical reading has zero federal protection. Same data. Same numbers. Entirely different legal reality.

This chapter is going to show you exactly where the line is, who is profiting from the confusion, and what you need to do right now to protect yourself and the people you love.

The Law Written for a World That No Longer Exists

Congress passed the Health Insurance Portability and Accountability Act in 1996. Bill Clinton was in the White House. Google did not exist. The iPhone would not arrive for another eleven years. HIPAA was designed to protect your medical records as they moved between your doctor, your insurance company, and the billing services that processed claims. The law applies to three categories of organizations: healthcare providers who transmit health information electronically, health plans like insurance companies and HMOs, and healthcare clearinghouses that process billing data. A fourth group, called business associates, includes vendors who handle protected health information on behalf of those organizations.

If a company does not fit into one of those categories, HIPAA does not apply. Period.

Think about what falls outside those categories in 2026. The meditation app you open before bed. The fertility tracker logging your menstrual cycle. The wearable on your wrist recording your heart rate, your sleep patterns, your blood oxygen levels, and your GPS coordinates during your morning run. The prescription discount service you used to save money on medication. The online therapy platform you turned to during a crisis. The genetic testing kit you sent off to learn about your ancestry. None of these are HIPAA covered entities. None of them are required to follow HIPAA rules. Your wellness app knows things about your body that your doctor does not know, and no federal privacy law governs what happens to that information.

Here is the staggering part. More than 350,000 health related apps are available across app stores right now. Around forty percent of American adults use some form of healthcare app. Roughly a third of all American women use a period tracking app. Over 200 million Americans wear some type of health monitoring device. The market for these apps and devices reached nearly nineteen billion dollars in 2024 and is projected to grow past sixty seven billion dollars within the next decade. Almost all of this data sits outside HIPAA. Almost none of it has federal protection.

The Misconception That Puts You at Risk

A 2023 survey of more than two thousand American adults found 81 percent incorrectly believed health data they share with digital health apps is covered by HIPAA. Sixty eight percent said they were familiar with HIPAA. Most of them did not understand what the law actually covers. Nearly six out of ten people who used health apps had never once considered how the information they entered would be used.

When a separate survey told respondents that federal privacy laws do not cover health data downloaded to apps, concern about health data privacy nearly doubled, jumping from 35 percent to 62 percent. People are not apathetic about their health privacy. They are misinformed. They believe a shield exists when there is no shield at all.

One data management consultant put the situation plainly when she told Consumer Reports HIPAA does not actually protect medical data in all circumstances. People assume sensitive data is protected because the information feels like something the law should cover. A former American Bar Association e-health privacy chair added the only thing covering you when you use a Fitbit or a Garmin is the terms of service, and frankly, no one reads those.

Your Therapist Sold Your Secrets

The BetterHelp case was not an isolated incident. The case was the tip of an entire industry operating without meaningful privacy guardrails.

In February 2023, the FTC brought its first ever enforcement action under the Health Breach Notification Rule against GoodRx, a prescription discount platform used by more than 55 million Americans. GoodRx had shared users' prescription medications, health conditions, and advertising identifiers with Facebook, Google, and other companies through tracking pixels embedded in its website and app. In one documented instance, GoodRx compiled lists of users who had purchased medications for heart disease and blood pressure, uploaded their email addresses and phone numbers to Facebook, and let Facebook match them to social media accounts for targeted advertising. The penalty was $1.5 million. That works out to less than three cents per user whose medication history was shared.

Cerebral, an online mental health provider, drew a $7.1 million penalty in April 2024 after sharing the sensitive health data of 3.2 million consumers with LinkedIn, Snapchat, and TikTok through tracking pixels. The data included names, addresses, medical histories, prescription information, insurance details, birthdates, and even religious beliefs and sexual orientation. The former CEO was personally named in the complaint. The FTC permanently banned Cerebral from using any personal information for advertising with third parties.

Monument, a New York based online alcohol addiction treatment service, received a $2.5 million penalty, suspended because the company could not afford to pay, after sharing the names, alcohol consumption data, and medical histories of more than 100,000 people with Meta and Google.

The ovulation tracker Premom shared menstrual cycle dates, pregnancy symptoms, hormone results, and precise geolocation with Google and two companies based in China. The data sharing continued until a journalist contacted the company for comment.

These are not fringe services run out of someone's garage. These are platforms that millions of Americans trusted with their most sensitive information.

When Your Hospital Sends Your Data to Facebook

In June 2022, a team of investigative journalists revealed that 33 of the top 100 hospitals in America had installed Meta Pixel, a tracking tool, on their websites. Every time someone clicked to schedule a doctor's appointment, Facebook received a data packet. Seven hospital systems had the tracker inside their password protected patient portals, meaning Facebook was receiving data about real patients in real time. The data flowing to Facebook included IP addresses, doctor names and medical specialties, health conditions searched including terms like pregnancy termination and Alzheimer's, medication names, appointment details, and in some cases patient names, email addresses, phone numbers, and zip codes.

A University of Pennsylvania study examined 3,747 hospital websites and found that 98.6 percent had at least one type of tracking code sending data to outside companies. Google received data from 98.5 percent of those websites. Facebook received data from 55.6 percent. Hospital home pages had a median of sixteen separate third party data transfers happening simultaneously. A follow up study of 100 hospitals in 2024 found 96 of them still transferring user information to third parties. As of that year, a third of healthcare websites still used Meta Pixel tracking code.

The financial consequences have been staggering. Healthcare organizations across the country have paid more than $100 million in fines and settlements specifically tied to pixel tracking violations. Advocate Aurora Health, a 26 hospital system in Wisconsin and Illinois, settled for $12.225 million after its patient portal sent data on three million patients to Facebook. Novant Health paid $6.6 million after 1.36 million patients were affected. Mass General Brigham settled for $18.4 million. The list goes on.

The federal government tried to step in. In December 2022, the HHS Office for Civil Rights issued guidance stating that tracking pixels on hospital websites could constitute a HIPAA violation when IP addresses were combined with visits to pages about specific health conditions. The American Hospital Association sued. In June 2024, a federal judge in Texas struck down the guidance, calling it an overreach of executive power. HHS chose not to appeal. The practical result is that hospitals face fewer federal restrictions on using tracking technologies on their public facing webpages.

A consolidated class action against Meta over hospital pixel tracking remains active in a California federal court. In April 2025, a judge ordered Mark Zuckerberg to sit for a deposition, finding he was the final decision maker on all consequential privacy decisions at the company. Meta appealed. The court also considered sanctions against Meta for deleting data relevant to the case.

Your Body on a Billboard: Wearables and the Data They Collect

Look at your wrist. If you are wearing a smartwatch or fitness tracker, the device knows your heart rate right now. The device knows your heart rate variability, your blood oxygen saturation, your sleep stages from last night, your stress level, your skin temperature, your respiratory rate, and every place you went today down to the GPS coordinate.

Modern wearables sample your heart rate data between ten and one hundred times per second. A Samsung Galaxy Watch study showed that photoplethysmography signals were sampled every 100 milliseconds and transmitted to a server every thirty minutes. One 2025 analysis found that the Fitbit companion app collects up to 21 different categories of data, nearly double the industry average.

Google acquired Fitbit for $2.1 billion in 2021. At the time, Fitbit had roughly 28 million active users. A Google executive promised the deal was about devices, not data. The European Commission imposed conditions: Google pledged not to use Fitbit health, fitness, or location data for Google ads in Europe, committed to maintaining a separate data silo, and agreed to preserve third party access for ten years. The United States imposed no comparable conditions. Google later paid nearly $400 million in a settlement after investigations showed the company continued tracking Fitbit user location data after users had turned off location tracking.

A 2025 analysis from University College Dublin ranked the privacy policies of seventeen wearable manufacturers. Apple scored among the lowest risk, with the shortest policy at 4,408 words, an emphasis on processing data on the device itself, and a stated commitment not to use health data for advertising. On the other end, Xiaomi received the highest risk designation with sixteen high risk ratings across twenty four criteria. Whoop, a brand popular among fitness enthusiasts, had the longest privacy policy at 12,125 words and clustered with the highest risk group.

The FBI Knows Your Heart Rate, Too

Wearable data is not just a privacy concern. Wearable data is a forensic tool.

In Connecticut, a man named Richard Dabate told police that a masked intruder shot and killed his wife Connie in their home in December 2015. Connie's Fitbit told a different story. The device showed she was moving around the house for approximately one hour after the time Richard claimed she had been killed. Digital evidence from a laptop and Facebook Messenger timestamps confirmed the timeline. Richard Dabate was convicted of murder in May 2022 and sentenced to 65 years in prison. The Connecticut Supreme Court upheld the conviction in 2025, specifically ruling that Fitbit data was scientifically reliable and properly admitted as evidence.

In San Jose, California, a woman named Karen Navarra was found dead in 2018. Her Fitbit Alta HR showed a significant spike in heart rate at 3:20 PM on September 8, followed by a rapid decline. By 3:28 PM, the device registered no heartbeat at all. Ring camera footage placed her stepfather's car in the driveway during that eight minute window. He was arrested for murder.

Police obtain wearable data through warrants or subpoenas served directly on device manufacturers. Your device does not need to be physically in law enforcement's hands. They go straight to the company.

The FDA Clears the Device, Not Your Privacy

The Apple Watch received FDA clearance for ECG and atrial fibrillation detection in 2018. Fitbit received ECG clearance in 2020. Samsung followed the same year. In 2024, the Apple Watch received clearance for sleep apnea detection. These devices are performing medical grade functions. They are generating the kind of data that a hospital would guard under HIPAA.

FDA clearance does not trigger HIPAA protection. The same atrial fibrillation reading that your cardiologist would treat as protected health information has no federal privacy protection when your Apple Watch records it and sends it to a server. The device crosses into medical territory. The privacy law does not follow.

When Your Mental Health Becomes a Product

Mozilla Foundation ran the most thorough audit of mental health app privacy practices. In 2022, their Privacy Not Included team reviewed 32 mental health and prayer apps. Twenty eight of 32 received a warning label. Twenty five failed minimum security standards. Only two apps earned a clean bill of health: PTSD Coach, developed by the Department of Veterans Affairs, and Wysa, an AI chatbot. Mozilla's lead researcher described the category as exceptionally creepy, saying these apps track, share, and capitalize on users' most intimate personal thoughts and feelings.

A year later, the picture had gotten worse. Nineteen of those 32 apps still carried the warning label, and 40 percent of them had degraded in their privacy practices since the previous year. Talkspace, one of the largest online therapy providers, buried a clause in its privacy policy allowing the company to use inferences about gender identity, sexual orientation, and depression for marketing purposes.

One of the most disturbing cases involved Crisis Text Line, a nonprofit providing free 24/7 text based crisis support for people in moments of acute mental health distress. The organization had spun off a for profit company called Loris.ai, retaining a 53 percent stake, and shared crisis conversation data with the spinoff. Loris.ai used that data to build customer service chatbot software. Users reaching out during their darkest moments had to agree to more than 4,000 words of terms of service before they could get help. A former board chair acknowledged knowing full well that no one would read those terms. The data sharing arrangement ended in January 2022 after journalists exposed the practice.

$275 for Five Thousand People's Mental Health Records

A Duke University researcher contacted 37 data brokers in 2023, asking to buy health data in bulk. Twenty six responded. Eleven were ready and willing to sell mental health records. One broker advertised the names and addresses of individuals with depression, bipolar disorder, anxiety, panic disorder, cancer, PTSD, and personality disorders, sorted by race and ethnicity. The price tag for 5,000 people's mental health profiles was $275. That is about five cents per person.

The pipeline works like this. Health apps and websites embed tracking tools, software development kits and tracking pixels, automatically transmitting user data to advertising platforms and data brokers. When you open a health app, data flows. When you visit a health website, data flows. When you complete a questionnaire about your symptoms, data flows. An advertising identifier links all of this activity to your profile. A 2018 study found more than 60 percent of Android apps tested shared data with Facebook the moment a user

opened the app, regardless of whether the user even had a Facebook account. An investigation of 50 telehealth websites found 49 of 50 shared user data with advertising platforms. Thirteen of those sites sent answers to health intake questionnaires to companies including Meta, Google, TikTok, and Snapchat.

Blue Shield of California disclosed in 2025 a Google Analytics misconfiguration had shared patient data for 4.7 million individuals with Google Ads. This was not a hack. A configuration error at one of the largest insurers in the state exposed millions.

Your Employer Wants Your Steps, Too

Employer wellness programs represent another gap. When your company runs a standalone wellness program outside its group health plan, the data collected does not fall under HIPAA. These programs gather biometric screening results, blood pressure, cholesterol, BMI, health risk assessments, smoking status, and wearable activity data. The American Medical Association has warned that there are no regulations stopping companies from using data like calorie intake, blood pressure, and weight to penalize patients.

John Hancock, the insurance company, announced in 2018 that it would only sell life insurance policies that track fitness and health data. No more traditional policies without monitoring. The Vitality program, now celebrating its tenth anniversary, collects steps, exercise intensity, annual health screening results, flu shots, and even healthy food purchases. Compatible devices include Apple Watch, Fitbit, Oura Ring, Garmin, and Whoop. Policyholders earn premium savings of up to 25 percent. UnitedHealthcare's UHC Rewards program, available to three million members, offers up to $1,000 per year for daily steps, sleep tracking, biometric screenings, and health surveys.

John Hancock's own FAQ states the company will not use Vitality data to change a policyholder's risk classification. Consumer advocates are watching carefully. The infrastructure for insurance surveillance is being built one step at a time. Right now, these programs offer rewards. The data keeps accumulating. Patterns of declining activity, irregular heart rhythms, poor sleep, and elevated stress sit on company servers alongside your name and policy number.

Your Cycle, Your Messages, Your Criminal Case

After the Supreme Court's Dobbs decision in June 2022 overturning Roe v. Wade, privacy experts raised immediate alarms about period tracking apps. Missed periods, pregnancy test results, and fertility data logged in these apps could be subpoenaed by law enforcement in states that criminalized abortion.

The warnings proved prescient. In Nebraska, a 17 year old named Celeste Burgess and her mother Jessica were investigated by local police. Officers served Facebook with a warrant, and Meta complied, turning over private messages in which Jessica coached Celeste on taking abortion pills. Both mother and daughter were convicted. This case demonstrated in stark terms that digital communications about reproductive health, completely outside HIPAA, can become criminal evidence.

As of mid 2024, no court has subpoenaed period tracker app data specifically. Experts point out that prosecution timelines move slowly. A 2023 academic analysis of 35 period tracking apps found that 16 of them had privacy policies explicitly stating they could disclose personal data to law enforcement in response to subpoenas.

Some apps responded to the post Dobbs reality. Flo introduced an Anonymous Mode in late 2022, built with an encryption protocol stripping names, email addresses, and technical identifiers from health data. The feature won an industry privacy award. The mode remains opt in, meaning you have to know about the feature and turn it on yourself. Clue, a Berlin based app governed by the European Union's privacy law, publicly stated the company would refuse to share data with anyone, even in response to a legal subpoena. Post Dobbs, Clue saw a 2,200 percent increase in downloads over a single weekend. Apple added end to end encryption for health data synced through its cycle tracking feature.

States Are Moving. Washington Is Not.

The most aggressive response to the HIPAA gap has come from individual states. Washington State passed the My Health My Data Act in April 2023, and the state attorney general's office called it the first privacy focused law in the country designed to protect health data that falls outside HIPAA.

The law defines consumer health data broadly. It covers past, present, and future physical or mental health status, including medication purchases, biometric data, reproductive and sexual health information, and, critically, health information that algorithms derive or extrapolate from data that is not itself health data. The law requires opt in consent before collecting health data, demands separate consent before sharing it, requires written authorization before selling it, and gives consumers a broad right to delete. It includes a private right of action, meaning individuals can sue directly. It bans geofencing around healthcare facilities entirely, with no exceptions, to prevent anyone from tracking who visits a doctor's office or clinic.

Nevada passed a similar law in 2023 with attorney general enforcement only and no private right of action. Connecticut added health data provisions to its existing privacy act the same year. Maryland went further in 2024, passing the first state law to ban the sale of sensitive health data entirely, regardless of whether the consumer consents. Virginia passed a law in 2025 specifically restricting the disclosure and sale of reproductive and sexual health information, with statutory damages of at least $500 per violation. California's attorney general reached a $1.55 million settlement with a health media company in July 2025, treating the sharing of health article titles with advertisers as an impermissible disclosure of sensitive health data.

More than twenty states now have some form of privacy law treating health data as a sensitive category requiring extra protections. An eight state enforcement consortium formed in 2025 to coordinate privacy actions across borders.

Federal legislation keeps stalling. The American Data Privacy and Protection Act passed a House committee 53 to 2 in 2022 and never received a full vote. The American Privacy Rights Act did not advance beyond subcommittee in 2024. In November 2025, Senator Bill Cassidy, chair of the Senate HELP Committee, introduced the Health Information Privacy Reform Act. The bill would create a new category called Applicable Health Information covering data from apps and wearables, require consent before selling health data, mandate a right to deletion, and require companies to tell consumers when HIPAA does not protect them. As of early 2026, the bill remains in committee.

The result is a country where your health data rights depend on your zip code. A resident of Washington State has legal protections that a resident of Alabama does not. A Californian has tools that someone in Ohio lacks entirely.

192.7 Million Americans Exposed in a Single Breach

The largest health data breach in American history struck on February 21, 2024, when hackers broke into Change Healthcare, a clearinghouse that processes fifteen billion healthcare transactions every year. The attackers used stolen credentials on a server that did not have multifactor authentication enabled. The breach affected 192.7 million people, nearly two thirds of the entire United States population. Healthcare operations across the country ground to a halt for weeks. The total cost exceeded $2.4 billion.

The breach involved a HIPAA covered entity. The event received federal scrutiny. Now imagine the breach affecting a health app used by ten million Americans. No HIPAA obligation to notify patients. No HHS investigation. No federal consequence unless the FTC decides to act, and the current FTC has signaled a more narrow enforcement approach under new leadership installed in 2025. The agency dismissed two Democratic commissioners in March 2025. Privacy enforcement observers expect the commission to pull back from the aggressive health data enforcement actions producing the BetterHelp, GoodRx, and Cerebral settlements.

Meanwhile, breaches keep accelerating. In 2024, 742 large health data breaches exposed a record 288.9 million individual records. IBM's 2025 report found that healthcare breaches averaged $7.42 million each, making healthcare the most expensive sector for breaches for the fifteenth consecutive year. In 2025, Yale New Haven Health System disclosed a breach affecting 5.5 million people and settled for $18 million in just seven months. Blue Shield of California's Google Analytics error affected 4.7 million. Aflac reported 13.9 million records compromised. The DaVita kidney care company was hit with ransomware. Frederick Health lost a million records.

What You Need to Do Right Now

The gap between what you believe about your health data and what the law actually says is one of the most dangerous privacy illusions in America. Here is what you need to do to protect yourself.

Audit every health app on your phone. Open your settings and look at every app related to health, fitness, menstrual tracking, medication, sleep, or mental wellness. For each one, go into the app's settings and look for privacy controls, data sharing options, and account deletion features. Turn off everything that is not essential to the app's core function. If the app does not give you clear controls, delete it and find an alternative that does.

Check your wearable's privacy settings. If you wear a Fitbit, Apple Watch, Garmin, Oura Ring, Whoop, or any other device, open the companion app and review what data you are sharing and with whom. Turn off features like location tracking during workouts if you do not need them. If your device offers on device processing instead of cloud syncing, choose local storage whenever possible.

Read the first three paragraphs of every health app's privacy policy. You do not need to read all 12,000 words. The first few paragraphs typically tell you whether the company shares data with third parties, whether it sells your information, and what rights you have to delete your data. If those first paragraphs do not answer those questions clearly, that tells you something important. You can also copy and paste the privacy policy into your favorite AI chatbot and ask it to tell you all about whether or not the company shares data with third parties, does it sell your information and what your rights are to block this from happening and delete the data.

Use anonymous or privacy modes when available. Flo's Anonymous Mode, Apple's end to end encryption for health data, and similar features exist because these companies know the risk. Turn them on. Do not assume the default settings protect you.

Know your state's law. If you live in Washington, Nevada, Connecticut, Maryland, Virginia, or California, you have specific health data rights. Use them. Request deletion of your data from companies that no longer serve you. File complaints with your state attorney general if a company ignores your request.

Separate your health identity from your advertising identity. Use a dedicated email address for health apps that you do not use anywhere else. This makes it harder for data brokers to connect your health profile to your broader digital identity.

Talk to your family, especially your kids. Young people share health data freely through fitness challenges, mood tracking apps, and AI chatbots. They deserve to know that the app promising to help them manage anxiety might be selling their emotional state to an advertising platform.

The Illusion Ends When You See the Truth

The law that most Americans believe protects their health data was written three decades ago for a world of paper charts and fax machines. It was never designed to cover the smartwatch on your wrist, the therapy app on your phone, or the genetic data you spit into a tube and mailed to a company that later went bankrupt.

Eight out of ten Americans believe their health app data is protected. It is not. Data brokers sell mental health profiles for five cents a person. Hospitals sent your appointment details to Facebook through invisible tracking pixels. Insurance companies now require fitness data for their policies. AI therapy bots collect your deepest fears with minimal oversight. And 15 million Americans watched their DNA go through a bankruptcy auction.

This is not a story about technology. This is a story about you. Your blood pressure, your menstrual cycle, your therapy sessions, your heart rate at three in the morning, your prescription for antidepressants, your genetic predisposition to breast cancer. Every piece of this information has value to someone. And right now, the law treats most of it as fair game.

You deserve better than a privacy illusion. You deserve actual protection. Until Congress acts, you are the last line of defense for your own health data. Start today.

Chapter 13: Your Reproductive Data Isn't Private

They Sold Your Secret for a Hundred and Sixty Dollars

In 2022, a journalist at Vice walked up to a data broker's website, typed in a credit card number, and paid $160. For that price, the journalist received one week of location data covering more than 600 Planned Parenthood clinics across the United States. The data showed where each visitor came from, how long each visitor stayed, and where each visitor went afterward. The data broker, SafeGraph, had classified Planned Parenthood as a trackable brand and Family Planning Centers as a searchable category. Anyone with a credit card and an internet connection was able to do the same thing. No warrant. No subpoena. No judge. Just a credit card and a few clicks.

Stop and sit with that for a moment. Your visit to a reproductive health clinic, a visit you believed was private, a visit protected by the walls of a medical facility, was for sale. Your arrival time, your departure time, your home address, all of it bundled into a data set and offered to the highest bidder. And here is the part that should send a chill down your spine. This was perfectly legal under current laws.

This chapter is about the most intimate data you generate as a human being. Your reproductive health information. Your menstrual cycle data. Your pregnancy status. Your fertility treatments. Your clinic visits. Your search history when you type a question about your own body into Google at two in the morning. This is data so personal that most people assume federal law protects it. Most people are wrong.

Since the Supreme Court's 2022 Dobbs decision overturned the constitutional right to abortion, reproductive health data has become the most legally vulnerable category of personal information in America. And the single federal rule designed to protect it lasted exactly eleven months before a Texas judge struck it down and the federal government walked away. What follows is the story of how we got here, who profits from the surveillance of reproductive decisions, and what you need to do to protect yourself and the people you love.

The Federal Protection That Vanished

In April 2024, the Biden administration published the HIPAA Reproductive Health Rule. The rule was the most ambitious federal attempt to shield reproductive health records from law enforcement. It prohibited hospitals, clinics, insurance companies, and other HIPAA covered entities from handing over protected health information in response to criminal, civil, or administrative investigations targeting someone for seeking, obtaining, providing, or helping with reproductive health care that was lawful in the state where it occurred. The rule covered abortion, contraception, IVF, fertility treatment, prenatal care, and gender affirming care. It required anyone requesting reproductive health records, whether law enforcement, a court, or an oversight agency, to sign a sworn statement declaring the request was not for a prohibited purpose.

The rule survived barely eleven months. On June 18, 2025, Judge Matthew Kacsmaryk of the Northern District of Texas vacated the rule nationwide. Kacsmaryk, the same judge who previously attempted to overturn FDA approval of mifepristone, ruled that HHS had exceeded its authority. The plaintiff was a Texas family medicine physician represented by the Alliance Defending Freedom. At least 17 states had filed challenges across four separate lawsuits. The Trump administration's HHS declined to defend the rule. When the appeal deadline passed on August 18, 2025, no government lawyer filed an appeal. Proposed intervenors,

including two cities and a physicians' organization, attempted their own appeal and then withdrew it. The Fifth Circuit dismissed the case on September 10, 2025. The rule is permanently dead.

What remains is the original HIPAA framework, which permits covered entities to disclose your health information to law enforcement. It does not require disclosure. Your doctor still has the discretion to say no. HIPAA's minimum necessary standard still applies, meaning providers should share only the minimum information needed to respond to a request. The American Medical Association, the American College of Obstetricians and Gynecologists, and the American Psychological Association all issued statements in late 2025 calling for protections of reproductive health records. None of those statements carry the force of law. The specific prohibition against disclosing reproductive health records for prosecution purposes no longer exists anywhere in federal law.

When Your Phone Testifies Against You

The cases are real. The names are public. The consequences were devastating.

In Mississippi in 2017, Latice Fisher, a Black mother of three, experienced what she described as an unexpected delivery that ended in a stillbirth. She went to the hospital seeking medical care. Police showed up. They asked for her iPhone. She handed it over voluntarily, the way most people would, believing she had nothing to hide. Investigators did not focus on medical evidence. They scrolled through her search history and found queries like "how to induce a miscarriage" and "buy Misoprostol Abortion Pill Online." No physical evidence confirmed she ever took any medication. No toxicology report supported the prosecution's theory.

The search history, the private questions she had typed into her phone in a moment of desperation, became the prosecution's case. A grand jury indicted her on second degree murder charges carrying up to 40 years in prison. Fisher spent weeks in jail, separated from her children, before a second grand jury declined to indict in March 2020. Her life was upended because she searched for information on her own phone. Civil rights attorney Cynthia Conti-Cook described the case as a blueprint for how digital evidence gives prosecutors a window into a woman's most private thoughts.

In Indiana in 2013, Purvi Patel was convicted of feticide based partly on text messages about ordering abortion pills from an overseas pharmacy and email exchanges with the supplier. A toxicologist found no trace of the drugs in Patel's body or in the fetus. She was sentenced to 20 years in prison before an appeals court vacated the feticide conviction in 2016, marking the first time a state feticide law had been used against a woman for attempting her own abortion. In Idaho, police used cellphone geolocation records to track a visit to a Planned Parenthood in Oregon, then charged a mother and son with kidnapping for helping someone travel across state lines for an abortion.

According to Pregnancy Justice, prosecutors brought 210 pregnancy related criminal cases in the single year following the Dobbs decision, the highest number ever recorded in a single year. Charging documents increasingly cite "researching or exploring the possibility of an abortion" as evidence of criminal intent. The digital tools prosecutors rely on include search warrants for phone contents and cloud backups, subpoenas for subscriber records, geofence warrants demanding data on every device within a geographic area, and direct purchases of location data from commercial data brokers. That last category requires no warrant at all.

The Hundred and Sixty Dollar Surveillance Machine

The data broker pipeline targeting reproductive health clinics works through a chain most Americans never see. Hundreds of smartphone apps, from weather apps to games to prayer apps, contain embedded code called Software Development Kits from data brokers. These SDKs silently transmit your phone's GPS coordinates, accurate to within ten feet, along with a unique Mobile Advertising ID that functions as a permanent digital fingerprint. Data brokers collect billions of these location pings every day and sell them to anyone willing to pay.

In October 2024, privacy researchers obtained a trial of Locate X, a surveillance tool sold to law enforcement by a company called Babel Street, simply by saying they planned to work with police. Using the tool, the researchers tracked a specific device from a home in Alabama to a reproductive health clinic in Florida and back again. They observed 700 unique phones at the clinic. Each device was traceable to its apparent home address. No warrant was required to access any of this information.

A company called Near Intelligence sold location data to an anti-abortion organization called the Veritas Society, which hired an advertising agency to draw digital geofences around reproductive health clinics and deliver more than 14 million targeted ads to people who visited those clinics across 48 states. Near's own chief privacy officer admitted the company had no technical controls to prevent this kind of targeting. Another data broker, Mobilewalla, specifically collected location data from women at pregnancy centers to build audience segments labeled "pregnant women" for advertisers, pulling from real time bidding auctions that broadcast personal data hundreds of billions of times each day.

The FTC sued data broker Kochava in August 2022, its first lawsuit against a geolocation data broker, after Kochava's data, covering 94 billion transactions per month across 125 million devices, tracked visits to reproductive health clinics. A settlement reached in late 2025 required Kochava to filter 2.1 million sensitive locations from its data sets.

Google announced in July 2022 that it would automatically delete location history entries near abortion clinics. A Washington Post investigation found the system failed roughly half the time, with the clinic name deleted in some cases and the route traveled still visible. Google made a more structural change in late 2023 by moving all location history to on device storage with a three month retention period and encryption for cloud backups, effectively ending its ability to respond to geofence warrants.

The Apps That Know Your Cycle Better Than You Do

More than 100 million women worldwide use period tracking and fertility apps to log their menstrual cycles, pregnancy status, sexual activity, fertility indicators, symptoms, and medications. This is some of the most intimate data a person generates. And almost none of it falls under HIPAA, because these apps are consumer products, not medical providers. They occupy a regulatory gap that leaves users exposed.

The FTC's 2021 enforcement action against Flo Health proved just how exposed users were. Flo, one of the most popular period tracking apps in the world, had shared menstruation data, fertility data, and pregnancy data with Facebook, Google, AppsFlyer, and Flurry. Flo had promised users their information would stay private. The FTC's consent order, finalized by a unanimous 5-0 commission vote, required Flo to obtain affirmative user consent before sharing health data, notify all affected users, and instruct third parties to

destroy the improperly obtained data. It was the first time the FTC had ever required a company to notify its own users of a privacy enforcement action.

To its credit, Flo responded by launching an Anonymous Mode feature in mid 2022. This feature uses a technology called Oblivious HTTP through Cloudflare so that no single party holds both a user's identity and their health data simultaneously. If law enforcement requests data on an Anonymous Mode user, Flo says it cannot fulfill the request because it does not have the information needed to connect anonymous data to a specific person.

Natural Cycles, a Swedish based FDA cleared contraceptive app, developed a similar Go Anonymous feature and stores user data within the European Economic Area under strict European privacy regulations. Clue, headquartered in Berlin, went further. Its co-CEOs publicly pledged they would never hand over private health data to any authority seeking to use it against their users, and stated that as a German company, no German court would permit such a request.

Not every app has reformed. Ovia, owned by diagnostics company Labcorp, sells data to employers' human resources departments. Its terms of use grant the company a royalty free, perpetual, and irrevocable license, throughout the universe, to use and exploit de-identified user data. Ovia retains your data for seven years after you stop using the app. FEMM, a fertility tracking app with more than 400,000 downloads, is funded by the Chiaroscuro Foundation, a Catholic charity backed by conservative hedge fund manager and anti-abortion activist Sean Fieler, who sits on FEMM's board. Two of FEMM's medical advisors have ties to a Catholic university in Chile where abortion access is severely restricted. The app collects menstrual and sexual activity data from users who have no idea about the ideological commitments of its funders.

The safest choice among major apps is Euki, the only app to earn Mozilla's Best Of rating. Euki stores all data locally on your device, collects no personal information, and includes a feature called a duress PIN. If someone forces you to open the app, entering 0000 displays false data. In Mozilla's 2022 investigation of 25 reproductive health apps, 18 received the Privacy Not Included warning label. Only Euki and Natural Cycles avoided the warning.

No publicly documented case has so far involved a subpoena specifically for period tracking app data. The legal tools to compel that data already exist. Apps storing data in the cloud are subject to warrants and subpoenas under the Stored Communications Act. Data stored on a seized phone is extractable by forensic tools like Cellebrite and GrayKey, which recover deleted app data from internal databases until the data is overwritten. The potentially incriminating data points in a fertility app are significant. A logged pregnancy that suddenly disappears. A missed period followed by a resumption of regular cycles. Symptoms consistent with medication abortion. Location data showing a clinic visit. Each of these data points tells a story, and prosecutors know how to read it.

A Country Split in Two

With federal protection gone, state legislatures have become the front line of reproductive data privacy. The result is a patchwork so extreme that the same data point receives strong protection in one state and zero protection in the state next door.

Washington's My Health My Data Act, signed in April 2023, is widely considered the strongest law in the nation. It created the broadest definition of consumer health data anywhere in America, covering reproductive and sexual health information, efforts to research or obtain reproductive services, and data derived or inferred from non-health information. That last category is critical, because it means an algorithm that infers your pregnancy status from your shopping patterns falls under the law's protection. The law applies to any business regardless of size, covers data collected in Washington regardless of where the consumer lives, prohibits geofencing within 2,000 feet of health care facilities, and provides a private right of action with damages up to $25,000.

California has assembled the most layered package of protections through multiple statutes. AB 254, signed in 2023, reclassified period tracking apps and fertility websites as providers of health care, subjecting them to the same privacy standards as a doctor's office. AB 352, also from 2023, requires electronic health record systems to include features restricting access to reproductive health data.

AB 1242, signed in 2022, prohibits California based communication companies from complying with out of state warrants for abortion services that are legal in California. And AB 45, signed by Governor Newsom on September 26, 2025 and effective January 1, 2026, bans geofencing within 1,850 feet of any family planning center and prohibits collecting, selling, or sharing personal information of anyone at or near those facilities. Violations carry penalties of $25,000 per violation, and the law gives individuals the right to sue for up to three times their actual damages.

Virginia provided a rare bipartisan moment when Republican Governor Glenn Youngkin signed SB 754 on March 24, 2025, effective July 1, 2025. The law prohibits obtaining, disclosing, selling, or spreading personally identifiable reproductive or sexual health information without consent. Its definition of protected data covers menstruation, pregnancy, contraception, fertility, sexual activity, and information derived or inferred from non-health data. It provides individuals the right to sue with minimum statutory damages of $500 per violation.

Six states now ban geofencing near health care facilities. Washington established a 2,000 foot radius effective July 2023. New York set an 1,850 foot limit effective July 2023. Connecticut imposed a 1,750 foot limit effective October 2023. Nevada followed with 1,750 feet effective March 2024. California added its 1,850 foot ban effective January 2026.

At the opposite end of the spectrum, states with total abortion bans and no reproductive data privacy protections include Alabama, Mississippi, Louisiana, Missouri, South Dakota, North Dakota, West Virginia, and Oklahoma. Idaho has a citizen enforcement mechanism allowing private civil lawsuits with $10,000 bounty rewards.

Texas, despite having a broad data privacy law, has the most aggressive enforcement posture. The Texas attorney general has sued a New York doctor for providing abortion pills to a Texas resident, testing whether shield laws in one state hold up against legal attacks from another. The data broker loophole means that even in states with strong shield laws, law enforcement in states with abortion bans is free to purchase location data showing your clinic visits without any warrant or judicial oversight. Montana became the first state to close this loophole in 2025, prohibiting law enforcement from purchasing data from brokers when they would otherwise need a warrant to obtain the same information.

The Federal Enforcer That Stepped Back

During the Biden administration, the Federal Trade Commission became the most active federal enforcer of reproductive data privacy, bringing a wave of cases that reshaped the rules for health apps and data brokers. Beyond the Flo settlement, the FTC brought the first ever enforcement of the Health Breach Notification Rule against GoodRx in February 2023, securing a $1.5 million penalty for sharing prescription and health data with Facebook, Google, and advertising platforms at the same time it displayed a fake HIPAA Secure badge on its website. A month later, the FTC reached a $7.8 million settlement with BetterHelp for sharing mental health intake data, including responses about depression and suicidal thoughts, with Facebook, Snapchat, and Pinterest. In May 2023, the agency penalized Easy Healthcare's Premom ovulation app for sharing fertility data and GPS coordinates with analytics firms based in China.

That enforcement energy has evaporated. Under the current FTC chair, the commission operates with only two sitting commissioners, and leading privacy attorneys predict a dramatic pullback from the health data enforcement theories developed over the past few years. The amended Health Breach Notification Rule, which explicitly covers health apps and connected devices, remains in effect. New enforcement actions pushing the boundaries of health data privacy appear unlikely.

Congress has not stepped in. The My Body My Data Act, reintroduced in June 2025, would create a national standard protecting reproductive and sexual health data collected by apps, search engines, and other entities outside HIPAA's reach. It has 22 co-sponsors from one party and zero from the other. It has not received a committee hearing. Federal legislation on reproductive data privacy is dead on arrival in the current political environment.

What Is Coming Next

The trajectory through 2027 points toward deeper fragmentation. At the federal level, the HIPAA rule is dead, the FTC is pulling back, and no new federal rulemaking on reproductive data privacy is expected. The administration proposed defunding Title X in its FY2026 budget and has restructured agencies that previously supported reproductive health research.

At the state level, the patchwork is growing. New York's Health Information Privacy Act, vetoed by Governor Hochul in December 2025, was reintroduced in February 2026 with changes addressing the governor's objections. If enacted, it would become the strongest state health data privacy law in the nation, covering reproductive data, location data, payment information, and health inferences. Virginia is advancing another bill to prohibit sales of precise geolocation data. Hawaii, Vermont, and New Hampshire are moving reproductive data bills forward. Legal analysts expect the pace of state health data legislation to accelerate through 2026 and into 2027.

The most significant legal development on the horizon sits at the Supreme Court. The Court has agreed to hear Chatrie v. United States, a case asking whether geofence warrants, which demand data on every device within a geographic area, violate the Fourth Amendment. The Fourth and Fifth Circuit courts of appeal split on this question.

A decision striking down geofence warrants would meaningfully protect people visiting reproductive health clinics. In Maryland, the state supreme court's February 2025 decision in Moira Akers v. State overturned a

conviction where internet searches about abortion had been used as evidence of criminal intent, setting an important precedent against weaponizing reproductive health search data.

The emerging frontier of risk is artificial intelligence powered behavioral inference. Modern algorithms already infer pregnancy from purchasing patterns, browsing behavior, and location data without any explicit health information being shared. You do not need to tell anyone you are pregnant. The algorithm figures it out from the supplements you buy, the websites you visit, the stores you walk into, and the changes in your daily routine. Researchers at the Harvard Petrie-Flom Center warned in November 2024 that AI's ability to predict miscarriage or stillbirth risks becoming a surveillance tool, casting suspicion on women who suffer natural pregnancy losses.

A woman who has a miscarriage could find herself flagged by an algorithm that determined she was pregnant and then determined she was no longer pregnant, all without her ever telling a single person about her pregnancy. Wearable devices passively collecting skin temperature, heart rate variability, and sleep patterns reveal pregnancy status without any user input. Your smartwatch knows your body is changing before you confirm it yourself. The reproductive technology industry, projected to exceed $50 billion in revenue, generates enormous quantities of data that falls entirely outside HIPAA's reach.

Protecting Yourself Starting Today

The law is fractured. The federal government has stepped back. Your state of residence determines whether your reproductive health data receives strong protection or no protection at all. And the data broker economy continues to collect, package, and sell your most intimate information every single day. Waiting for the legal system to fix this problem means waiting too long.

You need to take control of what is within your power right now. Use Signal with disappearing messages for any sensitive health conversations. Disable the Mobile Advertising ID on your phone, which is the tracking identifier that follows you from app to app. On an iPhone, go to Settings, Privacy, Tracking, and turn off Allow Apps to Request to Track. On Android, go to Settings, Privacy, Ads, and delete or reset your advertising ID. Leave your phone at home or power it off when visiting a reproductive health clinic.

Use a privacy focused browser like DuckDuckGo with NordVPn or the Tor Browser for any health research. Choose period tracking apps that store your data locally on your device, not in the cloud. Euki, Natural Cycles in anonymous mode, and Clue all offer stronger privacy protections than most alternatives. Use cash for reproductive health purchases. Keep your phone's operating system updated, because newer software is significantly harder for forensic tools to crack. And talk to the people you love about these risks, especially your daughters, your sisters, your nieces, and your friends.

HIPAA covers roughly 10 percent of the entities that handle reproductive health data. Think about that number. Ten percent. Your doctor's office, your hospital, your insurance company, these are covered. The other 90 percent, the period tracking apps, the search engines, the data brokers, the advertisers, the social media platforms, the wearable device manufacturers, they operate with little or no federal regulation. The fertility app on your phone knows more about your reproductive health than your doctor's billing system, and it has far fewer legal obligations to keep that information private. More than 400 people faced pregnancy related criminal charges in the two years following the Dobbs decision. Digital evidence played a prominent role in case after case. The data trail you leave behind every time you open an app, search for a symptom, or

walk into a clinic has become a weapon in the hands of prosecutors in states that have criminalized reproductive health decisions.

This is not a problem that will fix itself. This is not a problem that someone else will solve for you. This is a problem that requires every American to understand the risks, protect their own data, and demand that their elected representatives close the gaps that leave reproductive health information exposed. Call your state legislators. Ask them where they stand on reproductive data privacy. Ask them whether law enforcement in your state is allowed to buy your location data without a warrant. Ask them whether they support geofencing bans near health care facilities.

Make this a voting issue, because the people making these decisions are counting on you not paying attention. Your body is your own. Your health decisions are your own. And the data those decisions generate should be your own too.

Chapter 14: They Sold Your DNA Without Asking You

In October 2023, a hacker using the alias "Golem" posted one million stolen genetic profiles on a dark web forum. Every single profile belonged to a person of Ashkenazi Jewish descent. The price tag was one dollar per person. A few days later, the same hacker uploaded hundreds of thousands of profiles belonging to people of Chinese ancestry, offered at similar prices.

The stolen data came from 23andMe, the company where 15 million Americans had mailed their saliva, answered deeply personal health questions, and trusted a corporation to guard the most permanent information a human body produces. Within eighteen months, the company filed for bankruptcy. And the DNA of every one of those 15 million people became an asset in a corporate fire sale.

You need to sit with this for a moment. One dollar. Your genetic code, your ancestry, your health risks, your family connections going back generations, all of these things were worth less than a cup of gas station coffee to a criminal on the internet.

This story is about far more than a single company falling apart. The most personal information your body produces has no dedicated federal law protecting the people who hand their DNA over to private businesses. Companies holding your genetic material face almost no legal obligation to keep your data safe when they go broke. And your genetic data does not belong solely to you. Your spit in a tube exposed your parents. Your children. Your cousins and half siblings and distant relatives you have never met. None of them consented to a thing.

If you have ever taken a direct to consumer DNA test, or if anyone in your biological family has, this chapter will show you exactly what happened to all of the data, who wants access to your genetic information, and what you need to do about the situation right now.

What Your Spit Tube Actually Surrenders

When you mail a saliva sample to a company like 23andMe, AncestryDNA, or MyHeritage, you are sending them approximately three billion base pairs of genetic code. This is the complete biological blueprint of who you are, encoded in every cell of your body. Most of these companies run your saliva through something called an SNP genotyping array, which reads between 600,000 and 700,000 specific genetic markers. Those markers represent a tiny fraction of your full genome, roughly 0.03 percent. Through a computational process called imputation, the companies fill in the gaps and infer millions of additional data points about you, essentially building a far more detailed picture of your biology than the raw analysis alone would produce.

The genetic readout is only the starting line. These companies also collect your answers to detailed health surveys, questions about your medical conditions, your lifestyle habits, your family medical history, and your physical traits. 23andMe built what the company described as "billions of phenotypic data points" from the roughly 80 percent of customers who opted into research. On top of all of this, the companies store your name, your email address, your date of birth, your ethnicity, your billing information, and your browsing behavior tracked through cookies and marketing tools. A 2021 Consumer Reports investigation found these companies "over collect personal information" and deploy marketing trackers with the potential to reveal sensitive health conditions.

Here is the part most people miss when they think about a DNA test. You are not handing over something you get to take back. Your genetic code stays the same from the moment you are born until long after you die. No reset button exists. No replacement copy arrives in the mail. Your DNA passes through your bloodline for generations. And researchers have shown as few as 30 genetic markers are enough to single out one individual from the rest of the human population.

Think about what this means for your family. A landmark 2018 study found a genetic database covering just 2 percent of a target population provides enough information to find a third cousin match to almost any individual in the group. So when one person in your family decides to take a DNA test, the decision effectively enrolls your entire extended family in a genetic database without their knowledge. Your grandmother, your nephew, your second cousins. Nobody asked them. Nobody told them. The European Union recognized this problem by defining a new category called a "biological group," one whose members come into existence every time genetic data gets processed and whose members never chose to join.

And deletion is never as clean as the companies suggest. When you close your 23andMe account, the company still keeps your genetic information, date of birth, and sex for what the company calls regulatory compliance. Data already folded into research studies is gone from your control forever. MIT Technology Review asked 23andMe to explain exactly what "retained genetic information" means after account deletion. The company pointed them to a privacy policy and refused to say anything more. You spit in a tube thinking you were learning about your ancestry. What you actually did was hand over data the company now refuses to fully describe or fully delete.

From Six Billion Dollars to Bankruptcy Court

The story of 23andMe's collapse reads like a warning label nobody bothered to print. Anne Wojcicki founded the company in 2006. The company went public in June 2021 through a special purpose acquisition at a valuation of roughly 3.5 billion dollars. The stock briefly pushed the market value to six billion dollars, peaking near $17.65 per share. The company never earned a profit. Not once. Not in a single quarter of a single year since founding.

The business model was fatally simple. A DNA test kit is a one time purchase. Once you buy one, you are done. There is no subscription. There is no refill. As the novelty wore off and the market filled up, revenue dropped from about $299 million in fiscal year 2023 to approximately $255 million in fiscal year 2024. Net losses hit $667 million in the same period. A five year exclusive partnership with GlaxoSmithKline, which had included a $300 million investment and produced around 50 drug discovery programs, ended. Cash reserves drained from $314 million in mid 2023 to $127 million by September 2024.

Then the breach ripped through the company. Between April and September 2023, hackers used a technique called credential stuffing, where criminals take stolen passwords from other websites and try them on new accounts, to break into about 14,000 individual 23andMe profiles. Because of the company's DNA Relatives feature, which links customers to their biological matches, the hackers scraped data on roughly 6.9 million users. Nearly half the entire customer base. The data appeared on dark web forums with disturbing ethnic specificity, including the "Golem" posting targeting Jewish and Chinese Americans. The company took five months to detect the intrusion. When 23andMe finally responded, the company blamed its own customers for reusing passwords. The response drew widespread condemnation from cybersecurity experts, consumer advocates, and elected officials.

The fallout was swift. In September 2024, Nasdaq warned the stock would be delisted because the share price had fallen below one dollar. On September 17, 2024, all seven independent board members resigned on the same day, including Sequoia Capital's Roelof Botha and YouTube CEO Neal Mohan. Their joint resignation letter cited Wojcicki's concentrated voting power of 49 percent and the absence of any real plan to rescue the company. A reverse stock split in October briefly propped up the share price. In November, 200 employees lost their jobs, slashing the workforce by 40 percent and leaving roughly 300 people at a company once employing thousands.

On March 23, 2025, 23andMe filed for Chapter 11 bankruptcy, listing $277 million in assets and $214.7 million in debts. Wojcicki stepped down as CEO and immediately positioned herself as a buyer. What followed became the most consequential data privacy fight in the history of American bankruptcy law.

Your DNA on the Auction Block

The bankruptcy sale process turned the genetic data of 15 million Americans into a contested asset on an auction floor. Regeneron Pharmaceuticals won the initial bid at $256 million for substantially all of 23andMe's assets, including the full genetic database and the biobank of physical saliva samples. Wojcicki's newly created nonprofit, TTAM Research Institute (the initials stand for Twenty Three And Me), challenged the result, secured a reopened auction through her attorneys at Quinn Emanuel, and ultimately prevailed on June 13, 2025 with a $305 million offer.

The court appointed Consumer Privacy Ombudsman, Washington University professor Neil Richards, filed a remarkable report two days before the auction closed. He wrote he was unable to determine whether the sale of customer data was consistent with the company's own privacy policies. He recommended the court require the buyer to get fresh consent from every customer before the transfer. The court declined to follow his recommendation.

On June 27, 2025, Judge Brian C. Walsh approved the sale. He used a legal structure called an "equity toggle" placing the company's assets into a subsidiary, then selling the subsidiary's equity to TTAM. The court ruled this structure did not qualify as a "transfer" of data, which would have triggered consent requirements under state genetic privacy laws. The judge acknowledged the sale was, in his own words, "a scary proposition." He noted lawmakers had simply never prohibited this type of arrangement. TTAM completed the acquisition on July 14, 2025. The company was renamed Chrome Holding Co. in bankruptcy proceedings, and most of the Chapter 11 cases were closed on January 21, 2026.

During the bankruptcy proceedings, 1.9 million customers, roughly 15 percent of the total, rushed to request deletion of their genetic data. The flood of requests crashed 23andMe's website for hours. A class action settlement over the original data breach was finalized at up to $50 million, with a claims deadline of February 17, 2026 and payouts still pending as of March 2026. The UK's Information Commissioner fined the company an additional 2.31 million pounds for the breach.

Washington Wakes Up, Slowly

The 23andMe crisis crashed into an already growing panic in Washington about foreign governments systematically acquiring American genetic data. The collision produced a wave of federal action, some of the measures meaningful, some of the bills still stuck in committee.

The most significant regulatory response was the DOJ Bulk Data Rule, implementing Executive Order 14117, signed by President Biden on February 28, 2024. The final rule took effect on April 8, 2025 and prohibits or restricts bulk transfers of Americans' sensitive personal data to six designated countries of concern: China (including Hong Kong and Macau), Russia, Iran, North Korea, Cuba, and Venezuela. For genetic data specifically, the rule sets the lowest trigger threshold of any data category. Any transfer involving 100 or more U.S. persons' genomic data or biological samples falls under the prohibition. For comparison, personal health data and financial data each require 10,000 persons to trigger the same rule. The extraordinarily low number tells you how seriously the Department of Justice treats genetic information.

The rule applies regardless of whether the data has been anonymized, encrypted, or scrubbed of identifying details. The DOJ concluded even data stripped of names and addresses, when combined in large quantities, gives hostile governments the ability to identify specific individuals and carry out targeted operations. Deputy Attorney General Todd Blanche put the logic plainly: "If you're a foreign adversary, why would you go through the trouble of complicated cyber intrusions and theft to get Americans' data when you just buy the data on the open market?" Enforcement began on July 8, 2025 after a 90 day grace period, with additional compliance requirements including mandatory audits and annual reporting kicking in on October 6, 2025.

These concerns are grounded in real cases, not speculation. China's BGI Group, the Beijing Genomics Institute, has been designated by the Department of Defense as a Chinese military company operating in the United States. BGI subsidiaries were sanctioned in 2023 for using genetic analysis to support China's repression of Uyghur populations. BGI also harvested genetic data from over eight million pregnant women in Europe through prenatal tests marketed to hospitals, without disclosing the data was being routed to China and used in research with the People's Liberation Army. Another Chinese company, WuXi AppTec, which generates over 60 percent of revenue from U.S. clients, has sponsored events with China's military and jointly operated genetic collection sites with PLA entities. The BIOSECURE Act, signed into law on December 18, 2025, now prohibits federal agencies from contracting with biotechnology companies identified as national security threats, with BGI entities already on the designated list.

At the June 2025 Senate Judiciary hearing on the 23andMe bankruptcy, Senator Grassley warned foreign governments holding large genetic databases gain the ability to track military and intelligence personnel, monitor ethnic minority populations, plant DNA evidence to frame individuals, and conduct intelligence operations falling into what one analyst described as "gray zone" geopolitical conflict. A 2025 Bulletin of the Atomic Scientists assessment found ethnicity targeted bioweapons remain "exceptionally challenging" because human populations are genetically diverse and blended through centuries of intermarriage. The most immediate and realistic dangers center on surveillance, identification, and intelligence exploitation of American genetic data by adversary nations.

Congress introduced three bills directly in response to the 23andMe crisis. The Genomic Data Protection Act, introduced in March 2025 by Senators Bill Cassidy and Gary Peters, would give consumers the right to access, delete, and request destruction of their genetic data and biological samples within 30 days. The Don't Sell My DNA Act, introduced in May 2025 by Senators Chuck Grassley, John Cornyn, and Amy Klobuchar, would rewrite the federal Bankruptcy Code to explicitly treat genetic information as personally identifiable information and require written consent before any sale in bankruptcy.

As Senator Grassley stated at the June 2025 hearing, "Genetic data is the blueprint to a person. It is sensitive, it is personal, and in the wrong hands, it is dangerous." The American Genetic Privacy Act of 2025,

introduced in March 2025 by Representative Tim Burchett, would prohibit the sale or disclosure of DNA testing data to China or Chinese controlled entities. As of March 2026, all three bills remain in committee. Not one has reached a floor vote.

Your Protection Depends on Your Zip Code

Because Congress has not passed a dedicated federal genetic privacy law, the states have become the front line. The result is a patchwork of protections varying wildly depending on where you live. More than 15 states have enacted dedicated genetic privacy laws as of early 2026, with several more introducing bills explicitly in response to the 23andMe collapse.

Montana has arguably built the strongest genetic privacy framework in the country. The state's Genetic Information Privacy Act, first enacted in 2023 and expanded in May 2025, requires a warrant before any government agency accesses genetic data. The law mandates genetic data of Montana residents be stored within the United States. Express consent is required for collection and all secondary uses. The 2025 expansion also extended protections to neurotechnology data, making Montana one of only three states protecting neural information.

Indiana's law, signed in May 2025, includes a provision going further than any other state. The statute absolutely prohibits consumer genetic testing companies from sharing individual genetic data with insurers, employers, or risk assessment businesses. No exceptions. No consent workaround. Civil penalties run up to $7,500 per violation.

Texas broke new ground in 2023 by establishing individual property rights in genetic samples and data, treating your DNA as property you own, not merely as regulated information. The 2025 Texas Genomic Act added national security protections, banned genetic data transfers to adversary nations during bankruptcy, prohibited the use of genome sequencing equipment produced by foreign adversaries, and created a private right of action allowing individuals to sue for up to $5,000 per violation.

Illinois has the oldest dedicated genetic privacy statute, dating to 1998, and carries the steepest penalties: $15,000 per intentional violation and $2,500 per negligent violation. A wave of more than 50 complaints in 2023 alone activated a law dormant for 25 years.

California layers multiple protections through its Genetic Information Privacy Act, the CCPA, and CalGINA, which extends anti discrimination protections to housing, education, and mortgage lending.

Maryland's 2025 Genetic Testing Protection Act fills a critical hole in federal law by prohibiting insurers from discriminating based on genetic test results in life insurance, long term care insurance, and disability insurance.

Several states are pushing even further.

South Carolina's proposed SB 731 would establish a private right of action with no exemption for so called de identified data, reflecting the growing recognition genetic data simply is not something you strip of identifiers and call safe. Alabama's proposed HB 265 would make the unauthorized sale of genetic data a Class C felony. The trend is unmistakable. The 23andMe crisis lit a fire under state legislatures across the country, and the bills keep coming.

The Federal Law Leaving You Exposed

The Genetic Information Nondiscrimination Act, known as GINA, was signed into law in 2008. Most Americans who have heard of the statute assume their genetic information is fully protected. GINA prevents health insurers from using genetic information in underwriting decisions. The law also stops employers with 15 or more workers from using genetic data in hiring, firing, or promotions. Those are real protections, and they matter.

Here is what GINA does not cover. Life insurance. Disability insurance. Long term care insurance. These are exactly the types of coverage where a person's genetic risk profile carries the most weight, and where the financial incentive to discriminate runs the deepest.

The threat is not hypothetical. A woman identified as Jennifer Marie in a report by Fast Company was denied life insurance at 36 years old because she tested positive for a BRCA1 gene mutation, which is linked to elevated breast and ovarian cancer risk. Her denial letter specifically cited the genetic test result. Life insurers in most states face no legal barrier to requesting genetic test results, asking whether you have been tested, and factoring your answers into their underwriting decisions. The more you learn about your own genetic risks, the more vulnerable you become to discrimination in the very areas of insurance where coverage matters most.

The chilling effect is real and well documented. The NIH's National Human Genome Research Institute has found the absence of protection against genetic discrimination in life insurance has caused many people to avoid getting tested altogether. Others who do get tested refuse to share results with their own doctors, knowing once genetic information lands in a medical record, insurers reviewing future applications gain access.

Think about what this means for public health. People are choosing ignorance about their own cancer risk, their heart disease risk, their neurological risk, because our laws punish knowledge. They are making themselves sicker to avoid being uninsurable. Only a handful of states, including Florida, Oregon, and now Maryland, have extended nondiscrimination protections to life and long term care insurance. For the vast majority of Americans, the gap in GINA remains wide open.

When Your Cousin's DNA Test Helps Catch a Killer, and When Things Go Wrong

Genetic data also intersects with law enforcement in ways most Americans have never considered. The most famous case is the arrest of Joseph James DeAngelo, the Golden State Killer, in April 2018. Investigators uploaded crime scene DNA to GEDmatch, a free genealogy database, and found several third and fourth cousin matches. Four months of family tree reconstruction narrowed the search. Discarded DNA collected near DeAngelo's home confirmed the match. He pleaded guilty to 13 murders and received a sentence of life without parole. The technique has since helped solve more than 651 criminal cases. The power of the method is undeniable. So are the risks.

GEDmatch was acquired by Verogen in 2019, later purchased by QIAGEN for $150 million, and shifted to an opt in model for law enforcement searches. QIAGEN later discovered some forensic genealogy practitioners had systematically bypassed the opt in settings, accessing profiles of users who had specifically declined to participate and falsifying reports to hide the violations. FamilyTreeDNA took a different path

entirely, quietly granting the FBI access to search a database of more than a million profiles starting in late 2018 and defaulting users to opt out rather than requiring them to opt in.

The case of Michael Usry Jr. shows what happens when the system produces a wrong answer. In 2014, investigators wrongly targeted Usry as a murder suspect because his father's donated DNA partially matched crime scene evidence in a genealogy database. After weeks of investigation and mounting anxiety, Usry's own DNA did not match. The actual killer was eventually identified through a proper forensic genetic genealogy process. Usry had done nothing wrong. His father had done nothing wrong. Someone in their biological orbit had volunteered DNA, and the decision sent law enforcement knocking on the wrong door. Usry spent weeks under suspicion for a murder he had nothing to do with, all because of a partial genetic connection he did not create and did not know existed.

The Attorney General Who Told Californians to Delete Their DNA

Two days before 23andMe filed for bankruptcy, California Attorney General Rob Bonta issued an urgent consumer alert telling the state's 23andMe customers to consider deleting their data and requesting destruction of their biological samples. The alert cited Californians' rights under the state's Genetic Information Privacy Act and the California Consumer Privacy Act. The notice included step by step deletion instructions. Bonta stated plainly, "Given 23andMe's reported financial distress, I remind Californians to consider invoking their rights."

Bonta's office backed up the words with action. In 2023, the office partnered with the FTC to take enforcement action against CRI Genetics, resulting in a $700,000 penalty for deceptive marketing practices. In April 2025, Bonta announced a bipartisan consortium of privacy regulators, including attorneys general from seven other states, to coordinate genetic privacy enforcement nationwide. California was among five states maintaining active opposition to the TTAM acquisition even after the bankruptcy court approved the sale, arguing the transfer violated state genetic privacy law because no consumer gave fresh consent.

The multi state response was extraordinary. Within days of the bankruptcy filing, attorneys general in more than a dozen states, including New York, Connecticut, Virginia, Massachusetts, Iowa, Washington, Missouri, and Pennsylvania, issued their own consumer alerts urging customers to delete their data. On June 9 and 10, 2025, a bipartisan coalition of 27 states and the District of Columbia, led by New York Attorney General Letitia James, filed a lawsuit in bankruptcy court to block the genetic data sale. James declared 23andMe "cannot auction millions of people's personal genetic information without their consent." The FTC Chairman sent a letter insisting the company's privacy promises must remain binding on any purchaser. Twenty eight attorneys general and the federal government's top consumer protection agency all lined up on the same side. The bankruptcy court approved the sale anyway.

What You Need to Do Right Now

If you have ever sent your saliva to a genetic testing company, you need to take action today. Not next week. Not when you get around to reading the fine print. Today.

Start by requesting deletion of your data. For 23andMe, which is now operated by TTAM Research Institute: log into your account, go to Settings, scroll to "23andMe Data," click "View," then click "Permanently Delete Data" and confirm through the email they send you. To destroy your physical saliva sample, change your

storage preferences under "Preferences." To revoke your consent to research, go to "Research and Product Consents" in your account settings. For AncestryDNA: go to DNA Settings, click "Delete" next to "Delete DNA Test Results and Revoke Consent to Processing," or call their Member Services line at 801 705 7000. For MyHeritage: go to "Manage DNA Kits" in your account settings or contact their customer support team directly.

Before you delete anything, download a personal copy of your data first. Once you delete, the data is gone from your end permanently. The Consumer Reports Permission Slip app helps with deletion requests for some companies. Understand the hard truth here: data already shared with research partners in a form stripped of your name is beyond retrieval. One University of Iowa professor put the reality simply. "You are never going to find your information at whatever pharmaceutical companies have already received your data, because your name is no longer attached." The deletion you perform today removes your data from the company's active systems. The copies already out in the world are gone for good.

If you are thinking about taking a DNA test, know this before you spit in the tube. Clinical genetic testing through a healthcare provider falls under HIPAA protections; direct to consumer testing does not. Read the consent forms carefully, especially the research consent. Eighty percent of 23andMe customers signed away their research rights, many without understanding the scope of their agreement.

The Electronic Frontier Foundation recommends weighing whether the information you gain from a test justifies the permanent surrender of data defining your biological identity forever. Once your DNA is in the system, you do not get to decide what happens to the company holding your genetic blueprint ten years from now. Ask yourself: if the company goes bankrupt, gets acquired, or gets hacked, are you comfortable with a stranger owning the most personal information your body has ever produced?

Your legal rights depend on where you live. If you are in California, Montana, Indiana, Texas, Illinois, or Maryland, you have dedicated genetic privacy laws with real teeth. Indiana bans genetic testing companies from sharing your data with insurers and employers under any circumstances. Montana requires a warrant for law enforcement access. Illinois lets you sue for up to $15,000 per violation. If your state has no dedicated genetic privacy law, contact your state legislators and tell them you want one.

Tell your friends and family to do the same. Share this chapter. The 23andMe crisis proved voluntary corporate promises to protect your DNA are worth exactly as much as the company standing behind them. And when the company is gone, so are the promises.

Your Body. Your Data. Your Fight.

The 23andMe bankruptcy set a precedent echoing through every future corporate failure involving sensitive biological data. A bankruptcy court determined the genetic information of millions of people was eligible for transfer through a legal structure designed to sidestep state consent laws. The court appointed privacy expert told the judge he was unable to confirm the sale was consistent with the company's own privacy policies. The judge approved the sale anyway. Science journal warned the next case will likely involve a buyer completely unrelated to the original company, meaning even fewer protections for consumers.

No dedicated federal genetic privacy law exists in the United States. GINA, the closest thing we have, leaves life insurance, disability insurance, and long term care insurance completely uncovered. Three major bills

introduced in direct response to the 23andMe crisis remain stuck in committee as of March 2026. State laws provide meaningful protection if you happen to live in the right state, and almost nothing if you do not.

Genetic privacy is different from every other privacy issue in this book because the stakes are permanent and the damage radiates outward in every direction. You are not going to reset your DNA. Your children, your siblings, your parents, and your cousins are exposed when your data is exposed, and none of them signed a consent form. The analytical value of your genetic data grows every year as science advances, meaning information collected today will reveal far more about you under tomorrow's technology than anyone predicts right now.

The companies Americans trusted with this information did not survive long enough to keep their promises. The legal system has not decided whether genetic data deserves a category of protection as unique as the data itself. The decision is still being made, in courtrooms, in state legislatures, and in Congress.

You have a voice in every one of those places. Use your voice. Talk to your family. Talk to your representatives. Because your DNA is the one thing about you never changing, and once someone else owns your genetic blueprint, you are never getting the information back.

Chapter 15: Your Kids Are Being Watched

In Lawrence, Kansas, nine high school students walked into a federal courthouse in August 2025 and filed a lawsuit against their own school district. Their crime, according to the artificial intelligence software installed on every school issued device in the district, was creating art. Gaggle, a monitoring program running silently on school laptops and tablets across 1,500 American school districts, had flagged student artwork as "child pornography." The software also seized journalism files from the student newspaper staff, preventing four editions of the paper from going to print. These were not delinquents. These were artists, writers, and student journalists whose school district had handed their digital lives over to an algorithm that branded them as criminals.

Halfway across the country in Minneapolis, a transgender teenager named Logsdon Wallace sat down to write a school assignment about a painful chapter in his life. He wrote about past suicidal thoughts. He was trying to be vulnerable with his teacher. Gaggle flagged the assignment. Two days later, school officials contacted his parent. "I was trying to be vulnerable with this teacher," he said. "Now the school is contacting my counselor and is freaking out." The system designed to protect him punished him for being honest.

These are not isolated stories. Twenty seven million American students are monitored by a single company called GoGuardian, which tracks browsing histories, documents, search queries, and screen activity in real time, including outside school hours, including inside their own homes. In Baltimore, school police received GoGuardian alerts on nights, weekends, and holidays. Teachers viewed student webcam footage from bedrooms and living rooms without parental consent. A separate company called Gaggle scans billions of student communications annually across six million students, using a combination of artificial intelligence and hourly wage human content moderators to review school provided emails, documents, and chats around the clock.

And here is the part that should stop every parent reading this book in their tracks. A 2023 RAND Corporation study found only "scant evidence" that any of these surveillance tools effectively prevent youth suicide. That is the primary justification schools give for using them. The evidence says the justification does not hold up. Your children are being surveilled at a scale and intensity that would have been unthinkable a decade ago, and the system producing that surveillance has almost no accountability, almost no transparency, and almost no proof that the surveillance works.

This chapter is about what happens when an entire country decides to watch its children instead of protecting them.

The Data Extraction Starts at Birth

By the time an American child turns 13, an estimated 72,000 data points have been collected about that child. The extraction begins with connected baby monitors and smart toys. It accelerates through entertainment apps that harvest location data, browsing behavior, and voice recordings. And then the school system takes over, feeding academic records, behavioral assessments, disciplinary histories, disability information, and counselor notes into a sprawling network of technology vendors whose names most parents have never heard.

The average American school district now deploys 2,591 educational technology tools every year. In the 2016 to 2017 school year, that number was roughly 300. The pandemic drove an eightfold increase, and schools never scaled back. A staggering 96 percent of school recommended apps share student personal information with third parties. The data flowing out of those apps extends far beyond report cards. Browsing histories, search queries, keystroke patterns, documents created, emails sent, calendar entries, and in some cases biometrics and precise geolocation all flow through vendor systems that parents never agreed to and rarely know about.

Parents post an average of 1,300 photos and videos of their children before those children turn 13. Barclays, the financial services firm, projects that this kind of oversharing by parents will account for two thirds of identity theft targeting young people by 2030. Nine hundred fifteen thousand American children fall victim to identity fraud every year. Children are 51 times more likely than adults to have their identities stolen, partly because the theft goes undetected for an average of 12 years. A teenager applies for a first credit card or a student loan and discovers that a stranger has already destroyed the credit file.

One in every 50 American children gets victimized each year. The rate surged 40 percent between 2021 and 2024.

The Law Finally Moves, a Decade Late

On January 16, 2025, the Federal Trade Commission voted unanimously, five to zero, to finalize the first major amendments to the Children's Online Privacy Protection Act Rule since 2013. The updated rule took effect on June 23, 2025, with a compliance deadline of April 22, 2026.

The biggest change expands the definition of "personal information" to include biometric identifiers like fingerprints, retina patterns, voiceprints, gait patterns, and facial templates, along with government issued identifiers like Social Security numbers. The rule now requires separate verifiable parental consent before any company discloses children's personal information to third parties for targeted advertising. That means companies need your specific permission before they sell your child's data to advertisers. Companies must identify the specific third party recipients by name in their notices to parents. Three new consent verification methods were approved, including text based mechanisms, knowledge based authentication, and facial recognition matching against government identification.

Behind the scenes, the rule now requires every operator collecting children's data to maintain a written data retention policy that spells out why the data was collected, why the company needs to keep the data, and when the company will delete the data. Personal information cannot be retained indefinitely. Every operator must also maintain a written information security program with a designated coordinator, annual risk assessments, and regular testing. If you are a parent reading this, you now have the legal right to demand that any company collecting your child's data show you these written policies. Exercise that right.

The FTC also dropped several proposed provisions. Restrictions on push notifications to children were abandoned. A proposal to let educational technology providers obtain parental consent through schools was shelved. FTC Chair Andrew Ferguson identified the rule's failure to create a clear exception for age verification data collection as a significant gap. The FTC partially addressed that gap in February 2026 with a policy statement saying the agency would not pursue operators who collect data solely for age verification purposes, as long as those operators promptly delete the data and maintain reasonable security.

Congress Talks, Nothing Passes

No major federal children's privacy statute has been enacted since COPPA became law in 1998. The Kids Online Safety Act passed the Senate 91 to 3 in July 2024. Ninety one to three. That is as close to unanimous as the United States Senate gets on anything. The bill stalled in the House. Reintroduced in May 2025, the bill was folded into the broader Kids Internet and Digital Safety Act, which the House Energy and Commerce Committee passed 28 to 24 on March 6, 2026. The bill requires platforms to put safeguards in place by default for known minors, provide parents with tools to manage usage, and allow minors to disable addictive features and opt out of algorithmic recommendations.

COPPA 2.0, the Children and Teens' Online Privacy Protection Act, would raise the protected age from under 13 to under 17, ban targeted advertising to children and teens, create an "eraser button" for data deletion, and establish a Youth Marketing and Privacy Division at the FTC. The Congressional Budget Office estimated 164 million dollars in additional enforcement revenues over the next decade. The bill advanced from the Senate Commerce Committee in June 2025. Forty state attorneys general urged Congress to preserve state authority over children's privacy. Neither KOSA nor COPPA 2.0 has reached the President's desk.

The ACLU and the Electronic Frontier Foundation have raised concerns that KOSA could be used to suppress LGBTQ content under politically motivated enforcement. GLAAD reversed its support of the bill in 2025 after changes in FTC leadership. The result is that the patchwork of state laws continues to grow while Congress debates.

The FTC Goes After the Biggest Names

The Federal Trade Commission under both the Biden and Trump administrations has made children's privacy its top enforcement priority, and the results have been striking.

The Epic Games and Fortnite case remains the landmark. In December 2022, Epic Games agreed to pay 520 million dollars. The 275 million dollar COPPA penalty was the largest in FTC history. An additional 245 million dollars went to consumer refunds. The case addressed Fortnite's collection of children's data without parental consent and the game's default settings, which activated voice and text chat matching children with adult strangers. Internal documents showed employees urged the company to change those default settings as early as 2017. The company refused. Children were bullied, threatened, and sexually harassed through the platform. By June 2025, the FTC had distributed more than 200 million dollars in refunds to over 969,000 players.

The pace picked up in 2025. Cognosphere, the maker of Genshin Impact, paid 20 million dollars in January 2025 for collecting children's data without consent and running deceptive loot box schemes. The FTC deemed the game "directed to children" based on its anime style graphics and childlike characters. Disney paid 10 million dollars in September 2025 after failing to properly label child directed YouTube videos featuring Frozen, Toy Story, and Pixar properties as "Made for Kids," which allowed targeted advertising without parental consent. YouTube had flagged over 300 Disney videos as early as mid 2020. Disney kept its original designation policy in place for years after those warnings.

Illuminate Education became the first educational technology company to face federal enforcement in December 2025 after a breach exposed 10.1 million student records, including disability information, disciplinary records, and health data. The company stored student data in plain text despite promising schools that data was encrypted. Third party auditors flagged security vulnerabilities as early as January 2020. The company ignored every warning. When breaches occurred, some school districts did not receive notification for nearly two years. The FTC's consent order permanently banned the company from misrepresenting data security practices and required deletion of unnecessary student data within 90 days.

The Department of Justice, on FTC referral, sued TikTok in August 2024 in what remains the most aggressive pending children's privacy case. The complaint alleges TikTok knowingly allowed millions of children under 13 on the platform despite a 2019 consent order that carried a 5.7 million dollar penalty. TikTok built back doors allowing children to bypass age gates using Google and Instagram credentials. Account reviewers spent an average of five to seven seconds per review when deciding whether a user was a child. TikTok's motion to dismiss was largely denied in November 2025. The case is proceeding through discovery. The penalties sought go up to 51,744 dollars per violation per day.

Smaller cases tell an equally important story. A Chinese robot toy company called Apitor Technology paid 500,000 dollars after its companion app sent children's geolocation data to servers in China. NGL Labs paid 5 million dollars and its founders were personally banned from offering anonymous messaging apps to children after the company sent fake messages to drive engagement. The Sendit app, which knew 116,000 of its users were under 13, faces a pending lawsuit after sending provocative fake messages to children, including messages asking "have you done drugs?" When you see your child playing with a connected toy or downloading an anonymous messaging app, these are the kinds of companies behind those products.

The Surveillance Machine Inside Your Child's School

Let's go back to those school monitoring tools, because this is where the story gets personal for every family with a child in public school.

GoGuardian monitors 27 million students across roughly half of all American K 12 public schools. The company tracks browsing histories, documents, search histories, and screen activity in real time. The monitoring does not stop when the school bell rings. In Baltimore City Public Schools, school police received GoGuardian alerts on nights, weekends, and holidays. Teachers viewed student webcam footage from inside students' homes without consent. Gaggle scans billions of student communications annually across approximately six million students in 1,500 districts, using AI and hourly wage human content moderators to review school provided emails, documents, and chats 24 hours a day, seven days a week. Bark, Securly, and Lightspeed Systems round out the monitoring ecosystem, collectively watching tens of millions of students around the clock.

The Electronic Frontier Foundation's October 2023 investigation, titled "Red Flag Machine," found that GoGuardian's false positives heavily outweigh accurate detections. The system flagged Bible verses containing the word "naked." The system flagged Texas legislature pages about cannabis bills. The system flagged college application websites. The system flagged LGBTQ information. Research from the Center for Democracy and Technology found that 29 to 30 percent of LGBTQ students reported being outed as a result of school activity monitoring. Think about that. Nearly a third of LGBTQ students in monitored schools had their sexual orientation or gender identity exposed because of software their school installed.

In Vancouver, Washington, nearly 2,200 students, representing 10 percent of enrollment, were flagged by Gaggle in a single school year. Forty four percent of teachers reported students contacted by police as a result of monitoring. In Austin, Texas, GoGuardian alerts for "sexual content" go directly to the school district's police department. A 2025 study from the Center for Democracy and Technology found students flagged by monitoring tools were contacted by immigration enforcement.

Turn off school issued devices when your child finishes homework. Make sure your child never logs into personal accounts on a school device. Ask your school's administration for a written copy of the school's monitoring policy, and ask specifically which companies have access to your child's data. If the school cannot answer those questions, that silence tells you everything you need to know.

The Law That Protects Nothing

The Family Educational Rights and Privacy Act, signed into law in 1974, is supposed to protect student records. In practice, FERPA creates a false sense of security.

Through regulatory amendments in 2008 and 2011, the Department of Education, without a Congressional vote, expanded the "school official" exception to allow schools to designate virtually any educational technology vendor as a "school official" with a "legitimate educational interest." That designation allows schools to share your child's data with those vendors without your consent. The Electronic Privacy Information Center described these amendments as the cause of educational data flowing nearly unrestricted from schools to third parties.

FERPA's enforcement mechanism has never worked because the only penalty available, withdrawal of all federal funding from a school, is so severe that the government has never imposed the penalty in 50 years. Not once. There is no private right of action under FERPA. The Supreme Court ruled in 2002 that students and parents cannot sue schools for FERPA violations. Only 12 percent of school websites include any navigation to data privacy information. The Student Privacy Policy Office relies on voluntary compliance. And the executive order directing the closure of the Department of Education puts the primary FERPA enforcement body at risk of elimination entirely.

When educational technology companies fail or get acquired, student data follows the money. When ConnectEDU went bankrupt in 2014 with 20 million student records, the FTC had to step in to prevent the sale of that student data to a venture capital fund. When AllHere Education collapsed in 2024 after building an AI chatbot for the Los Angeles Unified School District, a whistleblower revealed student data had been processed on offshore servers in Japan, Sweden, and France. The CEO was arrested on fraud charges. Anthology, the parent company of Blackboard, filed for Chapter 11 bankruptcy in September 2025 with more than one billion dollars in debt, putting student data from hundreds of institutions into bankruptcy proceedings. If your child's school uses one of these platforms, ask the school what happens to your child's data if the company goes under. You deserve an answer, and right now, the law does not require the school to give you one.

Fifty States, Fifty Different Sets of Rules

The federal vacuum has produced a sprawling patchwork of state laws and an equally sprawling body of litigation challenging those laws. NetChoice, the technology industry trade group whose members include Google, Meta, Amazon, TikTok, and Snap, has filed lawsuits in at least 15 states and won most of them.

California's Age Appropriate Design Code Act has been mostly blocked by courts since its passage in 2022. In March 2026, the Ninth Circuit affirmed injunctions against five of six challenged provisions, finding terms like "best interests of children" and "materially detrimental" unconstitutionally vague. Laws in Arkansas, Ohio, and Louisiana have been permanently struck down on First Amendment grounds. Georgia and Utah's laws are preliminarily blocked.

Florida's HB 3 is the sole survivor. After a district court blocked the law, the Eleventh Circuit stayed the injunction in November 2025 in a two to one decision, finding the law likely content neutral and satisfying intermediate scrutiny. That makes Florida's law the only social media access law to survive appellate review as of March 2026.

A newer wave of laws aims to survive constitutional challenge by focusing on platform design rather than content restrictions. Maryland's Kids Code has been in effect since October 2024 and survived a motion to dismiss in November 2025. Nebraska, Vermont, and South Carolina enacted their own design codes in 2025 and 2026. New York's Child Data Protection Act took effect in June 2025, prohibiting operators from collecting or selling personal data of users under 18 unless strictly necessary. If you live in one of these states, learn the specific protections your state provides and use them. File complaints with your state attorney general when companies violate those protections.

The Age Verification Trap

The Supreme Court's six to three decision in Free Speech Coalition v. Paxton in June 2025 reshaped the legal terrain by upholding Texas's age verification requirement for adult content. The ruling applies narrowly to sexually explicit material obscene to minors. Courts continue to strike down age verification requirements for general social media platforms.

The deeper problem is what privacy advocates call the "age verification trap." To protect children's privacy, every system proposed so far requires collecting sensitive data about everyone, including adults. Facial age estimation technology from companies like Yoti processes over 850 million age checks worldwide, with a mean error of about 1.2 years for people around age 18. That technology requires a selfie. Government ID upload creates centralized databases that become targets for hackers. One vendor breach already exposed 70,000 government ID records. When Discord announced mandatory age verification in February 2026 requiring selfies or government identification, the backlash was severe enough to delay the rollout worldwide.

The FTC's February 2026 enforcement policy attempted to create a safe harbor for operators collecting data solely for age verification, as long as those operators promptly delete the data and maintain reasonable security. The technology industry remains deeply divided. Meta pushes for app store level verification, which would shift the burden to Apple and Google. Apple and Google prefer device level approaches using age range signals that never expose a child's birthdate to app developers. Roughly half of American states now mandate age verification for adult content, and at least 17 states have enacted social media access laws for minors. Most of those laws face legal challenges.

The Rest of the World Figured This Out

The United Kingdom's Age Appropriate Design Code, in force since September 2021, establishes 15 enforceable design standards for any online service likely to be accessed by children under 18. An independent assessment documented 44 platform changes to community standards enforcement, 43 changes to content and advertising safeguards, and 31 changes to privacy settings across major platforms as a direct result of the law. TikTok changed default privacy to private for users ages 13 to 15. Instagram made all accounts created by users under 18 private by default. YouTube disabled autoplay for minors. One major technology company told regulators that the UK Code had a greater effect than GDPR enforcement actions.

The European Union enforces children's protections with financial penalties that make American fines look like pocket change. Ireland's Data Protection Commission fined TikTok 345 million euros in 2023 for children's privacy violations, fined Instagram 405 million euros in 2022 for default public settings on children's accounts, and fined TikTok again for 530 million euros in 2025 for illegal data transfers to China. The EU's Digital Services Act requires private by default settings for minors and prohibits profiling based advertising targeted at children. Australia went the furthest in December 2025, enacting the world's first outright ban on social media for children under 16, with penalties up to 49.5 million Australian dollars.

The United States lacks every major protection that exists in those countries. No federal design based regulation for children. No federal privacy by default requirement for minors. No anti profiling rules for children. No data minimization requirements. No mandatory impact assessments for children's data. COPPA fines are modest compared to the EU's penalty of four percent of global annual turnover. When you hear people say "we do not know how to protect children online," remember that the United Kingdom, the European Union, and Australia have already shown exactly how.

62 Million Student Records and the Ransom That Solved Nothing

The PowerSchool breach of December 2024 stands as the defining data disaster of this era. Between December 19 and December 28, hackers accessed PowerSchool's customer support portal using a compromised employee credential. The system lacked mandatory multi factor authentication. The hackers pulled 62 million student records and 9.5 million educator records, including names, dates of birth, Social Security numbers, medical alert information, disciplinary records, and individualized education plans. PowerSchool serves 75 percent of the American education market.

PowerSchool paid the ransom. The hackers provided a video showing the data being deleted. On May 7, 2025, extortion emails containing samples of the stolen data arrived at schools in Canada and North Carolina. The data had not been deleted. Stolen transcripts appeared on dark web marketplaces priced between 50 and 300 dollars per record. Matthew Lane, a 19 year old Massachusetts college student who demanded 2.85 million dollars from PowerSchool, was sentenced on October 14, 2025 to four years in federal prison and ordered to pay 14 million dollars in restitution. Prosecutors acknowledged that money would likely never be collected. More than 100 school districts sued PowerSchool.

The breach exposed a systemic truth that every parent needs to understand. Schools retain student records for decades. There is no federal law requiring schools to delete old data. That decades long retention policy meant the breach captured historical records going back years, affecting students who had long since graduated. Education is now the most attacked sector globally, with 4,388 cyberattacks per week targeting

schools in the second quarter of 2025, a 31 percent year over year increase. If your child attends a school that uses PowerSchool or any similar platform, freeze your child's credit report at all three bureaus today. Do not wait. The breach has already happened.

Teaching Your Children What No Law Will Teach Them

Since the legal protections are not keeping pace with the threats, families need to fill the gap themselves. Common Sense Media's Digital Citizenship Curriculum, accessed by 1.3 million educators in more than 88,000 schools, provides the most widely adopted framework for teaching children about privacy, digital footprints, and information literacy from kindergarten through 12th grade. Developed in collaboration with Harvard's Project Zero, the curriculum teaches decision making frameworks rather than rigid rules, because the specific platforms and threats change faster than any list of do's and don'ts.

Age appropriate approaches matter. Children ages five to seven learn not to share personal information and to ask a parent before downloading anything. By ages 11 to 13, students start to understand how companies make money from their data, learn to evaluate terms of service, and practice managing privacy settings on platforms. High schoolers study algorithmic profiling, data broker ecosystems, and their legal rights. Leah Plunkett of Harvard Law School puts the point sharply. Privacy literacy needs to go beyond what children post. Children also need to understand what gets posted about them and what systems collect about them invisibly.

Here is what you and your family should be doing right now. Keep school issued devices turned off when your child is not using them for schoolwork. This prevents surveillance tools from activating during personal time. Never let your child log into personal email, social media, or other personal accounts on a school device. Freeze your child's credit report at Equifax, Experian, and TransUnion. Conduct regular privacy audits by searching your child's name online and checking what comes up. Model good privacy behavior yourself.

Forty one percent of parents say they would share less about their children online if they could start over. Start over now. Every photo, every check in, every proud post about a child's accomplishments adds another data point to a profile that will follow that child into adulthood.

The System Needs to Change, and You Need to Act

American children grow up inside a data extraction machine that operates at every stage of their lives. Entertainment apps collect behavioral data through addictive design. Schools funnel academic, behavioral, and emotional information through thousands of vendor relationships governed by a federal law that has never been enforced. Data brokers aggregate all of the data and offer lists of nearly six million high school students along with data on children as young as two. Monitoring software runs around the clock, flags LGBTQ students for who they are, sends police to children's homes for after hours internet searches, and produces no verifiable evidence of preventing the harms schools cite to justify the surveillance.

The 2025 COPPA amendments and the FTC's aggressive enforcement represent genuine progress. The core architecture of childhood surveillance remains intact. COPPA protects only children under 13. FERPA's enforcement mechanism has never been used. Neither KOSA nor COPPA 2.0 has become law. The states pushing hardest face systematic legal challenges from the industry trade groups whose members profit from the current system. The United Kingdom, the European Union, and Australia have already demonstrated that

enforceable, design based regulation works. Platforms changed their behavior when those countries confronted them with real penalties.

America knows how to protect its children's data. The models exist. The question is whether we will demand that our elected leaders act on what the rest of the world has already proven. Contact your members of Congress and tell them to pass COPPA 2.0 and the Kids Online Safety Act. File complaints with your state attorney general when companies violate your state's children's privacy laws. Show up at your next school board meeting and ask the superintendent, on the record, exactly which companies have access to your child's data and what happens to that data if those companies go bankrupt. Demand written answers. The only way this changes is if parents refuse to accept silence as an answer.

Your children deserve better than this. And deep down, you already know that.

Chapter 16: It's Your Phone Number And It's Putting Everything at Risk

Robert Ross, a former Apple engineer and Silicon Valley angel investor, was sitting at his desk in San Francisco on an ordinary Friday afternoon in October 2018 when something small and strange happened. His iPhone stopped working. The screen still glowed, the apps still loaded, the Wi-Fi still connected. The cellular signal, the bars in the top corner of his screen, simply vanished. "No Service." He assumed a tower was down or a software glitch needed a restart. He set the phone aside.

Then a notification appeared. A withdrawal request from one of his financial accounts. He had not made a withdrawal. He picked up the phone again, tried to call his bank, and realized he had no ability to place a call. No ability to send a text. No ability to receive the security codes his bank needed to verify his identity. In the next twenty minutes, someone on the other side of the country drained $500,000 from his Coinbase account and $500,000 from his Gemini account. One million dollars. Gone. His daughters' college fund. Ninety percent of his net worth. Wiped clean in the time most people spend scrolling through morning emails.

Rob Ross did not get hacked through some sophisticated zero-day exploit or cutting-edge malware. No one broke into his home. No one stole his laptop. A criminal simply called his wireless carrier, pretended to be him, and asked customer service to transfer his phone number to a different SIM card. That was the entire attack. A phone call. A convincing lie. And then every account protected by a text message verification code became an open door.

This crime has a name. SIM swapping. And the phone number you carry around in your pocket, the one tied to your bank, your email, your retirement accounts, your social media, your health insurance portal, and your tax filings, is the single most dangerous piece of personal information in your digital life.

What SIM Swapping Looks Like From the Inside

Your wireless carrier assigns your phone number to a tiny chip inside your phone called a SIM card. When you make a call, send a text, or receive a verification code from your bank, the carrier's network routes everything to the device holding your SIM. SIM swapping works because carriers allow customer service representatives to reassign your number to a new SIM card, and criminals have figured out exactly how to make that happen without your knowledge or permission.

The attack comes in two flavors. In a traditional SIM swap, the criminal contacts your carrier directly, either by phone or by walking into a retail store with a fake ID, and claims to be you. The story is always some version of the same thing: lost phone, damaged phone, need a new SIM. The representative runs through a few security questions, the criminal answers them using information purchased from data brokers or harvested from social media, and the swap goes through. Your phone dies. Their phone lights up with your number.

The second version is called port-out fraud. Instead of staying on the same carrier, the criminal opens an account with a different wireless provider and requests that your number be "ported," or transferred, to the new account. The federal government actually mandates that carriers allow number porting so consumers have the freedom to switch providers. Criminals turned that consumer protection into a weapon.

Both versions follow the same playbook once the number moves. The criminal opens your email app, clicks "Forgot Password," and requests a verification code sent by text. The code arrives on their phone. They reset your email password. From your email, they do the same thing to your bank. To your investment accounts. To your cryptocurrency exchange. To anything and everything connected to that phone number. A Princeton University study in 2020 tested five major carriers and found that attackers succeeded on 80 percent of first attempts, because carriers relied on security questions whose answers were easy to find online or purchase from data brokers.

The entire attack, from the first call to your carrier to the last dollar leaving your account, takes less than an hour. For Rob Ross, the financial devastation took twenty minutes.

The People Inside the Phone Companies Who Help Make This Happen

The scariest part of SIM swapping is not the criminals calling from the outside. The scariest part is the employees helping from the inside.

The FBI identified insider corruption as one of the three primary methods criminals use to pull off SIM swaps. Criminal networks recruit carrier employees through Telegram, offering anywhere from $300 to $5,000 per swap. In 2024, T-Mobile and Verizon employees across the country received unsolicited text messages from unknown numbers offering cash for help executing swaps. One message, obtained during an investigation, read simply: "I got your number from the T-Mo employee directory. I'm looking to pay someone up to $300 per sim swap done."

One criminal wrote a software script to search Twitter for posts from people who mentioned working at cell phone carriers. He would then send direct messages offering bribes. He later estimated that about ten percent of the employees he contacted agreed to participate. He bragged to an interviewer that at one point in time, he had a person at every major carrier on his payroll.

Jonathan Katz, a 42 year old store manager at a wireless carrier in New Jersey, used his manager-level system access to execute SIM swaps for $1,000 each plus a cut of the profits. He pleaded guilty in 2024 to five counts of computer fraud. In Manhattan, District Attorney Alvin Bragg brought charges against a ring that included AT&T and T-Mobile retail workers who used their access to steal $435,000 from victims between October 2021 and July 2022. In the case of a gang called "The Community," three carrier insiders were charged, including a Verizon employee who earned a total of $3,500 in bribes and two AT&T contractors in Tucson. Those contractors became the first telecom employees in American history to face federal charges for participating in SIM swap fraud.

Think about what this means for a moment. The people you trust to manage your wireless account, the people sitting behind the counter at the store where you bought your phone, are being actively recruited to sell out your identity. The carrier is not the wall protecting you. In many cases, the carrier is the unlocked door.

How a Phone Number Became the Key to Your Entire Life

Here is the question worth sitting with: How did a ten-digit phone number become the master key to everything you own?

In the mid-2000s, companies needed a way to add a second layer of security beyond passwords. Passwords alone were failing. People reused them across accounts, chose weak ones, and fell for phishing emails. The solution the industry landed on was sending a text message with a one-time code to your phone. Type in the code, prove you have the phone, get into your account. SMS two-factor authentication was born.

The approach spread everywhere because sending a text requires nothing special. No app to download. No device to buy. No technical knowledge. Every phone sold in the last two decades receives text messages. Banks adopted SMS verification. Email providers adopted SMS verification. Social media platforms, cryptocurrency exchanges, government portals, healthcare logins, retirement account dashboards, tax filing services. All of them. The phone number became the default proof of identity across the entire American digital economy.

Today, 56 percent of organizations worldwide support SMS-based verification codes. Among people who have turned on two-factor authentication, 41 percent rely on text messages as their verification method. Coinbase, one of the largest cryptocurrency exchanges in the world, revealed in late 2022 that 95 percent of all account takeovers on its platform involved customers who relied on SMS-based authentication.

The deeper problem goes beyond authentication. Phone numbers also serve as the default account recovery mechanism for most services. Even when someone switches to a more secure authentication method like an app-based code generator, most platforms keep SMS as a backup option in case you lose access to your primary method. That backup option creates a permanent vulnerability. A criminal who takes over your phone number bypasses every other security measure you have in place.

The National Institute of Standards and Technology, the federal agency responsible for cybersecurity guidelines, first warned the country about this problem in July 2016. A draft of their authentication standards called SMS verification "deprecated" and said future versions of the guidelines would no longer allow its use. The final version, published in 2017, softened the language slightly and labeled SMS a "restricted authenticator." The industry's response, by and large, was to keep doing exactly what they were doing. Nine years later, in 2026, SMS text message verification remains the dominant authentication method across the American financial system.

The FCC Finally Stepped In. The Carriers Pushed Back.

On November 15, 2023, all five FCC commissioners voted unanimously to adopt new rules specifically targeting SIM swap and port-out fraud. The order required carriers to verify a customer's identity through secure authentication before processing any SIM change or number transfer. Carriers had to notify customers immediately when a change was requested. Every carrier had to offer free account locks allowing customers to freeze their numbers against unauthorized transfers. Carriers had to train employees on fraud detection and retain records of SIM change requests and fraud incidents for three years.

These rules made sense. They addressed a real problem. And the wireless industry immediately began pushing to delay them.

The original compliance deadline was July 2024. Industry groups petitioned for more time. The FCC pushed the deadline back, tying full compliance to a federal paperwork review process. As of March 2026, the exact

final compliance date remains unconfirmed in publicly available federal records. The rules exist on paper. Whether every carrier has fully adopted them remains an open question.

The FCC has also imposed serious financial penalties on carriers for security failures. T-Mobile agreed to pay $31.5 million in September 2024 to settle investigations into data breaches in 2021, 2022, and 2023. The 2021 breach alone exposed the personal information of approximately 76 million customers. The 2022 breach involved an illegal SIM swap of a T-Mobile employee's account that gave hackers access to internal company systems. TracFone, a Verizon subsidiary, paid $16 million in July 2024 for vulnerabilities that allowed unauthorized port-outs. AT&T paid $13 million for a vendor cloud breach. The Second Circuit Court of Appeals in September 2025 upheld a $46.9 million FCC penalty against Verizon for failing to protect customer information, a ruling with direct implications for carrier responsibility in SIM swap cases.

These fines are a start. For companies generating tens of billions in annual revenue, they are also a cost of doing business. The real question is whether the new rules will actually prevent your phone number from being stolen the next time a criminal calls customer service with your name and your Social Security number.

The Losses Are Staggering and the Stories Are Devastating

The FBI's Internet Crime Complaint Center has tracked SIM swapping as a distinct category since 2018. The numbers paint a grim picture. From 2018 through 2020, the center received about 320 complaints totaling roughly $12 million in losses. In 2021, complaints surged to over 1,600 with losses reaching $68 million. The 2022 peak brought more than 2,000 complaints and $72.7 million in reported losses. Recent years show a decline in reported complaints, with about 980 in 2024 and $26 million in losses.

Those numbers, as bad as they are, dramatically understate the real damage. The FBI itself has found, through other investigations, that only about 20 percent of cybercrime victims report the crime to law enforcement. Many SIM swap victims report the resulting bank fraud or identity theft to their financial institution rather than to the FBI. The true scope of the problem is almost certainly multiples of what the official statistics show.

The scale of individual attacks keeps growing. On November 11, 2022, the same day the cryptocurrency exchange FTX filed for bankruptcy, a 26 year old from suburban Chicago allegedly orchestrated the largest SIM swap theft in history. Robert Powell, operating under the online name "ElSwapo1," directed a 23 year old co-conspirator named Emily Hernandez to fly to Texas, walk into a wireless store with a fake ID bearing her photograph and an FTX employee's personal information, and execute a SIM swap. Over $400 million in Bitcoin disappeared from FTX within hours. Blockchain investigators traced portions of the stolen funds through mixing services where they were combined with money from criminal organizations linked to Russia. Hernandez pleaded guilty to wire fraud. She had been paid $2,500 for her role in a $400 million theft.

In May 2025, federal prosecutors unsealed a RICO conspiracy indictment, the same organized crime statute originally written for the mafia, against Malone Lam and 12 co-defendants for stealing over $263 million in cryptocurrency. In a single August 2024 attack, the group stole more than 4,100 Bitcoin from one victim in Washington, D.C. The stolen money funded nightclub tabs as high as $500,000 in a single evening, a fleet of 28 exotic cars with price tags between $100,000 and $3.8 million, and cash mailed inside Squishmallows stuffed animals. Nine defendants had pleaded guilty by early 2026.

In January 2024, a man named Eric Council Jr. used a portable ID card printer to create a fraudulent driver's license, SIM-swapped an employee associated with the Securities and Exchange Commission, and used the hijacked number to post a false announcement on the SEC's official social media account claiming that Bitcoin exchange-traded funds had been approved. Bitcoin's price spiked over $1,000 within minutes before crashing more than $2,000 when the SEC issued a retraction. The real ETF approval happened the very next day. Investigators found Google searches on Council's laptop including "how can I know for sure if I am being investigated by the FBI." He received 14 months in prison.

The Data Brokers Who Make the Crime Possible

Every SIM swap starts with information. The criminal needs your full name, your date of birth, your Social Security number, your address, your carrier account PIN, and your answers to security questions. Where does all of that come from? Largely from the data broker industry covered earlier in this book.

Approximately 2,400 data brokers operate in the United States, and they collect an average of 1,000 data points on each person with an online presence. They sell names, addresses, phone numbers with carrier identification, dates of birth, Social Security numbers, known family members, email addresses, and answers to the exact security questions carriers use to verify identity. Everything a SIM swapper needs arrives in a single purchase.

Security journalist Brian Krebs documented a cybercriminal who ran a Telegram bot selling access to a data broker's records for $8 to $40 per lookup. The criminal averaged 100 lookups per day through Telegram and 400 per day through direct system access. Most of his customers came from SIM swapping forums. The National Public Data breach in April 2024 exposed up to 2.9 billion records containing names, addresses, phone numbers, dates of birth, and Social Security numbers. The company shut down entirely by December 2024 under the weight of lawsuits.

This is the connection that does not get talked about enough. SIM swapping is not a standalone crime. The data broker industry, operating legally and selling your personal details for pennies, provides the raw material that makes every SIM swap possible. When your information sits in a database that anyone with a credit card or a cryptocurrency wallet can access, the question is not whether someone will try to steal your phone number. The question is when.

California's DELETE Request and Opt-Out Platform, called DROP, launched on January 1, 2026. The platform gives California residents a single place to request removal from all registered data brokers in the state. If you live in California, go to the DROP platform today and submit your request. If you live anywhere else, search for each major data broker individually and submit removal requests directly. Every piece of personal information you remove from these databases makes a SIM swap harder to execute.

What You Need to Do Right Now to Protect Yourself

Every major wireless carrier in America now offers free tools to lock your phone number in place. Not one of these tools comes turned on by default. You have to activate them yourself, and doing so takes about five minutes.

If you use T-Mobile, open the T-Life app or log in to T-Mobile.com. Turn on SIM Protection, which blocks unauthorized SIM swaps on your line. Then turn on Port Out Protection under Services, which prevents your number from being transferred to another carrier. Then set a 6 to 15 digit account PIN through the app, the website, or by calling 611. That PIN will be required for all account changes.

If you use AT&T, open the myAT&T app and activate the Wireless Account Lock. AT&T launched this feature in July 2025, and the lock blocks 12 different types of account changes including SIM swaps, eSIM swaps, and number porting. Even AT&T employees cannot override this lock while the feature is active. Also update your passcode, because the default is the last four digits of your Social Security number.

If you use Verizon, open the My Verizon app and turn on both Number Lock and SIM Protection. Number Lock prevents port-outs to other carriers. SIM Protection prevents internal SIM changes. Both need to be active at the same time because they address different risks. Verizon also imposes a 15-minute cooling-off period after SIM Protection is turned off before any SIM change goes through.

Do this today. Right now. Before you finish this chapter. Call your family members tonight and walk them through the same steps. Your parents, your adult children, your siblings. Anyone who has a phone number connected to a bank account needs these protections turned on.

The second step is to move away from text message verification codes entirely. In December 2024, the FBI and the Cybersecurity and Infrastructure Security Agency issued joint guidance telling Americans point-blank: do not use SMS as a second factor for authentication. This guidance came after the Salt Typhoon cyberattack, in which Chinese government-affiliated hackers penetrated the core networks of AT&T, Verizon, T-Mobile, and Lumen Technologies and gained access to unencrypted text messages, including authentication codes.

The strongest replacement is a hardware security key, a small USB device like a YubiKey (about $50) or Google Titan key (about $30) that you plug into your computer or tap to your phone when logging in. The next strongest option is a passkey, a free, built-in feature on most modern phones and computers that ties your login to your device and your fingerprint or face. Approximately 69 percent of consumers now have at least one passkey, and over one billion passkeys have been created worldwide. Microsoft made passkeys the default for new accounts in May 2025. The third option is an authenticator app like Google Authenticator, Microsoft Authenticator, or the open-source Ente Auth, all of which generate codes on your device rather than sending them over the cellular network.

Log in to your most important accounts this week, your bank, your primary email, your investment accounts, and check your security settings. If the option exists to switch from SMS verification to an authenticator app, a passkey, or a hardware key, make the switch. Every account you move off SMS is one more door that a SIM swap cannot open.

One frustrating reality remains. Many of the largest financial institutions in America, including Chase, Wells Fargo, Citibank for most account functions, and Vanguard, still offer only SMS-based verification with no app or passkey alternative. If your bank falls into this category, call them. Tell them you want stronger authentication options. File a formal complaint. The more customers demand better security, the faster these institutions will move.

The Legal Tide Is Turning

For years, carriers treated SIM swap losses as somebody else's problem. Victims filed lawsuits. Carriers buried them in arbitration. Courts dismissed claims. The legal ground is shifting in a meaningful way.

The Ninth Circuit Court of Appeals ruled in September 2024 that victims of SIM swap fraud have the right to sue their wireless carrier under Section 222 of the Federal Communications Act. That case, brought by Michael Terpin after losing $24 million through an AT&T SIM swap, was the first appellate ruling to establish that carriers owe a legal duty to protect customers from this specific type of fraud.

In March 2025, an arbitrator ordered T-Mobile to pay $33 million to Joseph "Josh" Jones, a victim who lost over $38 million in cryptocurrency after his T-Mobile account was compromised in February 2020. Jones had set up heightened security on his account with an 8-digit PIN. The attackers bypassed the PIN and executed the swap. T-Mobile attempted to seal the arbitrator's findings, and attorneys accused the company of trying to hide evidence of systemic security failures. That $33 million award is the largest SIM swap-related recovery on record.

The Terpin case against AT&T was set for trial on March 3, 2026. If a jury finds AT&T liable, the precedent will send a clear message to every carrier in America: letting a criminal walk off with a customer's phone number carries a financial cost that makes the necessary security investments look cheap by comparison. That kind of financial accountability is the one force that consistently changes corporate behavior.

Your Phone Number Is Not Just a Number

Everything you have read in this chapter comes down to one uncomfortable truth. Your phone number has been turned into something its designers never intended. A ten-digit sequence originally meant to route voice calls now functions as the password to your financial life, your personal communications, your health records, your government accounts, and your digital identity.

The carriers know this. The banks know this. The federal government has known this since at least 2016 when NIST first warned the country to move away from SMS-based security. And still, in 2026, the default setting on your wireless account leaves your number unprotected. The default setting on most major bank accounts still uses text messages as the only form of two-factor authentication. The data broker industry still sells the personal information that makes SIM swaps possible for less than the price of a fast food meal.

You did not create this system. You did not choose to have your entire financial life hinge on whether a customer service representative at a wireless store can tell the difference between you and a criminal with a convincing story and a fake ID. You did not ask for your Social Security number, your date of birth, and your mother's maiden name to be available for purchase on the open market.

You do have the power to lock your carrier account down today. You do have the ability to switch your most important accounts off SMS verification this week. You do have the right to demand that your bank offer real security instead of a system the FBI has told you to stop using.

Rob Ross, the Apple engineer who lost a million dollars in twenty minutes, went on to found a nonprofit called StopSIMCrime.org dedicated to raising awareness and pushing for stronger protections. He did not accept the loss as inevitable. He fought back.

Five minutes. That is all the time you need to lock your phone number down. Pick up your phone, open your carrier's app, and turn on every protection available to you. Then move your most important accounts off text-message authentication. Then tell the people you love to do the same. The criminals who steal phone numbers count on one thing above all else: the assumption that ordinary people will never bother to take these steps.

Prove them wrong.

Chapter 17: Ransomware and Data Breaches Can Destroy Your Life

You open your mailbox on a Tuesday afternoon. Between the electric bill and a credit card offer, you find a letter from a company you do not remember doing business with. The envelope looks official. The letter inside is two pages long. The first line reads: "We are writing to inform you of a recent security incident." Your stomach drops. You scan the rest of the letter. It tells you that your name, your date of birth, your Social Security number, your health insurance policy number, and your banking information were "potentially accessed by an unauthorized third party." The letter does not tell you when the breach happened. The letter does not tell you how the attackers got in. The letter offers you twelve months of free credit monitoring and suggests you "remain vigilant."

You fold the letter. You set it on the kitchen counter. And you do nothing. Because this is the third one you have received this year.

If that scenario sounds familiar, you are in the majority. Eighty percent of American consumers received at least one breach notification letter in 2025 alone. Forty percent received between three and five separate notices in a single year. Since 2005, more than twelve billion breach notification letters have been sent in the United States, a number that exceeds the entire population several times over. The data breach is no longer a rare event that happens to someone else. The data breach is a permanent feature of American life, and the consequences stack up silently over years and decades like interest on a debt you never agreed to carry.

This chapter is about what happens when the companies holding your most sensitive information lose control of it. This chapter is about the scale of the problem, the specific failures behind the worst breaches in recent history, the criminal industry that now profits from those failures, and the gap between what the law promises you and what the law actually delivers. Most of all, this chapter is about what you need to do right now, today, to protect yourself and your family, because the system designed to protect you is broken.

Three Thousand Breaches a Year and Counting

The numbers tell a story that defies common sense. In 2025, the Identity Theft Resource Center documented 3,322 data breaches across the United States, a new all time record and a 79 percent increase over just five years. Three consecutive years, 2023 through 2025, exceeded 3,000 compromises annually, a threshold never previously approached. In 2024, six massive breaches drove victim notifications to 1.35 billion, more than four times the entire U.S. population. That means on average every American received four separate notifications in a single year that their data had been stolen.

The upward climb has been relentless. In 2017, there were 1,506 reported breaches. That number dipped slightly in 2018, climbed again in 2019, and dropped during the pandemic disruption of 2020. Then everything accelerated. Breaches surged 68 percent in 2021. They doubled again by 2023. They have stayed above 3,100 every year since. The attacks have also become more secretive. In 2020, nearly all breach notification letters explained what caused the attack. By 2025, only 30 percent did. Seven out of ten notification letters now tell you your data was stolen and refuse to explain how the attackers got in.

The financial damage is staggering. The average cost of a data breach in the United States hit $10.22 million in 2025, an all time high and more than double the global average. Healthcare remains the most expensive industry for breaches, a distinction it has held for fourteen straight years. The FBI recorded $16.6 billion in

cybercrime losses across 859,532 complaints in 2024, a 33 percent jump from the prior year. And the average breach takes 241 days from intrusion to containment. That means a breach discovered today started nearly eight months ago.

The history books are filled with breaches so large they reshaped entire industries. The Yahoo breach of 2013, not disclosed until late 2016, ultimately encompassed all three billion user accounts. The Equifax breach of 2017 exposed the Social Security numbers, birthdates, and addresses of 147.9 million Americans, roughly 40 percent of the entire population. Four Chinese military hackers were eventually indicted. Equifax paid $1.38 billion in total costs. Capital One lost 106 million records in 2019 after a former cloud services employee walked through a misconfigured firewall. T Mobile has suffered at least eight known breaches between 2018 and 2023, including one affecting 76.6 million people.

The drumbeat never stops. In January 2025, PowerSchool, a K through 12 education technology platform, disclosed that attackers used a stolen contractor login to access 62 million student records and 10 million teacher records. Blue Shield of California exposed 4.7 million records through a Google Analytics misconfiguration. DaVita, the dialysis provider, suffered a ransomware attack that compromised 2.7 million patient records. And researchers identified a compilation of 16 billion stolen credentials aggregated from malware across thousands of sources, the largest credential dump ever assembled.

One Missing Checkbox Shut Down American Healthcare

The February 2024 ransomware attack on Change Healthcare stands as the single most consequential cyberattack in the history of American healthcare. Not because the attackers were geniuses. Because the failure was breathtakingly simple.

Change Healthcare processes roughly 15 billion healthcare transactions every year, totaling more than $1.5 trillion in claims. The system connects 1.6 million health professionals, 70,000 pharmacies, and 8,000 healthcare facilities. When UnitedHealth Group acquired Change Healthcare in October 2022 for $13 billion, the American Hospital Association later described the combined entity as the predominant source of more than 100 critical functions that keep the entire U.S. healthcare system running.

On February 12, 2024, an affiliate of the ALPHV/BlackCat ransomware gang gained access through a Citrix remote access portal. The portal lacked multi factor authentication. That is a basic security control, the digital equivalent of a deadbolt on your front door. The attackers moved through the network for nine days. They stole an estimated six terabytes of data. On February 21, they deployed ransomware, encrypting systems and forcing Change Healthcare to shut down its entire network. Billing systems, pharmacy claims processing, prior authorizations, and electronic prescribing collapsed simultaneously across the country.

The impact hit patients immediately. Thousands of pharmacies could not process prescription claims for weeks. Military pharmacies worldwide went dark. Patients faced a choice between paying full out of pocket prices or going without their medications. A survey of roughly 1,000 hospitals found that 94 percent reported financial damage, 74 percent reported direct harm to patient care, and 33 percent said the attack disrupted more than half of their revenue. Eighty percent of physician practices lost revenue from unpaid claims. Fifty five percent of practice owners used their own personal savings to cover office bills and payroll. Industry estimates placed provider losses at roughly $100 million per day.

Around March 1, 2024, UnitedHealth Group paid approximately $22 million, about 350 Bitcoin, to the ransomware gang. CEO Andrew Witty later described the decision as one of the hardest he had ever made. The payment accomplished nothing. The gang's leadership pocketed the entire $22 million, refused to share it with the affiliate who actually conducted the attack, posted a fake FBI seizure notice on their dark web site, and vanished. The unpaid affiliate retained all the stolen data and launched a second extortion attempt through a new criminal operation called RansomHub, publishing partial stolen files as proof. UnitedHealth reportedly refused to pay a second time.

On May 1, 2024, Witty testified before the Senate Finance Committee and the House Energy and Commerce Subcommittee for more than four hours. The testimony revealed damning facts. Witty acknowledged the server through which attackers entered was not protected by multi factor authentication. Senator Ron Wyden called the failure "cybersecurity 101." Senator Thom Tillis brought a copy of Hacking for Dummies to the hearing. Senator John Barrasso asked why a small, financially struggling hospital in Wyoming managed to turn on multi factor authentication when a company generating nearly $100 billion in quarterly revenue did not. Witty disclosed that some of the company's technology dated back 40 years, and that the primary and backup systems were not isolated from each other, so both were directly compromised in the same attack.

The final scope of the breach dwarfed every initial estimate. Change Healthcare first filed a report with the Department of Health and Human Services listing 500 affected individuals. By October 2024, that number grew to 100 million. By January 2025, it reached 190 million. The confirmed final total, reported July 31, 2025: 192.7 million individuals. That is roughly two thirds of the American population and by far the largest healthcare data breach in U.S. history. The exposed data included health insurance information, medical records, prescriptions, diagnoses, Social Security numbers, and banking details.

UnitedHealth Group reported $3.09 billion in total breach related costs for 2024. As of March 2026, 78 class action lawsuits have been consolidated into federal multi district litigation in Minnesota. The Nebraska Attorney General filed the first state lawsuit, and a motion to dismiss was denied in November 2025. The HHS Office for Civil Rights investigation remains ongoing. No penalties have been announced.

The deeper lesson reaches beyond any single security failure. Change Healthcare's dominance created a single point of failure for the entire healthcare payment system in America. Nearly every hospital in the country ran transactions through this one system. When one company controls that much critical infrastructure, one missing security checkbox brings the whole thing down.

A Home Office in Florida Held Your Social Security Number

If the Change Healthcare breach showed what happens when a giant company fails at basic security, the National Public Data breach revealed something equally alarming: a shadow industry of data brokers collecting sensitive records on virtually every American adult with almost no security, almost no regulation, and almost no accountability.

National Public Data, operated by a Florida company called Jerico Pictures, was a background check data broker run by a single individual, a retired sheriff's deputy, from a home office in Pompano Beach. The company had fewer than 25 employees, annual revenue of roughly $1.15 million, and equipment consisting of two HP desktop computers, one laptop, and five servers. From that home office, NPD had compiled personal records on hundreds of millions of people.

A threat actor began probing NPD's systems in late December 2023 and successfully stole the entire database. The stolen data was posted on a dark web forum with a $3.5 million asking price. When no one paid that price, another hacker leaked the bulk of the database for free in July 2024. The headline figure was 2.9 billion records, a number confirmed by the 277 gigabyte file size. Security researchers determined that massive duplication inflated the count, with multiple entries per person for every address they had ever lived at over roughly 30 years. The realistic estimate of affected living individuals fell between 170 million and 272 million, still a catastrophic number representing the majority of American adults. Security journalist Brian Krebs discovered that NPD's sister site hosted a publicly accessible file containing source code and plaintext admin usernames and passwords. All users shared the same six character default password. Data was stored unencrypted.

Jerico Pictures filed for Chapter 11 bankruptcy in October 2024, listing total assets between $25,000 and $75,000. The company shut down permanently in December 2024. At least 20 class action lawsuits were filed. Given the company's insolvency, with less than $75,000 in assets against potential liabilities in the billions, meaningful financial compensation for victims is functionally impossible. The California Privacy Protection Agency filed a $46,000 enforcement action against NPD for failing to register as a data broker. That $46,000 fine captures the absurd mismatch between the tools regulators have and the scale of the damage they are supposed to address.

Ransomware Became a Business

Ransomware has evolved from crude software that locked your files and demanded a few hundred dollars into a professionalized criminal industry with organizational charts, revenue sharing agreements, customer service operations, and business models that mirror legitimate software companies.

The tactical shift has been dramatic. Before 2019, ransomware simply encrypted your files and demanded payment for a decryption key. Then criminal groups started stealing data before encrypting it, threatening to publish the stolen information publicly if the ransom went unpaid. That is called double extortion. By 2024, 96 percent of ransomware cases involved data theft alongside encryption. Some groups added a third layer of pressure, launching attacks against victims' websites, directly contacting victims' customers, or threatening individuals whose data was stolen. The most striking recent shift is toward extortion without encryption at all. By 2025, only 50 percent of ransomware incidents involved encryption. Criminal groups realized stolen data creates permanent leverage. You recover encrypted files from a backup. You never recover data that someone else already has.

The criminal business model runs on something called Ransomware as a Service. Developers create and maintain the ransomware code, the infrastructure, the leak sites, and the negotiation portals. Affiliates purchase access and carry out the attacks. Specialists called initial access brokers focus on gaining network footholds and selling them. Revenue splits typically give 70 to 80 percent to affiliates and 20 to 30 percent to the operators. Some of these operations charge as little as $40 a month. As of 2025, more than 85 distinct operations were running simultaneously, and 124 named ransomware groups were being tracked, a 46 percent increase from the year before.

Law enforcement has scored significant wins. In February 2024, Operation Cronos, led by the UK's National Crime Agency and the FBI, seized 34 servers belonging to LockBit, the dominant ransomware group globally, responsible for 25 percent of all attacks. The operation froze more than 200 cryptocurrency accounts and

recovered over 1,000 decryption keys. LockBit's leader was publicly identified as a Russian national, though he was not apprehended. LockBit's ransom payments dropped 79 percent in the second half of 2024. The ALPHV/BlackCat group, the gang behind the Change Healthcare attack, collected the $22 million ransom and then executed an exit scam, vanishing permanently. The Cl0p gang's 2023 exploitation of the MOVEit file transfer system compromised more than 2,700 organizations and exposed data on roughly 93 million individuals.

Total ransomware payments dropped from a record $1.25 billion in 2023 to $813 million in 2024, a 35 percent decline. Only 28 percent of victims paid in 2025, the lowest rate on record. The decline comes from law enforcement actions seeding distrust in criminal networks, from improved backup practices, and from growing evidence, demonstrated most dramatically by the Change Healthcare case, that paying does not guarantee the return of your data or prevent further extortion. The economic toll remains immense. Global ransomware damages reached an estimated $57 billion in 2025. The FBI ranked ransomware as the most pervasive threat to critical infrastructure, and half of all ransomware attacks in 2025 struck critical infrastructure sectors.

For historical perspective, the May 2021 Colonial Pipeline attack remains the moment ransomware entered the national conversation. A criminal group shut down the largest refined oil pipeline in the United States, 5,500 miles carrying 45 percent of the East Coast's fuel supply, for five days. Gas stations across the Southeast ran dry. The president declared a state of emergency. Colonial paid $4.4 million. The attack triggered a cybersecurity executive order, new pipeline security directives, and the Cyber Incident Reporting for Critical Infrastructure Act of 2022.

After the Breach: Your Data Never Stops Working Against You

A data breach is not a single event. The breach is the starting gun on a chain reaction that persists for years, compounding harm through stolen credentials, identity theft, synthetic fraud, and a thriving underground economy in personal information.

When attackers steal username and password combinations from one service, they automatically test those same combinations against hundreds of other platforms. This is called credential stuffing. Because 65 percent of people reuse passwords across multiple accounts, even a tiny success rate, as low as one tenth of one percent, translates into hundreds of thousands of compromised accounts when launched at scale. Akamai recorded 26 billion credential stuffing attempts per month in 2024.

Specialized malware called infostealers harvested 1.8 billion credentials in 2025 alone, an 800 percent surge. More than 54 percent of ransomware victims in recent years had their stolen login credentials appear on criminal marketplaces before the ransomware attack hit, sometimes with less than 48 hours between the credential theft and the ransomware deployment.

The FTC recorded 1,135,270 identity theft reports in 2024, a 9.5 percent increase, with total fraud losses exceeding $12.5 billion. The most common type was credit card fraud, with nearly 89 percent involving new accounts that criminals opened using stolen personal data. Through the first three quarters of 2025, identity theft reports had already exceeded the entire 2024 total.

A growing form of fraud called synthetic identity theft combines real stolen data, typically Social Security numbers, with fabricated names and addresses to build entirely fictitious people. Lenders reported $3.3 billion in exposure to synthetic identities by the end of 2024, an all time high. Children are especially at risk because their Social Security numbers have no credit history, and the fraud often goes undetected for more than a decade. Deloitte projects synthetic identity fraud losses will reach $23 billion by 2030.

On dark web marketplaces, your Social Security number sells for $1 to $6. A complete identity package, your name, Social Security number, date of birth, and address, costs $20 to $100. Credit card numbers with security codes run $10 to $40. Complete medical records sell for up to $500, ten times the price of a credit card, because medical records contain enough information for insurance fraud, prescription fraud, and identity theft all at once. American identity packages are the cheapest on the market due to decades of oversupply from breaches. Your Social Security number never expires and it does not change. A breach from five years ago keeps generating new victims indefinitely as datasets get recombined, repackaged, and resold.

Identity theft victims spend an average of 100 to 200 hours resolving the damage, and some spend more than 400 hours. Fraudulent accounts sent to collections damage credit scores for up to seven years. In most states, insurance companies factor your credit score into your premiums for auto and homeowners coverage, which means identity theft that damages your credit directly raises your insurance costs. Corrupted background check databases cause job rejections when breached data gets flagged during employment screening. In February 2025, a breach at DISA Global Solutions, a third party employment screening company, exposed Social Security numbers and financial information for 3.3 million people, the exact data companies use to decide whether to hire you.

Companies Hoard Your Data Like Toxic Waste

Every breach in this chapter was made worse by a common pattern: companies collected and stored far more of your personal data than they needed for any legitimate purpose. The volume of data managed by companies grew tenfold between 2016 and 2021, from 1.45 petabytes to 14.6 petabytes. Every additional record stored is one more record that gets stolen when attackers get in.

Security technologist Bruce Schneier put it plainly: data is a toxic asset, and we need to treat it as we would any other source of toxicity. The Federal Trade Commission has started agreeing. The FTC now treats excessive data retention as an independent violation of the law. In 2024, an enforcement action against Blackbaud, a nonprofit software provider, marked the first time the FTC alleged that keeping personal data longer than necessary was, by itself, an unfair business practice.

Blackbaud had stored Social Security numbers, medical information, and religious affiliations from former customers for years after the business relationship ended. When the company was breached in 2020, all that unnecessary data was exposed. Similar actions against CafePress, which stored more than 180,000 unencrypted Social Security numbers indefinitely for no business reason, and InMarket Media, which retained five years of precise location data, established a clear pattern. The FTC now routinely requires companies to establish mandatory retention schedules, delete unnecessary data, and build real security programs.

California's privacy law requires that data collection, use, retention, and sharing be reasonably necessary and proportionate to the purposes for which data was collected. The European Union's GDPR makes data minimization a legally binding principle with fines reaching 20 million euros or four percent of global annual

revenue. In the United States, only California and Maryland have what privacy advocates consider meaningful data minimization rules. For the rest of the country, companies face no legal obligation to limit how much of your data they stockpile.

Breach Notification Letters: Too Little, Too Late

All 50 states now have breach notification laws. That sounds reassuring until you look at what those laws actually require and how they actually work.

Twenty states set numeric deadlines for notifying you after a breach, ranging from 30 to 60 days. Thirty one states use vague language like "without unreasonable delay," a standard flexible enough to accommodate months or even years of silence. The Change Healthcare breach happened in February 2024. Final notifications were not completed until October 2025, twenty months later. The city of Long Beach, California, took nearly 18 months to notify more than 300,000 residents. The ITRC found that 70 percent of breach notification letters in 2025 withheld how the attack happened, up from 58 percent the prior year. That means you know your data was stolen, you just do not know how, and you have no way to assess your actual risk.

California is leading the way forward. SB 446, signed by Governor Newsom on October 3, 2025, and effective January 1, 2026, requires companies to notify affected California residents within 30 calendar days and the Attorney General within 15 days. When Social Security numbers are compromised, companies must provide at least 12 months of free identity theft prevention services. The law also requires notification letters to follow a standardized format, ending the practice of burying critical information in dense legal text.

The standard remedy companies offer after a breach, 12 to 24 months of credit monitoring, is widely recognized as inadequate. Credit monitoring alerts you after suspicious activity happens. It does not prevent anything. It covers credit fraud and misses medical identity theft, tax fraud, and synthetic identity creation. It expires after a year or two, even though stolen Social Security numbers remain usable forever. An NPR investigation found that when you sign up for "free" monitoring, the terms and conditions sometimes authorize the monitoring company to share your consumer data broadly, meaning the so called remedy creates new privacy risk. Credit freezes, which have been free under federal law since 2018, provide far stronger protection. Companies rarely mention them in breach notification letters because freezes do not generate revenue for credit bureaus.

No federal breach notification law exists, despite proposals dating back to 2003. Healthcare providers must notify within 60 days under HIPAA. Financial institutions face a 30 day deadline under the Safeguards Rule. The Cyber Incident Reporting for Critical Infrastructure Act of 2022 requires reports to the Cybersecurity and Infrastructure Security Agency within 72 hours, and its final rule has been delayed to May 2026.

What You Need to Do Right Now

The legal system has shifted the burden of protection almost entirely onto you. That is a failure of policy, and it is infuriating. It is also the reality. The good news is that the most effective tools available to you are free.

Freeze your credit at all three major bureaus: Equifax, Experian, and TransUnion. Federal law has required all credit freezes to be free since 2018. A freeze blocks anyone from opening new credit in your name. It lasts

indefinitely. It has zero impact on your credit score. Online or phone freeze requests must be processed within one business day, and unfreezes within one hour. Do not stop at the three major bureaus. Identity thieves know about lesser known credit reporting agencies and use them as backdoors. Freeze your credit at Innovis, which is used for identity verification. Freeze at ChexSystems, which banks check when you open a new account. Freeze at NCTUE, which telecom and utility providers check. And freeze at LexisNexis, which insurance companies use. A credit freeze is federally regulated under the Fair Credit Reporting Act. A credit "lock," which is a separate product the credit bureaus sell, is a contractual agreement that often contains arbitration clauses and class action waivers. The freeze gives you legal protection. The lock gives the credit bureau revenue.

Check whether your data has already been exposed. The website haveibeenpwned.com, created by security researcher Troy Hunt, tracks more than 15 billion compromised accounts across 900 plus breached sites and lets you search by email address for free. The Identity Theft Resource Center offers a Breach Alert tool that monitors up to five companies and sends you alerts when they are breached. The FTC's IdentityTheft.gov generates personalized recovery plans, official Identity Theft Reports, and pre filled dispute letters.

Be skeptical of identity theft protection services. Products like LifeLock, Aura, and Identity Guard primarily offer monitoring, which means they tell you after something has gone wrong. As one expert told Consumer Reports: no one can protect you from becoming a victim of identity theft. That protection simply does not exist. These services do not remove your data from the dark web. They do not guarantee reimbursement. Free tools, including credit freezes, the annual free credit report at AnnualCreditReport.com, and the Have I Been Pwned database, often provide equal or better protection.

Stop reusing passwords. Sixty five percent of people use the same password across multiple accounts, and that single habit is the number one way attackers turn one breach into access to your bank, your email, and your medical records. Use a password manager. Enable multi factor authentication on every account that offers it. These two steps alone would have prevented the Change Healthcare breach, the most devastating healthcare cyberattack in American history.

The Legal System Is Failing You

The Supreme Court's 2021 decision in TransUnion LLC v. Ramirez created an enormous barrier for anyone trying to hold companies accountable after a data breach. In a five to four ruling, the Court held that consumers whose data was stolen cannot sue in federal court unless they prove the stolen data was actually used by a third party. If identity theft has not materialized yet, courts call the harm speculative. Once it does materialize, proving which specific breach caused the harm is nearly impossible, since most people have had their data stolen in multiple breaches.

Data breach class action filings exceeded 1,800 in 2025, a 200 percent increase since 2022. When settlements are reached, the payouts to individuals are insulting. The Equifax settlement was headlined at $700 million. In practice, only $31 million went to the $125 cash payments consumers were promised. With roughly 5.95 million claimants, most individuals received approximately $5.21 for the exposure of their Social Security number, birthdate, and address. If all 147 million affected people had filed claims, each would have received about 21 cents.

Other democracies have found better answers. The European Union's GDPR requires organizations to notify their supervisory authority within 72 hours of a breach, with fines reaching 20 million euros or four percent of global revenue. Cumulative GDPR fines have exceeded 5.88 billion euros across more than 2,245 enforcement actions. Australia overhauled its entire privacy framework within weeks after two major breaches in 2022 affected a third of its population, raising maximum penalties to 50 million Australian dollars or 30 percent of adjusted turnover. The United States remains a global outlier: 50 plus state laws with varying definitions and timelines, sector specific federal rules without a unifying framework, and a primary reliance on lawsuits that pay victims pennies.

Your Data Is Already Out There. Your Next Move Is Not.

Here is what the numbers add up to. Three consecutive record breaking years of data breaches. Victim notifications reaching 1.35 billion in a single year. Average U.S. breach costs at an all time $10.22 million. FBI tracked cybercrime losses surging 33 percent annually. Ransomware groups numbering 124 and operating like franchises. And a legal system that pays you $5.21 when a company loses your Social Security number to hackers.

The worst breaches share the same features: companies hoarding data they do not need, basic security controls that were never turned on, monopolistic concentration that creates single points of failure, and legacy systems that persist because upgrading them costs more than absorbing the occasional breach. These are not mysterious hacking techniques. These are governance failures. These are choices.

The system that created this mess will not fix itself on its own timetable. And it will not fix itself for your benefit. That is why you have to act now. Freeze your credit at every bureau, including the ones most people do not know about. Check haveibeenpwned.com and see where your data has already been exposed. Stop reusing passwords. Turn on multi factor authentication everywhere. Talk to your family. Talk to your parents. Talk to your kids.

The companies holding your data have demonstrated, year after year, that they will not protect it. The government has demonstrated, decade after decade, that it will not make them. The only person left to protect you is you.

And you are more than enough for the job.

Chapter 18: Your Data Is Being Sold to the Government

On March 18, 2026, FBI Director Kash Patel sat before the Senate Intelligence Committee and said something that should alarm every American. Senator Ron Wyden asked him a direct question: will you commit to not buying Americans' location data without a warrant? Patel's answer was equally direct. The FBI, he said, purchases commercially available information, and those purchases have produced valuable intelligence. When Wyden pressed harder, Patel responded that the FBI uses all tools available to accomplish its mission.

Read those words again. The Director of the Federal Bureau of Investigation confirmed, under oath, in front of the United States Senate, that your government is buying your location data. Your travel patterns. Your daily routines. The places you visit, the doctors you see, the houses of worship you enter, the protests you attend. All of this data is available for purchase on the open market, and the FBI is buying. No judge signed off. No probable cause was established. No warrant was obtained.

Wyden called the practice an outrageous end run around the Fourth Amendment. He warned that combining these commercial data purchases with artificial intelligence creates a surveillance machine the framers of the Constitution never imagined and never would have permitted.

This is not a story about some obscure government program buried in classified documents. This is happening in the open. Agencies across the federal government, from immigration enforcement to intelligence, from the IRS to the Secret Service, are spending tens of millions of your tax dollars each year to buy the same intimate details about your life that they would need a warrant to obtain through any other means. And right now, as you read this sentence, no federal law prohibits the practice, and no court has ruled against the practice.

Your Fourth Amendment rights have a loophole. The government found the loophole. And the government is driving a fleet of surveillance trucks straight through the loophole.

How a 1979 Phone Call Created Today's Surveillance Machine

To understand how the government got away with buying your most private data, you need to understand a legal doctrine born nearly fifty years ago from a robbery suspect's rotary phone.

In 1979, the Supreme Court decided a case called Smith v. Maryland. Police had installed a device on a robbery suspect's phone line to record the numbers he dialed. The Court ruled 5 to 3 that no warrant was needed. The reasoning was simple and, at the time, seemed logical. When you dial a phone number, you voluntarily share that number with the phone company. Because you chose to hand over that information to a third party, you assumed the risk that the government might get access to the information too. The legal world calls this the third party doctrine. If you share your information with a company, the government argues you gave up your right to privacy in that information.

Justice Thurgood Marshall saw the danger. He wrote in dissent that people have no meaningful choice in the matter. You have to use a telephone to participate in modern life. Expecting someone to give up their telephone to avoid surveillance is unreasonable. Marshall warned that the majority's reasoning would let the

government decide for itself how far Fourth Amendment protections reach. For four decades, his warning went unheeded.

Then came the smartphone era. And with the smartphone era came a case that changed everything, and simultaneously changed nothing.

Carpenter Closed One Door and Left Another Wide Open

In 2018, the Supreme Court took up the case of Timothy Carpenter, a man convicted of a string of robberies. The FBI had obtained 12,898 location data points from his cell phone carrier over 127 days. That is an average of 101 data points every single day, compiled without a warrant. The Court, in a 5 to 4 decision written by Chief Justice John Roberts, ruled that this kind of cell phone location tracking amounts to a search under the Fourth Amendment, and that the government needs a warrant to get those records from your wireless carrier.

Roberts recognized that cell phones are indispensable to participation in modern society. He acknowledged that your phone logs your location automatically, without any action on your part beyond simply turning the device on. He understood that this data creates a detailed chronicle of a person's movements, revealing family, political, professional, religious, and sexual associations. The Court got the problem exactly right.

Then the Court deliberately created the escape hatch. Roberts wrote that the decision was narrow. The majority said nothing about whether the government is free to simply buy the same type of data from a commercial data broker instead of compelling a wireless carrier to hand the data over. Roberts even mentioned in his opinion that wireless carriers sell aggregated location records to data brokers. He acknowledged the pipeline. He chose not to address the pipeline.

Federal agencies read that silence as an invitation. If the Court said the government needs a warrant to force a phone company to turn over your location data, the government's lawyers reasoned that buying the same data on the open market from a data broker is a voluntary transaction, a commercial purchase, a routine business deal. In their view, no Fourth Amendment interest is triggered because the data is available to anyone willing to pay. One senior DHS attorney compared buying your location data to buying a widget.

Think about that comparison for a moment. Your movements through every hour of every day, your visits to your doctor, your child's school, your therapist, your place of worship, your attorney's office. All of that, in the eyes of your own government, is a widget.

The Agencies Shopping for Your Life

The scale of government data purchasing is staggering, and the FBI is far from the only buyer. Virtually every major federal law enforcement and intelligence agency has purchased Americans' personal data from commercial brokers. The receipts are public, obtained through years of Freedom of Information Act litigation, congressional investigations, and leaked internal documents.

Immigration and Customs Enforcement

ICE has built the largest commercial surveillance operation of any federal agency. The relationship stretches back to at least 2017, when the Department of Homeland Security began buying location data from a company called Venntel. Venntel advertised that its systems process more than fifteen billion location signals every day from over 250 million mobile devices. ICE purchased Venntel licenses for $190,000 in 2018. Contracts with Babel Street, another data vendor whose Locate X product lets users draw a digital boundary on a map and see every mobile device that entered or exited that area, grew from $1.1 million in 2018 to nearly $3 million by 2020.

The spending has accelerated dramatically. In September 2025, ICE signed a $2.3 million no bid contract with a company called PenLink for cell phone location tracking and social media surveillance tools. PenLink had previously merged with an Israeli surveillance firm called Cobwebs Technologies, a company that Meta banned from its platforms in 2021 for targeting activists and journalists. By January 2026, ICE's total spending on PenLink tools exceeded $5 million, and the agency's 2025 surveillance budget was ten times its entire surveillance spending over the previous thirteen years combined. In February 2026, ICE signed a blanket purchase agreement with Palantir carrying a one billion dollar ceiling.

ICE agents used Venntel data to identify a suspected smuggling tunnel from Mexico to a fast food restaurant in Arizona, then staged pretextual traffic stops to conceal the role the purchased data played. Every ICE investigative analyst had access to the Venntel system. When seventy two Democratic lawmakers demanded a DHS Inspector General investigation in March 2026, ICE cancelled a scheduled congressional briefing one day before it was supposed to happen, with no explanation. If your representative in Congress asks you to support legislation restricting warrantless data purchases, your support makes a tangible difference in closing this gap.

The FBI's Broken Promise

The FBI's arc on this issue tells you everything about government accountability on surveillance. In March 2023, Director Christopher Wray told the Senate Intelligence Committee that the FBI had purchased commercial location data for a specific national security pilot project, and that the project had not been active for some time. That was the assurance. The FBI tried the data. The FBI stopped using the data.

Three years later, under Director Patel, the FBI confirmed it resumed purchasing Americans' data. Patel described those purchases as consistent with the Constitution. He offered no specifics about what data is being acquired, how many Americans are affected, or what safeguards exist. Meanwhile, the FBI's existing contracts tell their own story. In March 2022, the bureau signed a deal for 5,000 licenses of Babel X social media tracking software, a contract worth up to $27 million over five years.

From the NSA to the IRS, Everyone Is Buying

The Defense Intelligence Agency confirmed in a January 2021 memo to Senator Wyden that the agency purchases commercially available smartphone location data. DIA analysts queried American location data five times in the preceding two and a half years. In March 2026, the DIA Director confirmed those purchases continue.

The NSA's role was declassified in January 2024 after Wyden spent nearly three years pushing for disclosure. The NSA confirmed that the agency buys various types of commercially available information, including

Americans' internet browsing records, obtained without warrants. You read that correctly. The National Security Agency is buying records of what websites Americans visit.

U.S. Special Operations Command paid a secretive firm called Anomaly Six nearly $590,000 for what the contract called a Commercial Telemetry Feed. An Air National Guard unit in Iowa that flies armed Reaper drones also purchased access to commercial location tracking tools.

The Secret Service purchased access to Babel Street location data for over $600,000 in 2019 and another $229,000 in 2021. Internal Secret Service personnel raised concerns that the data contained personally identifiable information, directly contradicting the agency's public claims that the data was anonymous.

The IRS Criminal Investigation division purchased Venntel subscriptions granting 12,000 queries per year, with agency lawyers arguing no warrant was needed because phone users had voluntarily granted access to location collecting apps. The DEA committed more than $10 million to PenLink surveillance tools.

Customs and Border Protection purchased over $5 million in data broker contracts and, in one three day span in 2018, acquired approximately 113,654 location data points from commercial sources. In March 2026, an internal DHS document confirmed for the first time that CBP's location data was sourced directly from real time advertising auctions, the same auctions that serve you banner ads on your phone.

Every time you read about another agency making these purchases, remember: you are paying for this. Your tax dollars fund every one of these contracts. You are financing the surveillance of yourself. That reality alone should motivate you to contact your senators and representatives and ask them where they stand on the Fourth Amendment Is Not For Sale Act and its successor legislation.

What the ACLU Uncovered in January 2026

On January 12, 2026, the ACLU published a new batch of documents obtained from DHS through ongoing Freedom of Information Act litigation originally filed in December 2020. These documents provided the most detailed look yet at the legal and operational machinery behind warrantless data purchasing.

The most significant document was a two page internal ICE legal memo laying out the agency's rationale for buying Americans' data without a warrant. The memo argued that purchased location data is fundamentally different from the cell phone records at issue in the Carpenter case because the data is commercially available rather than compelled from a carrier. This is the clearest statement of the government's legal theory that has surfaced through litigation. Additional documents revealed that the DHS Privacy Officer raised direct concerns about how advertising identifiers get linked to specific individuals during analysis, and about how long the data is retained. Internal records also showed that the DHS Office of Science and Technology purchased Venntel access without completing a required Privacy Threshold Assessment.

The document release implicated six separate DHS components: ICE, CBP, the Secret Service, the Coast Guard, DHS headquarters offices, and the Office of Science and Technology. DHS had announced in 2024 that the department was ending data broker contracts. The ACLU's documents showed that ICE had already resumed purchases through a different vendor.

Congressional reaction was immediate. Senators Mark Warner and Tim Kaine demanded a DHS Inspector General investigation. By March 2026, seventy two lawmakers had joined that call. Maine's legislature passed a landmark state privacy bill partly inspired by the revelations. You have the ability to demand the same from your state legislators, and if you live in a state that has not taken action, your voice is needed now.

The Law Congress Passed and Never Finished

A bill called the Fourth Amendment Is Not For Sale Act was first introduced in April 2021 by Senator Ron Wyden, a Democrat from Oregon, and Senator Rand Paul, a Republican from Kentucky. Let the bipartisan nature of that pairing sink in. On the question of whether the government should need a warrant to access your personal data, members of Congress from both parties agree. The bill attracted twenty Senate co sponsors spanning the ideological spectrum, from Senator Bernie Sanders to Senator Mike Lee.

The bill's core idea is straightforward. Government agencies would need to get the same court orders to obtain data from data brokers that they already need to get data from phone companies and tech firms. The bill would also prohibit agencies from purchasing data obtained through deception or hacking, and it would make any evidence obtained in violation of the law inadmissible in court.

In July 2023, the bill passed the House Judiciary Committee unanimously. On April 17, 2024, the full House of Representatives passed the bill 219 to 199. One hundred and twenty three Republicans and ninety six Democrats voted yes. Both the Speaker of the House and the Minority Leader voted in favor.

The bill then moved to the Senate, where it died. Offered as an amendment to the FISA Section 702 reauthorization, it failed 31 to 61, unable to reach the sixty vote threshold needed to survive. Senate leadership from both parties insisted that all amendments fail to avoid sending the FISA bill back to the House before its expiration deadline. The intelligence community lobbied hard against the measure, calling the restrictions devastating to operations. Law enforcement groups lobbied against the measure as well. The standalone Senate version expired without a committee vote on January 3, 2025.

As of March 2026, the bill has not been reintroduced in the new Congress. A related bill, the Government Surveillance Reform Act, was introduced on March 13, 2026 by Senators Wyden and Mike Lee. The data broker loophole in federal law remains wide open. Your representatives need to hear from you on this. A phone call, a letter, an email to your senators asking them to co sponsor surveillance reform legislation costs you five minutes and moves the needle more than you might think.

Geofence Warrants and the Supreme Court's Next Big Decision

Geofence warrants represent a different species of government surveillance, and the Supreme Court is about to weigh in on them for the first time. Here is how a geofence warrant works. Instead of identifying a suspect and then seeking that suspect's data, law enforcement draws a virtual boundary around a location, say a crime scene or a city block, and demands data on every single device present in that area during a specific time window. Google maintained a database called Sensorvault that archived location data from approximately 592 million users, logging locations on average every two minutes, sometimes with precision down to a few meters.

Google received its first geofence warrant in 2016. Requests exploded from there, increasing 1,500 percent from 2017 to 2018, then another 500 percent from 2018 to 2019. By 2020, Google was receiving over 11,500 geofence warrants per year, accounting for roughly one quarter of all the warrants the company receives.

Jorge Molina, a twenty three year old in Avondale, Arizona, learned what these warrants mean for an ordinary person. In December 2018, police arrested him for murder. A geofence warrant had flagged a device linked to his Google account near the crime scene. Officers told him his phone placed him at the scene one hundred percent, without a doubt. He spent six days in jail. He lost his job. He lost his car. He dropped out of school. Then police discovered the actual suspect was his stepfather, who had been using an old phone still logged into Molina's Google account.

In Gainesville, Florida, a thirty year old restaurant worker named Zachary McCoy became a burglary suspect because his daily cycling route happened to pass near a burglarized home. He borrowed $7,000 from his parents to hire a lawyer who got the warrant quashed.

The federal courts are deeply divided on whether geofence warrants pass constitutional muster. The Fifth Circuit ruled in August 2024 that geofence warrants are categorically prohibited by the Fourth Amendment, calling them modern day general warrants and comparing them to the exact sort of general, exploratory rummaging the Fourth Amendment was designed to prevent. The Fourth Circuit reached the opposite conclusion, ruling that a defendant had no reasonable expectation of privacy in Google location data because he had voluntarily shared that information with Google.

The Supreme Court agreed on January 16, 2026 to hear the case, Chatrie v. United States, with oral argument set for April 2026. This will be the first time the Court examines how the Carpenter decision applies to location data held by a commercial technology company rather than a wireless carrier. The ruling, expected by summer 2026, will shape Fourth Amendment protections for a generation. Meanwhile, Google announced in December 2023 that the company would move location history data to on device storage by default and encrypt cloud backups so that even Google cannot read them.

Those changes make it far harder for the company to respond to geofence warrants going forward. No state has enacted a full ban on geofence warrants, though New York's Reverse Location Search Prohibition Act has been reintroduced in 2025, and Montana became the first state in 2025 to prevent police from purchasing citizens' data that would ordinarily require a warrant.

Keyword Warrants: When Your Google Searches Become Evidence

If geofence warrants track where you go, keyword warrants track what you think. Law enforcement submits a request to Google asking for a list of every person who searched for specific terms during a given time period, sometimes within a particular geographic area. The warrant does not target a known suspect. The warrant targets everyone who typed certain words into a search bar.

The first state supreme court to address the constitutionality of keyword warrants was Colorado. An arson fire in Denver's Green Valley Ranch neighborhood killed five members of a Senegalese immigrant family, including a toddler and an infant. With no suspects after two months, Denver police obtained a reverse keyword warrant seeking the IP address of anyone who had searched for the home's address in the fifteen days before the fire. Google produced data on 61 searches from 8 accounts, leading to the arrest of three

teenagers. The Colorado Supreme Court recognized that people have a constitutionally protected privacy interest in their search histories, with direct implications for free speech and free thought, and assumed without deciding that the warrant was constitutionally defective. The court allowed the evidence under a good faith exception.

Keyword warrants require searching databases encompassing over one billion Google users processing up to 100,000 queries per second. Google's changes to location history storage have no effect on keyword warrants. Your searches remain stored on Google's servers, and law enforcement continues to demand access to them.

The Pennsylvania Supreme Court currently has a keyword warrant case pending, with major civil liberties organizations filing briefs arguing these warrants are unconstitutional dragnets. If the idea of the government having access to your search history concerns you, this is the moment to switch to a privacy focused search engine and to review the data Google stores about you. Go to myactivity.google.com and see for yourself what is there. Delete what you do not want stored. Turn off Web and App Activity if you are willing to trade some convenience for privacy.

The Bigger Surveillance Machine

Data purchases, geofence warrants, and keyword warrants are just one part of a much larger surveillance ecosystem. Tower dumps allow law enforcement to request records of every phone that connected to a specific cell tower during a given period. A single tower dump captures data from hundreds to nearly 1,700 phones. In February 2025, a federal magistrate judge in Mississippi became the first judge to deny a tower dump request outright, writing that the Fourth Amendment does not permit law enforcement to rummage through troves of data and then decide for themselves whether probable cause exists.

A company called Fog Data Science sells location data to local police departments for as little as $7,500 per year, putting mass surveillance capabilities within the budget of small town law enforcement. Stingray devices impersonate cell towers to capture every phone in range. Automated license plate readers, like those operated by Flock Safety in Vallejo, California, identified more than 400,000 vehicles in just thirty days. Data fusion platforms from companies like Palantir integrate all of these surveillance streams into unified dashboards, giving analysts a single screen view of a person's movements, associations, and digital life.

The intelligence community's own assessment of the problem is chilling. In June 2023, the Office of the Director of National Intelligence declassified a report from a Senior Advisory Group Panel that concluded commercially available data includes information on nearly everyone that is of a type and level of sensitivity that historically could only have been obtained through search warrants, wiretaps and surveillance. The report warned that this data could identify every person who attended a protest or rally, and that in the wrong hands, the data could enable blackmail, stalking, harassment, and public shaming. Most remarkable of all, the ODNI admitted that the intelligence community itself did not even know which agencies were buying Americans' personal data or the full extent of the purchases.

Real People, Real Consequences

As mentioned earlier in the book, the Muslim Pro prayer app had been downloaded more than 98 million times. Millions of Muslims around the world used the app to check prayer times and find the direction of

Mecca. In November 2020, journalists revealed that the U.S. military was buying granular location data extracted from Muslim prayer apps, including Muslim Pro and Muslim Mingle.

The data flowed through two pipelines. One was Babel Street's Locate X tool, purchased by U.S. Special Operations Command. The other was a company called X Mode Social, which had embedded its tracking software directly inside the apps and sold the resulting data to defense contractors. X Mode's CEO confirmed the company tracked 25 million American devices every month. The Council on American Islamic Relations captured the outrage perfectly: using an app to check prayer times should not lead a Muslim to become a victim of government surveillance.

Apple and Google banned X Mode from their app stores. The company changed its name to Outlogic, was acquired by another firm, and in January 2024 became the target of the FTC's first ever ban on selling sensitive location data.

The surveillance also reaches reproductive health. A data broker called Near Intelligence sold geolocation data to a nonprofit affiliated with Wisconsin Right to Life. That organization drew digital boundaries around nearly 600 Planned Parenthood clinics across 48 states and served 14.3 million targeted ads to women who visited those clinics between 2019 and 2022. Near's Chief Privacy Officer admitted the company had no technical controls to prevent targeting of sensitive location visitors until the summer of 2022.

A separate data broker called Mobilewalla tracked approximately 17,000 Black Lives Matter protesters in June 2020, analyzing their race, ethnicity, and home towns. The FTC brought enforcement actions against both companies.

Then there is the security risk created when surveillance data itself gets stolen. On January 4, 2025, Gravy Analytics, the parent company of Venntel, the primary vendor selling location data to federal agencies, discovered a massive data breach. A hacker posted 17 terabytes of stolen data on a Russian language cybercrime forum. The leaked sample contained 30 million records from just the first days of January 2025, drawn from 3,455 Android apps including dating apps like Grindr and Tinder. The full stolen database is estimated to contain over 200 billion records. The same data the government purchased to surveil Americans ended up in the hands of foreign cybercriminals.

When you reduce the amount of data apps collect about you, by reviewing app permissions on your phone and denying location access to any app that does not genuinely need your precise location, you reduce the data available to both commercial brokers and the hackers who steal from them.

What You Need to Know Right Now

The Fourth Amendment to the United States Constitution was written specifically to prevent this kind of government overreach. The framers had lived under British general warrants, which allowed government agents to search anyone, anywhere, for anything. They crafted the Fourth Amendment to require the government to go to a judge, demonstrate probable cause, and obtain a warrant before searching your private life.

The data broker loophole makes that protection meaningless for hundreds of millions of Americans. The government needs a warrant to compel your wireless carrier to hand over your location records. That same

government walks down the street to a data broker and buys information that is equally detailed, equally personal, and equally revealing, for a fraction of the cost and with zero judicial oversight.

Congress has tried to close this gap. The Fourth Amendment Is Not For Sale Act passed the House of Representatives with bipartisan support. The bill died in the Senate. A successor bill was introduced in March 2026. Whether this gap gets closed depends on whether enough Americans demand action from their representatives. Call your senators. Call your House member. Tell them you expect the same constitutional protections to apply regardless of whether the government compels your data or purchases your data.

The Supreme Court's upcoming decision in Chatrie v. United States is poised to reshape the landscape. If the Court extends the Carpenter reasoning to location data held by commercial technology companies, federal agencies will face a much harder legal argument for their data purchases. If the Court rules narrowly, the legislative path becomes even more urgent.

Every day you carry a smartphone, you generate data that describes your life in granular detail. That data flows from your phone to app developers, from app developers to data aggregators, from data aggregators to data brokers, and from data brokers to government agencies. This pipeline operates without your meaningful consent, without judicial oversight, and without any federal law preventing the final transaction.

The government is buying what the Constitution says it needs a warrant to obtain. The question is whether you and enough of your fellow Americans will insist that your government follow its own Constitution. The answer to that question will shape the privacy rights of every generation that follows.

Chapter 19: Algorithms Are Using Your Data to Overcharge You

Two shoppers walk into the same grocery store on the same Tuesday afternoon. They pick up the same brand of cereal, the same gallon of milk, the same pack of chicken thighs. They check out at the same register. One pays $114. The other pays $124. Neither one knows the other got a different price. Neither one was told.

This is not a hypothetical. In December 2025, a team of 437 volunteer shoppers spread across four American cities placed identical orders through Instacart, all at the same stores, all at the same time. Roughly 75 percent of the products they ordered were priced differently depending on who was buying. A 20 item basket at a Seattle Safeway ranged from $114.34 to $123.93, an 8.4 percent spread that, over the course of a year, would cost one family about $1,200 more than another family shopping in the exact same place.

The reason for the price gap had nothing to do with coupons, memberships, or sales. Instacart was running millions of secret pricing experiments on real customers using an AI tool called Eversight, a platform the company acquired in 2022. The algorithm looked at your personal data, estimated how much you were willing to spend, and charged you accordingly. You never saw the other price. You never knew a different price existed.

And Instacart is not the only one doing this.

Welcome to the world of surveillance pricing, where the store watches you before you walk in, calculates what you are worth, and adjusts its prices in real time. Welcome to the world of hidden dossiers, where companies you have never heard of maintain secret files about your banking history, prescription medications, insurance claims, rental record, and retail return habits, and those files determine whether you get an apartment, a job, or a fair insurance rate. One in five Americans has errors in these files. Most people have no idea the files exist.

This chapter is going to show you exactly how these systems work, who is profiting from them, and what you need to do about it right now.

The Federal Government Confirmed Your Suspicions, Then Stopped Investigating

On July 23, 2024, the Federal Trade Commission voted 5 to 0, including both Republican commissioners, to issue orders to eight companies demanding answers about how they use your personal data to set individualized prices. The targets were not the stores themselves. They were the middlemen who sell the surveillance pricing technology to the stores: Mastercard, Revionics, Bloomreach, JPMorgan Chase, Task Software, PROS, Accenture, and McKinsey and Co. These companies advertise the ability to track your behavior and calculate the highest price you will tolerate.

Six months later, on January 17, 2025, the FTC released a staff report confirming the scope of the problem. These eight intermediaries serve at least 250 retail clients, including grocery chains, apparel stores, beauty companies, home goods retailers, convenience stores, and hardware chains. Their tools collect precise geolocation, demographics, browsing patterns, shopping history, mouse movements on webpages, products abandoned in shopping carts, and search activity. One cosmetics company was targeting promotions to consumers based on their specific skin type and skin tone. Another company identified new parents through recent purchases and showed them higher priced baby thermometers.

The report found something even more troubling. Several of these tools merge data from multiple sources into unique profiles for individual shoppers. They assess your personal price sensitivity. They predict whether you are an impulse buyer. Some of them determine whether you qualify for food assistance. These are not broad marketing tools. They are individualized extraction machines designed to find the maximum amount you will pay and charge you that amount.

Then the investigation went quiet. Andrew Ferguson took over as FTC Chair on January 20, 2025, and shut down the public comment period on surveillance pricing within days, despite an original deadline of April 17. Commissioner Alvaro Bedoya responded publicly: Chairman Ferguson shut the American people out. No full final report has been published. The study appears dead.

Congressional pressure has continued. Senator Mark Warner led a bipartisan push in December 2025 urging the FTC to publish findings and act. Senator Ruben Gallego introduced the One Fair Price Act to ban the use of personal data in price setting. As of March 2026, no federal enforcement action targeting surveillance pricing has happened.

Twenty Years of Charging You More Because They Could

Surveillance pricing did not arrive overnight. Corporations have been testing how much they are able to get away with for more than two decades.

In 2000, Amazon ran one of the first documented experiments, varying prices on 68 DVD titles over five days. Customers in online discussion forums compared notes and discovered the differences. One customer found that deleting his browser cookies dropped his price immediately. Amazon refunded an average of $3.10 to 6,896 customers. CEO Jeff Bezos called the experiment a mistake and promised Amazon would never test prices based on customer demographics.

In 2012, Orbitz was steering Mac users toward pricier hotels. The company's own chief scientist confirmed the data: Mac users spent $20 to $30 more per night and were 40 percent more likely to book four and five star hotels. The underlying logic was income based. Mac owners earned an average of $98,560 compared to $74,452 for PC owners. The algorithm was reading your device and using your income bracket against you.

That same year, the Wall Street Journal tested over 42,000 ZIP codes on the Staples website and found a pattern that should disturb every American. Areas located more than 20 miles from a competitor saw higher prices 67 percent of the time. Areas near rivals saw higher prices only 12 percent of the time. The most painful detail was this: the ZIP codes getting the discounts averaged $59,900 in household income. The ZIP codes paying full price averaged just $48,700. Lower income communities paid more.

The Princeton Review charged higher prices for SAT prep courses in ZIP codes with high percentages of Asian American residents and stopped only after the practice was exposed by journalists.

The Instacart investigation in December 2025 brought this history into the present. Consumer Reports, the Groundwork Collaborative, and More Perfect Union documented a secret AI pricing experiment affecting shoppers at Albertsons, Costco, Kroger, Safeway, Sprouts, and Target. Seventy five percent of grocery products were priced differently for different customers. Some products had five different price points at the same time. Instacart halted the practice on December 22, 2025, after the investigation became public, saying

the company missed the mark for some customers. A survey of 2,240 adults found 72 percent of Instacart users did not want the company charging different prices for any reason.

Here is what you need to understand about the difference between regular dynamic pricing and surveillance pricing. When an airline raises ticket prices on a popular route because demand is high, that is supply and demand at work. The price responds to the market. Surveillance pricing is different. The price responds to you. The algorithm looks at your browsing history, your location, your purchase patterns, your device, and estimates what you personally will pay. Then the algorithm charges you that amount. The market does not set the price. Your data does.

A 2026 survey of 2,000 Americans found 62 percent are concerned about personalized pricing. Sixty six percent said they would stop shopping at a retailer that charged them more based on personal data. Only 7 percent actively support the practice. The public does not want this. The public just does not know how widespread this already is.

The Algorithm That Raised Your Rent

Surveillance pricing is not limited to what you buy at a store. An algorithm has been setting your rent too.

On August 23, 2024, the Department of Justice filed an antitrust lawsuit against RealPage, Inc., a Texas company owned by private equity firm Thoma Bravo, alleging its AI Revenue Management software violated federal antitrust law. The theory was straightforward. Competing landlords fed their proprietary rental data, including rates, occupancy levels, and lease terms, into a shared algorithm. That algorithm generated pricing recommendations that aligned rents across competitors. Landlords accepted these recommendations 80 to 90 percent of the time. The software included auto accept features and built in discouragement for landlords who wanted to set lower prices.

On January 7, 2025, the DOJ expanded the lawsuit to include six of the nation's largest landlords, joined by state attorneys general from California, North Carolina, Colorado, Connecticut, Minnesota, Oregon, Tennessee, and Washington. The named companies were Greystar Real Estate Partners, the nation's largest landlord with roughly 950,000 units, LivCor, a Blackstone subsidiary, Camden Property Trust, Cushman and Wakefield, Willow Bridge Property Company, and Cortland Management. Together these companies operate more than 1.3 million rental units across 43 states and Washington, D.C. The DOJ alleged renters in some markets paid 5 to 7 percent more than they would have in a competitive market. One landlord told RealPage that rents rose more than 25 percent within 11 months of adopting the software.

The DOJ reached a proposed settlement with RealPage on November 24, 2025, and consumer advocates called the terms deeply inadequate. No financial penalties. No admission of wrongdoing. RealPage agreed to stop using nonpublic competitively sensitive data in daily rent recommendations, eliminate auto accept features, and submit to a court appointed monitor for three years. The settlement runs seven years.

The private litigation has produced larger results. A federal court in Tennessee granted preliminary approval of 26 settlements with 27 defendants totaling $141.8 million. Greystar contributed $50 million, plus a separate $7 million multistate settlement with nine state attorneys general. State attorneys general in New Jersey, Washington D.C., California, Maryland, Kentucky, and Arizona filed their own lawsuits. California Governor Newsom signed AB 325, the Preventing Algorithmic Price Fixing Act, on October 6, 2025, explicitly

prohibiting algorithmic collusion under state antitrust law. San Francisco, Philadelphia, Minneapolis, Seattle, Jersey City, and other cities have banned algorithmic rent pricing.

One Wrong Name on a Screen and You Are Living in Your Car

The tenant screening industry generates approximately $1 billion annually and touches nearly every renter in America. Companies like SafeRent Solutions, TransUnion Rental Screening Solutions, and RealPage's LeasingDesk pull data from credit bureaus, criminal databases, eviction courts, and public records, then run that data through algorithms that produce a score or a simple accept or reject recommendation. A human being rarely looks at the results.

The error rates in this system are staggering. A 2024 criminology study compared official state criminal records against private sector background checks for 101 individuals. Sixty percent had at least one false positive error on regulated background checks. Seventy four percent of criminal charges listed on unregulated reports did not match official state records. One participant who had only two drug convictions from 30 years earlier found more than 50 erroneous charges attributed to him, including aggravated assault, robbery, gun possession, and child abuse. The problem is that these algorithms match records by names and aliases rather than by fingerprints or verified identifiers.

The FTC found in a landmark 2013 study that one in five consumers had at least one error on their credit reports. Five percent had errors severe enough to affect loan terms. That translates to roughly 42 million Americans carrying inaccurate information in their files.

The real stories behind these numbers will make your blood boil.

Marckus Williams is a Black man in Indianapolis who rebuilt his life after incarceration and founded a grocery store in a food desert. When he applied to rent a home from Tricon Residential in November 2022, his screening report showed three prior convictions. Two had been expunged. The third was not a conviction at all. Tricon applied a blanket ban with no individualized review. Williams ended up living in his car for about a month over Christmas and New Year's. He said it kind of broke me a little bit. His class action lawsuit alleges that Tricon's policy disqualifies Black applicants at a rate 5.32 times greater than white applicants.

Carmen Arroyo's son Mikhail was severely injured in a 2015 accident, leaving him unable to speak, walk, or care for himself. When Carmen applied to move Mikhail from a nursing home into her apartment, an automated screening tool flagged a disqualifying criminal record: a dismissed shoplifting charge from 2014. The screening company refused to give the family a copy of the underlying information and gave the landlord only a bare accept or decline recommendation. Mikhail had to remain in the nursing home for approximately one additional year. The Department of Justice filed a supporting brief in the case.

Marco Antonio Fernandez, a U.S. Navy servicemember with top secret security clearance, returned from a yearlong deployment in South Korea in 2018 and applied for an apartment near Fort Meade, Maryland. The landlord rejected him. The screening algorithm had confused him with Mario Fernandez Santana, a Mexico resident on a federal drug trafficking watch list. A completely different person with a different date of birth. His lawyers noted that this inaccurate reporting will follow him for the rest of his career.

An Oregon woman found her screening report loaded with burglary, narcotics charges, and bail jumping. None of the charges were hers. The report combined criminal records from five different women who were different races and had different birthdates, including one who was an active inmate.

The CFPB identified the core problem: screening companies appear inclined to include negative information even when that information might be inaccurate. Name only matching, where the system searches databases by first and last name alone, produces frequent false hits. The error risk falls disproportionately on Hispanic, Black, and Asian Americans because of less surname diversity. Over 12 million Latinos share just 26 surnames.

Federal enforcement has produced some accountability. In October 2023, the FTC and CFPB required TransUnion to pay $15 million for inaccurate tenant screening reports. In 2018, the FTC fined RealPage $3 million for attributing false criminal records to tenants with similar names. In 2020, the FTC fined AppFolio $4.25 million for including records over seven years old. The Louis v. SafeRent Solutions case produced a $2.275 million settlement after a court found that the company's algorithm assigned disproportionately lower scores to Black and Hispanic applicants using housing vouchers.

The Secret Files Most Americans Do Not Know Exist

When most people think about their credit file, they think about Equifax, Experian, and TransUnion. Those are only the beginning. Dozens of specialty consumer reporting agencies compile detailed files about you that determine major outcomes in your life, and most Americans have never heard of them.

ChexSystems tracks your banking history and maintains a score that determines whether you are allowed to open a checking account. LexisNexis operates the CLUE database, which records every insurance claim you have ever filed and is consulted by virtually every insurer in America.

MIB Group collects medical condition data from applications for life, health, and disability insurance. Milliman IntelliScript purchases your prescription drug history from pharmacy benefit managers: every medication, every dosage, every refill, every dispensing pharmacy, and every prescribing doctor. The Retail Equation tracks your retail return patterns and flags you if the algorithm decides you return too many items. The Work Number, owned by Equifax, has current payroll data covering more than 136 million jobs and is routinely consulted by landlords, creditors, and government agencies.

The CFPB publishes an annual list of these companies. The most recent version, released January 30, 2025, expanded to include sports betting companies for the first time. The list also documented that employment screening reports now include social media data and that auto insurers collect driving behavior data through GPS and mobile phone telematics.

The errors in these files carry devastating consequences. A woman named Crawford had an excellent credit score of 788 out of 850 with no criminal history or evictions. A tenant screening company gave her a score of just 685 out of 1,000, roughly a D grade. Her apartment complex demanded an extra month's rent as a security deposit. Judy Ann Sego was listed as deceased by LexisNexis. She was alive. She disputed the notation. LexisNexis verified it and continued reporting her as dead, destroying her access to credit. The LexisNexis Accurint class action covered 200 million class members.

The Rights You Have and the Watchdog They Are Trying to Kill

The Fair Credit Reporting Act, enacted in 1970 and updated in 2003, gives you rights that most Americans do not know about. You are entitled to one free report per year from each nationwide credit bureau and each specialty consumer reporting agency. The Big Three bureaus have offered free weekly reports through AnnualCreditReport.com since the pandemic, and that benefit continues as of early 2026. Under the 2019 Equifax breach settlement, you are also entitled to up to six free Equifax reports per year through December 2026.

If any company denies you credit, housing, insurance, or employment based on a consumer report, that company must tell you and provide the name, address, and phone number of the agency that supplied the information. You have the right to dispute inaccurate information, and the reporting agency must investigate within 30 days. You have the right to sue for violations, with statutory damages of $100 to $1,000 per willful violation, plus actual damages, punitive damages, and attorney's fees. Negative information must generally be removed after seven years, with ten years for bankruptcies.

For specialty agencies, there is no central portal. You must identify each agency from the CFPB's published list at consumerfinance.gov, contact them individually, provide identification, and request your file. Key contacts include ChexSystems at 800 428 9623, LexisNexis at 866 897 8126, and MIB at 866 692 6901. When disputing errors, consumer attorneys universally recommend sending disputes by certified mail with return receipt rather than relying on online portals that limit documentation. You should also file a complaint with the CFPB at consumerfinance.gov/complaint. Credit reporting has consistently been the number one complaint category at the Bureau.

Recent litigation has produced real results. A Wells Fargo class action has a pending settlement of $56.85 million. TransUnion settled a dispute handling case for $23 million covering approximately 485,000 consumers. CoreLogic paid $5.695 million for incorrectly listing consumers as deceased. The J.B. Hunt employment background check class action settled for $5 million in August 2025.

The Supreme Court's 2021 decision in TransUnion LLC v. Ramirez narrowed the path for FCRA class actions. In a 5 to 4 ruling, the Court held that only plaintiffs who suffered a concrete injury have standing to sue in federal court. The case involved TransUnion's erroneous flagging of 8,185 consumers as potential matches to a terrorist watchlist. The jury awarded $60 million. The Court limited recovery to the 1,853 consumers whose reports were actually shared with third parties, leaving 6,332 consumers whose inaccurate files sat in a database without a federal remedy. Justice Thomas noted in dissent that state courts are not bound by the same standing requirements, opening a path for state level FCRA cases.

The agency tasked with enforcing these rights is fighting for survival. After President Trump fired CFPB Director Rohit Chopra in January 2025, Russell Vought took over as acting head. Within weeks, DOGE operatives gained access to CFPB systems. On February 8, 2025, Elon Musk posted CFPB RIP on social media, notably while his platform was launching a digital banking service the CFPB would regulate. On February 10, Vought issued a stop work order and closed CFPB headquarters. The agency fired approximately 200 probationary employees. Vought canceled roughly $100 million in contracts and requested zero dollars from the Federal Reserve.

A federal judge intervened on March 28, 2025, reinstating employees and requiring continued operations. The administration responded by sending layoff notices to approximately 1,400 of the CFPB's 1,700 workers. On July 4, 2025, the One Big Beautiful Bill Act slashed the CFPB's budget by nearly 50 percent. A second federal judge ordered funding to continue on March 13, 2026. The Supreme Court had already ruled 7 to 2 in May 2024 that the CFPB's funding mechanism is constitutional.

The damage is already showing. Since January 2025, Experian's consumer dispute relief rate collapsed from approximately 20 percent to under 1 percent. TransUnion's dropped roughly 50 percent. Over 2.7 million credit reporting complaints remain unresolved. The CFPB dropped at least four pending enforcement cases and withdrew dozens of advisory opinions, including the critical 2021 guidance on name only matching that had been protecting renters from false criminal record hits.

These Algorithms Hit Some Americans Harder Than Others

Algorithmic pricing and screening do not affect all Americans equally. When algorithms set prices using ZIP codes, credit scores, browsing patterns, and shopping history, they replicate and amplify existing economic disparities. Those disparities track closely with race.

The Staples case made the math visible. The ZIP codes getting the best prices averaged $59,900 in household income. The ZIP codes paying the most averaged just $48,700. Lower income neighborhoods, disproportionately Black and Hispanic, paid more. The Consumer Federation of America found that property insurers charge homeowners with lower credit scores $1,996 more per year. Black homeowners carry an average credit score of 612 compared to 725 for white homeowners. Scholars call this digital redlining: using data driven systems to enforce the same geographic patterns of exclusion that redlining created in the 1930s.

In tenant screening, the disparities are measurable. Black applicants are disqualified under blanket criminal history bans at 5.32 times the rate of white applicants. Black women are overrepresented in eviction filings by nearly 200 percent. The SafeRent settlement established that screening algorithms failing to account for the financial value of housing vouchers discriminate against Black and Hispanic tenants under the Fair Housing Act.

The European Union requires companies to explain algorithmic decisions that significantly affect individuals, mandates human review, and classifies AI used in insurance pricing as high risk with penalties reaching 35 million euros or 7 percent of global revenue. The United States has no federal AI law, no right to algorithmic explanation, and no mandatory transparency requirements for pricing algorithms.

Colorado passed an AI Act requiring impact assessments and consumer notification for consequential AI decisions, with an effective date pushed to June 30, 2026. Illinois became the first state to create a disparate impact standard for AI hiring tools effective January 1, 2026. Thirty seven lawsuits were filed in the first month. New York became the first state to enact a surveillance pricing disclosure law, requiring businesses to post: This price was set by an algorithm using your personal data.

In the first seven months of 2025, 51 bills across 24 states were introduced targeting algorithmic pricing. That is up from just 10 in all of 2024. States are moving. The question is whether the federal government will try to stop them. A December 2025 executive order directed the DOJ to create a task force to potentially sue states with AI laws the administration considers too aggressive. Whether an executive order has the power to

override duly enacted state legislation without congressional authorization is a constitutional question headed for the courts.

What You Do Starting Today

You do not have to wait for Congress. You do not have to wait for the FTC. You do not have to wait for anyone. Here is your action plan.

First, pull your free credit reports from all three bureaus through AnnualCreditReport.com. Do this today, not next week. Review every line. Flag every error. Then go to consumerfinance.gov and download the CFPB's published list of specialty consumer reporting agencies. Contact ChexSystems, LexisNexis, MIB, The Work Number, and The Retail Equation. Request your file from each one. You are entitled to a free copy each year.

Second, when you find an error, and statistically one in five of you will, dispute the error by certified mail with return receipt requested. Do not rely on the online dispute portals. They limit what documentation you are allowed to submit. Send a letter with copies of supporting documents. Keep a record of everything. If the agency does not respond within 30 days, or responds inadequately, you have the right to sue.

Third, file a complaint with the CFPB at consumerfinance.gov/complaint and with your state attorney general. Even with the CFPB under assault, complaints create a record. State attorneys general in California, New York, Illinois, Colorado, and many other states are actively enforcing consumer protection laws. Your complaint adds to the evidence they need.

Fourth, pay attention to what your state legislature is doing. If you live in New York, surveillance pricing disclosure is already the law. If you live in California, the DELETE Act portal at https://privacy.ca.gov/drop/ lets you request deletion of your personal data from every registered data broker in one click. If your state has not passed similar laws, call your representatives and tell them you want the same protections. The 51 bills introduced across 24 states in 2025 happened because voters demanded them.

Fifth, make the algorithm's job harder. Use a VPN or a privacy focused browser when shopping online. Clear your cookies before comparing prices. Check prices in a private or incognito window alongside your regular browser. If the prices differ, you are seeing surveillance pricing in action.

The algorithms sorting your life right now were built to extract maximum value from your data. They were not built to be fair. They were not built to be accurate. They were built to make money for the companies that deploy them. Your personal information is the raw material, and these systems convert that information into prices designed to take as much from you as the data says you will tolerate.

You deserve to know what is in your files. You deserve to pay the same price as the person standing next to you. You deserve a fair shot at the apartment, the job, and the insurance rate you have earned. These are not partisan issues. These are American issues. And every single one of you has the power to start fighting back today.

Chapter 20: You Clicked Agree and Signed Away Your Rights

A team of researchers created a fake consent form for a fictitious social media service. They buried a clause deep inside the agreement. The clause said the company would receive naming rights to the respondent's firstborn child. You read those words correctly. Click "I agree," and a company you never heard of gets to name your baby. Ninety eight percent of the people who saw the form clicked "I agree." They signed away naming rights to a child who did not exist yet for a service no one had ever used.

And here is the part nobody in the privacy world has been able to explain away. Eleven percent of those respondents told the researchers they "thoroughly read" user agreements before signing. Every single one of them missed the clause. One hundred percent of the self described careful readers handed over the right to name their future child without noticing.

This was a study, not a real company. Nobody lost anything. The researchers wanted to prove a single point, and they proved the point beyond any reasonable doubt. The entire system of online consent, every privacy policy, every cookie banner, every "I agree" button you have ever clicked, depends on the assumption you read and understood what you were agreeing to. You did not. Almost nobody does. And every company collecting your data knows this perfectly well.

Every privacy violation described in the previous eighteen chapters of this book rests on one legal fiction. Data brokers packaging your life into a dossier for pennies. AI companies training on your personal conversations. Health apps sharing your diagnoses with advertisers. Insurance companies buying your driving data. Government agencies purchasing your location history without a warrant. All of these depend on a single claim: you consented. This chapter tears the claim apart and shows you exactly how the trick works.

The Design Tricks Stealing Your Choice

There is a name for what companies do to your consent. In 2010, a London UX designer named Harry Brignull started cataloging the specific tricks websites and apps use to manipulate your decisions. He called them dark patterns. He built an entire library of these tricks, naming and shaming the companies behind each one. The term stuck, and the concept eventually made its way into law. California now defines a dark pattern as "a user interface designed or manipulated with the substantial effect of subverting or impairing user autonomy, decision making, or choice." The FTC calls them "manipulative design tricks and psychological tactics."

These are not random design mistakes. They are deliberate choices, built by teams of engineers and psychologists, tested on millions of users, refined for maximum effectiveness at getting you to do what the company wants. The FTC identified four main categories in a 2022 report called Bringing Dark Patterns to Light. The categories include designs creating false belief, designs concealing information, designs leading to unauthorized charges, and designs obscuring your privacy choices. Companies layer multiple dark patterns on top of each other to multiply the effect.

Here's What Each Looks Like In Your Everyday Life

Trick consent manipulates the screen so you believe you are making one choice when you are actually making another. You are checking out on a major retailer's website. A brightly colored button says "Get free delivery."

You click because who does not want free delivery. You have now enrolled in a $14.99 per month subscription. The price and the auto renewal terms were below the visible screen area on your phone. You would have needed to scroll down to see them. The company knew most people would never scroll.

Forced action makes you do something you do not want to do in order to get something you need. You visit a website. A wall pops up, blocking everything. You are unable to read the article, check the weather, or see the recipe until you click "Accept All." You are being forced to hand over your data as the price of admission. A study of 100 major e commerce sites found one in ten uses this full wall approach, blocking all access until you agree to tracking.

Misdirection and obstruction work together. Misdirection steers you toward the choice the company prefers through visual tricks: bigger buttons, bolder colors, prominent placement. Obstruction makes the privacy protective choice unreasonably difficult. A 2024 global sweep of over 1,000 websites and apps, conducted by 26 enforcement authorities around the world, found nearly 40 percent of websites created obstacles for people trying to protect their privacy or access privacy information. One third of the sites repeatedly asked users to reconsider their decision to delete their account. The sites were not asking because they cared. They were asking because every additional prompt convinces a percentage of people to give up.

Confirmshaming uses guilt to keep you in line. You try to decline something and the button says "No, I don't want free shipping" or "No thanks, I prefer paying full price." These phrases are engineered to trigger a moment of hesitation, and the hesitation is all the company needs. You pause. You second guess. You click the other button. The company wins.

A tip for right now: the next time you see a pop up or banner on any website, look at the buttons. If one option is big, bright, and easy to click, and the other option is small, gray, and hard to find, you are looking at a dark pattern. The company is steering you, and now you know how to spot the manipulation.

Cookie Banners Are a Performance, Not a Choice

You have seen cookie banners thousands of times. They appear at the bottom or top of nearly every website you visit. They ask you to accept cookies, manage preferences, or close the notification. They are supposed to give you a real choice about whether websites track your activity. In practice, they are the single most visible arena for dark pattern manipulation on the internet today.

The standard design has become a masterclass in asymmetry. A large, brightly colored "Accept All" button sits next to a tiny, muted "Manage Preferences" text link. The accept button demands your attention. The preferences link fades into the background. Users overwhelmingly click the prominent option, and the website records the click as informed consent. The consent is neither informed nor meaningful. The design made sure of this before you ever saw the screen.

More than 80 percent of cookie banners offering only two options, accept or manage, contain visual nudging to push you toward acceptance. Only 7 percent of cookie notices on major e commerce sites mention your option to opt out. Only 5 percent mention the ability to disable cookies entirely. The information you need to make a real choice is withheld by design.

The exhaustion is deliberate. If you decide to be the person who clicks "Manage Preferences," you typically face two, three, or more additional screens. You toggle individual categories. You scroll through lists of advertising partners. You confirm your selection. Meanwhile, accepting everything takes one click. This asymmetry is not an accident. The company designed the experience to tire you out so thoroughly you would accept the default.

In March 2025, California's Privacy Protection Agency issued its first enforcement order, a $632,500 fine against Honda for exactly this kind of design on their website. Opting in to advertising cookies required one click on "Allow All." Opting out required two separate steps. If you came back to the site after declining cookies, a new "Allow All" button had appeared, as if your previous choice never happened. The CPPA ordered a complete redesign and required Honda to consult a user experience designer to fix the problems.

A clothing retailer called Todd Snyder faced enforcement next for requiring photo IDs to process opt out requests, a practice the law explicitly prohibits. Then Tractor Supply received a $1.35 million fine, the largest CPPA penalty to date, for running a "Do Not Sell" link on their website leading to a form doing absolutely nothing. The form existed. You were able to fill the form out. Submitting changed nothing behind the scenes. The tracking technologies kept firing. The company also ignored Global Privacy Control signals entirely, treating automated privacy requests from consumers as if those requests did not exist.

Here is your takeaway: a "Do Not Sell My Information" link on a website means nothing unless the company has actually connected the link to the technical systems controlling data sharing. Many have not. The presence of the link gives you the appearance of choice without the substance.

The 7,000 Word Document Nobody Reads

The average privacy policy on a major American website runs approximately 6,938 words. At a normal reading speed, each one takes about 29 minutes. If you sat down and read the privacy policies for the twenty most popular websites in the country, you would spend more than nine hours. If you read the privacy policies for the 96 websites a typical person visits in a single month, you would need 46.6 hours. Longer than a full work week. Every single month.

Researchers at Carnegie Mellon calculated the national cost of this system. If every American actually read every privacy policy they encountered, the lost productivity would total approximately $781 billion per year. Each person would need to spend 76 working days doing nothing except reading privacy policies. The entire model of notice and consent depends on you doing something requiring more time than many Americans spend on vacation in a year. The system assumes you will do this for free, on your own time, with no professional training in legal language. The system was built knowing you would never do this.

The policies are not written for you to understand. An analysis of 75 privacy policies found 80 percent scored below 50 on a standard readability scale, where 60 to 70 means most adults are able to follow along easily. A third scored below 40, requiring university level reading skills to comprehend. The privacy notices from major technology companies averaged 27,000 words with a readability score roughly equivalent to Stephen Hawking's A Brief History of Time. Hawking was trying to explain the universe in a way people would appreciate. These companies are trying to make sure you do not understand what they are doing with your data.

Pew Research surveyed over 5,000 American adults in 2023 and found 56 percent frequently click "agree" without reading the policy. Another 22 percent skip the policy sometimes. Only 9 percent of adults always read a privacy policy before agreeing. Sixty nine percent of Americans say privacy policies are "something to get past." They are right. The policies are designed precisely for this purpose.

A study of nearly 20 public websites in the United Kingdom found only 1 in 200 visitors, half of one percent, even opened the privacy notice page. The ones who did spent an average of 48 seconds looking at the document. At normal reading speed, 48 seconds gets you through about 5 percent of the text. Five percent. And these were people who deliberately clicked on the link.

So when a company tells a court or a regulator you agreed to its data practices because you clicked "I agree," you now know the truth. The agreement was a performance. The consent was a fiction. And the company counted on exactly this.

A practical step you should take today: install a browser extension called Terms of Service Didn't Read, abbreviated ToS;DR. This free tool grades the terms and privacy policies of major websites on a simple A through E scale and flags the worst clauses in plain language. You will learn more in 30 seconds from the extension than you would in 30 minutes of reading the actual policy.

Billion Dollar Consequences

For years, companies calculated the profits from manipulative design far exceeded any possible penalty. The math changed in September 2025.

On September 25, 2025, the Federal Trade Commission secured a $2.5 billion settlement with Amazon over its Prime subscription practices. The settlement landed days into what was expected to be a month long jury trial in Seattle. The FTC voted unanimously. This was the largest settlement in FTC history for a dark patterns case, and the internal evidence emerging during the trial was staggering.

Amazon had created a cancellation process for Prime the company internally called "The Iliad Flow." The name was a reference to Homer's epic poem about the Trojan War, a conflict lasting ten years. Nobody at Amazon chose the name by accident. The Iliad Flow forced consumers through a four page, six click, fifteen option gauntlet of discount offers, benefit reminders, emotional appeals, and guilt inducing language. "Are you sure? You'll lose free shipping." Fifteen options. Six clicks. Four pages. All designed to make you give up and keep paying $14.99 a month.

The FTC's evidence at trial showed every time Amazon simplified the cancellation process, Prime cancellations went up and new sign ups went down. Amazon's response each time was to reverse the simplification and add the obstacles back in. The company tracked the revenue impact of each individual hurdle placed in front of customers trying to leave.

Internal documents painted a picture of a company understanding exactly what the operation involved. Amazon employees described the unwanted subscriptions as "an unspoken cancer." They called the enrollment process "a bit of a shady world." One executive was referred to internally as the "chief dark arts officer." When the FTC sought documents during the case, Amazon withheld 70,000 of them by claiming

attorney client privilege. A judge forced Amazon to review the claims. The company withdrew 92 percent of them. The judge sanctioned Amazon for bad faith.

The settlement requires Amazon to pay $1 billion in civil penalties, the largest ever imposed for an FTC rule violation. Another $1.5 billion goes to consumer refunds for approximately 35 million people enrolled without clear consent. Amazon must now eliminate buttons saying "No, I don't want free shipping," clearly disclose Prime's price and auto renewal terms during sign up, and make cancellation as simple as enrollment. An independent third party monitor will oversee compliance for the next decade.

Amazon was not the only company paying a massive price. Epic Games, maker of the Fortnite video game, paid $520 million to settle FTC charges in 2022. Of the total, $245 million addressed dark patterns in billing, where confusing button layouts on smartphones caused players to make accidental purchases, and the company locked accounts of people who disputed charges through their credit card companies. Epic ignored more than one million user complaints. The remaining $275 million covered children's privacy violations.

The FTC also secured $8 million from Care.com in 2025 for making membership cancellation deliberately difficult, $18.5 million from Publishers Clearing House for misleading consumers about sweepstakes entries, and $3 million from Credit Karma for fake "pre approval" notifications designed to trick consumers into applying for credit cards.

The FTC tried to formalize these protections through a Click to Cancel Rule, finalized in October 2024, requiring companies to make cancellation as easy as sign up across the board. A federal appeals court struck the rule down in July 2025 on procedural grounds. The current administration has not revived the rulemaking. Enforcement still happens case by case, settlement by settlement. The Amazon case proves the penalties are growing. The question is whether companies will change their behavior before they face their own billion dollar reckoning.

The Privacy Signal Most Americans Have Never Heard Of

There is a tool already built into certain web browsers telling every website you visit not to sell or share your personal information. You set the tool once. The tool works automatically on every site. You do not need to fill out a single form, click a single opt out link, or read a single privacy policy. The tool is called Global Privacy Control, or GPC.

GPC was developed in 2020 by a coalition of privacy researchers and organizations, including the Electronic Frontier Foundation and a former chief technologist of the FTC. The concept is straightforward. Your browser sends a signal to every website you visit. The signal says: do not sell or share my personal information. Under the laws of a growing number of states, the signal carries the same legal weight as if you had clicked the opt out button on the website itself.

As of January 2026, twelve states require businesses to honor GPC or similar universal opt out signals. California, Colorado, Connecticut, Montana, Nebraska, New Hampshire, New Jersey, Minnesota, Maryland, Delaware, Oregon, and Texas all mandate compliance. If you live in one of these states and you have GPC enabled in your browser, every website you visit is legally required to treat your signal as a binding opt out request. Many companies ignore the signal anyway, and regulators are starting to punish them for doing so.

The first major enforcement action for ignoring GPC came in 2022, when the California Attorney General fined Sephora $1.2 million. In 2025, Tractor Supply paid $1.35 million for the same kind of violation. Healthline Media settled for $1.55 million after sharing sensitive health data with advertisers and ignoring opt out requests sent through GPC. In September 2025, California's privacy agency teamed up with the attorneys general of Colorado and Connecticut for a coordinated multistate sweep targeting businesses failing to honor GPC signals. The enforcers are getting organized, and the message is clear.

The biggest structural change arrives in January 2027. Governor Newsom signed the California Opt Me Out Act in October 2025, making California the first state requiring every web browser to include a built in opt out signal. When this law takes effect, Google Chrome, Apple Safari, Microsoft Edge, and every other browser offered in California must give users an easy way to tell every website: do not sell or share my data. Because most browser companies are based in California, the practical effect will reach every American. Privacy advocates expect a massive increase in automated opt out signals once this law goes live.

Right now, you do not need to wait for 2027. Browsers like Brave, DuckDuckGo, and Firefox already offer GPC. You also have the option of installing the Privacy Badger extension from the Electronic Frontier Foundation, which sends the GPC signal from Chrome and other browsers. Turn GPC on today. Every website you visit will receive your opt out signal automatically. No forms. No clicking. No reading. One setting protects you everywhere.

The System Was Never Built for You

Here is the core problem with every privacy policy, cookie banner, and consent form you have ever encountered. The entire system puts the burden on you. You are supposed to read thousands of words of legal language, understand the implications, track which companies have your data, exercise your rights on every individual website, and repeat this process hundreds of times a month. You are expected to be your own privacy lawyer, your own data auditor, and your own enforcement agency. Nobody elected you to the position. Nobody trained you for the work. Nobody pays you for the time.

Sixty seven percent of Americans say they understand little to nothing about what companies do with their personal data. The number has been climbing for years. Seventy three percent feel they have little or no control over what happens with their information. Eighty one percent believe AI will be used in ways making them uncomfortable. People care deeply about their privacy and feel completely powerless to protect their privacy at the same time.

This is not a failure of personal responsibility. This is a failure of system design. The notice and consent model was built to protect companies, not consumers. When a company makes you click "I agree" before accessing a service, the company is not asking for your informed permission. The company is building a legal defense. The click becomes the company's evidence in court. "The consumer agreed. The consumer consented. The consumer had a choice." You had no choice. The choice was engineered out of the process before you ever saw the screen.

Some states are starting to fix the problem by attacking from the other direction. Instead of asking you to opt out of data collection, regulators are limiting what companies are allowed to collect in the first place. California's privacy law already requires data collection to be "reasonably necessary and proportionate." Maryland's new privacy law, effective in 2025, goes further and prohibits the sale of sensitive personal

information outright. These laws shift the burden back to where the burden belongs: on the companies doing the collecting.

The principle of symmetry is also gaining ground in every enforcement action and advisory coming out of state regulators. The rule is simple: rejecting data collection must be exactly as easy as accepting data collection. One click to accept means one click to reject. Same button size. Same color. Same visual prominence. Same number of screens. The Honda settlement, the Todd Snyder settlement, the Amazon Prime consent order, and enforcement actions in France, Sweden, and Belgium all enforce this principle. If a company makes the "yes" button bigger and brighter than the "no" button, the company is violating the law.

Researchers have proposed standardized privacy labels modeled on nutrition labels. The idea is simple. Instead of reading a 7,000 word document written by lawyers for lawyers, you would see a one page summary showing exactly what data the company collects, who receives the data, how long the company keeps the data, and whether the company sells the data. Your browser would read these labels automatically and enforce your preferences without you lifting a finger. This vision is not here yet. The technology exists. The regulatory will to mandate the labels does not. At least not in 2026. The direction of travel is clear, and every enforcement action and every new state law moves the country closer.

Your Agreement Was Never Real. Your Power Is.

Every chapter in this book has described a different way your privacy disappears. Data brokers sell your life for pennies. Smart devices record you inside your own home. Scammers clone your family members' voices. Companies sort you into economic categories you never see. Government agencies buy your location data without a warrant. And the thread connecting every single one of these violations is the same: someone, somewhere, points to a consent form you clicked and says you agreed.

You did not agree. You were processed. You were designed around. You were exhausted into clicking a button so you were able to get on with your day. And now you know how the entire trick works.

The fact you know is the beginning of something companies do not want. An informed consumer makes different choices. An informed consumer spots the asymmetric button, activates the GPC signal, installs the browser extension, demands the opt out, and tells their friends and family to do the same. An informed consumer stops being easy to manipulate.

Enable Global Privacy Control in your browser today. Use Brave, DuckDuckGo, or Firefox, which already support GPC, or install the Privacy Badger extension from the Electronic Frontier Foundation for Chrome. When a cookie banner appears, look for the smallest, hardest to find option on the screen. The small option is almost always the one protecting your privacy. Click the small one. When a company makes you jump through extra steps to cancel a subscription or opt out of data sharing, file a complaint with the FTC at ReportFraud.ftc.gov or with your state attorney general. These complaints become the evidence regulators need to bring the next enforcement action.

Every dark pattern depends on your ignorance. Every manipulative banner assumes you will not notice the asymmetry. Every unreadable privacy policy bets you will click "agree" and move on.

Starting today, prove them wrong. Then turn the page to the next chapter, where we put every tool you have learned in this book into a concrete action plan you are able to finish in 30 minutes.

Chapter 21: Take Back Your Data and Privacy in Thirty Minutes

In February 2024, a woman in Pittsburgh opened a letter from her health insurance company and felt the floor drop out from under her. The letter informed her that Change Healthcare, the company processing her medical claims, had suffered a data breach. Her name, her Social Security number, her diagnoses, her prescription history, her bank account information for direct deposits, all of it now sat in the hands of criminals. She called the number on the letter. She sat on hold for forty five minutes. When she finally reached a representative, the person on the other end offered her two years of free credit monitoring. Two years of watching to see if someone destroyed her financial life, in exchange for the permanent loss of her most personal medical information.

She was one of 192.7 million Americans affected by that single breach. Two out of every three people in this country. And here is the part that should make your blood boil. The hackers got in through a remote access portal that did not require multi factor authentication. A door without a deadbolt. A company trusted with the health records of nearly two hundred million people left the digital equivalent of a sticky note on the lock that said "come on in."

If you have read all of the chapters of this book, you now understand how deep this problem runs. You know that your phone tracks your location with one meter accuracy, that data brokers sell your personal dossier for pennies, that your car reports your driving habits to your insurance company, that your DNA sits in a database that could outlive the company storing it, and that the United States still does not have a single federal law protecting your privacy. You understand the system. Now you need to do something about it.

This chapter is your action plan. Every step is specific. Every tool is named. Every website address and phone number is included. You do not need a law degree. You do not need to be a technologist. You need thirty minutes and the decision to stop waiting for someone else to protect what belongs to you.

Your Phone in Thirty Minutes: The Privacy Audit That Changes Everything

Your smartphone is the single greatest surveillance device ever invented, and you carry it everywhere you go. A 2024 study of the fifty most popular Android apps found that each one requests an average of eleven dangerous permissions, including access to your camera, your microphone, your contacts, and your precise location. Research from NordVPN found that eighty seven percent of Android apps and sixty percent of iPhone apps request permissions they do not need to perform their basic function. Your weather app does not need to know your exact GPS coordinates. Your shopping app does not need access to your microphone. Your news app does not need your contact list.

The first thing to do on your iPhone is turn off app tracking. Go to Settings, then Privacy and Security, then Tracking, and switch off "Allow Apps to Request to Track." This single toggle prevents apps from following your activity across other companies' apps and websites. When Apple first introduced this feature, consent rates for tracking dropped to roughly twenty five percent. The advertising industry lost billions. Meta's stock dropped an estimated thirty percent. That tells you everything you need to know about how much your data is worth.

Next, open Safety Check. Go to Settings, then Privacy and Security, then Safety Check. Apple originally designed this tool for domestic abuse survivors, and it remains one of the most powerful privacy features on

any phone. Safety Check walks you through a five step review of who you share information with, which apps have access to your data, and how your device security is configured. If you are in a situation where you need to cut off access immediately, its Emergency Reset mode revokes all sharing, resets your Apple ID password, and disables Find My sharing instantly. A Quick Exit button returns you to the Home Screen if someone walks in while you are using it.

After Safety Check, turn on Advanced Data Protection for iCloud to enable end to end encryption. Turn off Personalized Ads under Settings, Privacy and Security, Apple Advertising. Enable Stolen Device Protection under Settings, Face ID and Passcode. Review the App Privacy Report under Settings, Privacy and Security, App Privacy Report, which shows you exactly which apps accessed your camera, microphone, and location over the past seven days. Those green and orange dots that occasionally appear in your status bar are telling you something. Green means your camera is active. Orange means your microphone is active. Pay attention to them.

On Android, the Privacy Dashboard gives you similar power. Go to Settings, then Security and Privacy, then Privacy Dashboard. You will see a pie chart showing which apps accessed sensitive permissions over the past twenty four hours, expandable to seven days on Android 15 and later. Quick Settings toggles let you disable camera and microphone access for all apps with a single tap. Android also automatically revokes permissions from apps you have not used in several months.

On both platforms, the single most consequential permission to change is precise location. When an app has precise location access, it knows where you are within roughly one meter. Approximate location narrows it to only two to ten kilometers. Navigation and ride sharing apps genuinely need precise location. Weather apps, social media, shopping apps, and news apps do not. Go through your apps right now and switch every one of them to approximate location unless you have a specific reason not to.

Your Browser Is Leaking: Fix It Today

Your browser choice matters more than most people realize. After years of promising to phase out third party tracking cookies, Google announced in April 2025 that Chrome would not even roll out a consent prompt. Cookies remain enabled by default. Every website you visit on Chrome drops tracking files that follow you across the internet, building a profile of your interests, your habits, your purchases, and your health concerns.

Safari and Firefox have blocked third party cookies by default for years. Brave blocks ads, trackers, and fingerprinting the moment you install it. If you do one thing after reading this section, switch your default browser to Firefox or Brave. In Firefox, set Enhanced Tracking Protection to Strict and enable Global Privacy Control. In Safari, enable Hide IP Address from Trackers. If you insist on staying with Chrome, manually block third party cookies in settings and know that Chrome's Manifest V3 framework limits the effectiveness of ad blockers.

Freeze Your Credit: The Five Minute Step That Stops Identity Theft Cold

A credit freeze is the single most protective action most Americans skip. It has been free by federal law since September 2018. It takes roughly five minutes per bureau. It has zero impact on your credit score. And as of right now, only an estimated ten to twenty percent of Americans have done it. Credit card fraud was the largest category of identity theft in 2024, with 449,076 reports to the FTC. Total identity theft reports hit

1.135 million that year, up nearly ten percent from the year before, with losses reaching 12.5 billion dollars. Every one of those numbers represents a real person whose life was disrupted because a criminal opened an account in their name.

A credit freeze stops that from happening. It prevents lenders from accessing your credit report, which means no one, including a criminal with your Social Security number, gets approved for new credit in your name.

You need to freeze at all three major bureaus. At Equifax, visit equifax.com/personal/credit-report-services/credit-freeze or call 888-298-0045, create a myEquifax account, and select Place a Security Freeze. At Experian, visit experian.com/help/credit-freeze or call 888-397-3742, and toggle the Frozen button after creating an account. At TransUnion, visit transunion.com/credit-freeze or call 888-909-8872 and click Add Freeze in the TransUnion Service Center.

When you need to apply for credit, ask the lender which bureau they pull from and temporarily thaw only that bureau. Federal law requires thaws to take effect within one hour online or by phone. Set a date range so the freeze automatically restores itself.

The Bureaus Most Americans Forget

Freezing only the big three leaves significant gaps that criminals regularly exploit. ChexSystems, at chexsystems.com or 800-887-7652, tracks checking and savings account history and is used by eighty percent of banks and credit unions. Without this freeze, someone can open bank accounts in your name. NCTUE, at nctue.com or 866-349-5355, tracks telecom and utility accounts. Without this freeze, someone can open phone or utility service under your identity. Innovis, at innovis.com or 800-540-2505, is the fourth credit bureau consulted when the other three are unavailable. LexisNexis, at consumer.risk.lexisnexis.com/freeze or 800-456-1244, maintains extensive public records, insurance claims, and background data. Freezing LexisNexis now also covers SageStream reports. All of these freezes are free.

One more distinction matters here. A credit freeze is federally regulated under the Fair Credit Reporting Act, always free, and protected by law. A credit lock is a private product governed by terms of service that a company writes and controls. Experian charges 24.99 dollars per month for its CreditLock product. The Consumer Financial Protection Bureau confirmed in September 2025 that credit locks are no more effective than credit freezes. The freeze is what you want. Do not pay for what the law already gives you for free.

Take Your Data Back from the Brokers Who Stole It

The data broker industry is worth roughly 294 billion dollars. At least 750 known brokers operate in the United States, and the real number is likely in the thousands. Acxiom alone claims files on 2.5 billion people with more than 3,000 data points per person. These companies know your income, your religion, your health conditions, your daily habits, and your home address. They sell that information to advertisers, insurance companies, employers, scammers, and government agencies.

You have the right to see what they have on you and to demand they delete it. For your LexisNexis consumer disclosure report, visit consumer.risk.lexisnexis.com/request or call 888-497-0011. This report contains real estate records, liens, judgments, bankruptcy records, professional licenses, historical addresses, and insurance claims history. For your insurance claims history specifically, request a CLUE report through LexisNexis,

which covers seven years of auto and home insurance claims and is used by ninety five percent of insurance companies. For a separate property loss database, request your A-PLUS report from Verisk at fcra.verisk.com or call 800-627-3487. Both are free annually under federal law.

For people search sites, manual opt outs work roughly seventy percent of the time according to a 2024 Consumer Reports study. At Spokeo, visit spokeo.com/optout, search for your listing, paste the URL, and verify via email. Removal happens within 24 to 72 hours. At Whitepages, visit whitepages.com/suppression-requests and verify via phone call. At BeenVerified, visit beenverified.com/f/optout/search and verify via email. At Radaris, visit radaris.com/control/privacy and complete CAPTCHA and email verification. The critical thing to understand is that these opt outs are generally not permanent. Brokers re scrape public records and rebuild your profile. You need to repeat the process or use an automated service.

Automated removal services handle this maintenance for you. The Consumer Reports 2024 study found Optery led with a sixty eight percent removal rate, starting at 3.99 dollars per month. EasyOptOuts, the Wirecutter budget pick, covers roughly one hundred sites for 19.99 dollars per year. Incogni by Surfshark covers more than 420 brokers for 7.99 dollars per month. DeleteMe costs 10.75 dollars per month with human reviewed quarterly reports. Each service has different strengths, and any of them provides dramatically more protection than doing nothing.

California's DROP Platform: One Click Deletion from 500 Brokers

On January 1, 2026, California launched the most aggressive privacy tool in American history. The Delete Request and Opt Out Platform, known as DROP, is the first government operated system allowing consumers to submit a single deletion request that reaches every registered data broker simultaneously. If you are a California resident, this tool was built for you. And if you are not, it shows where the rest of the country needs to go.

To use DROP, visit privacy.ca.gov/drop and authenticate through the California Identity Gateway or Login.gov. Submit your personal identifiers, including your name, date of birth, phone number, email, and optionally your mobile advertising IDs. That single request automatically goes to all 545 registered data brokers in California. Starting August 1, 2026, brokers must retrieve and process these requests every 45 days, complete their determinations within 90 days, and continue deleting any newly collected data about you indefinitely. If a broker fails to comply, the penalty is 200 dollars per consumer per day.

The California Privacy Protection Agency created a Data Broker Enforcement Strike Force in November 2025 and immediately started using it. Background Alert was ordered to shut down through 2028. ROR Partners was fined 56,600 dollars for building profiles on 262 million Americans without registering. Rickenbacher Data was fined for selling data on millions of people living with Alzheimer's and drug addiction. This agency has teeth, and it is biting.

As of early 2026, nineteen to twenty states have consumer privacy laws with data deletion rights. California, Maryland, Minnesota, and Oregon have the strongest protections. Twelve states now require businesses to honor Universal Opt Out Mechanisms like Global Privacy Control. No other state has built a centralized deletion platform like DROP. If your state representative has not heard about what California is doing, tell them.

Stop Using Text Messages for Security: Hardware Keys and Passkeys

If you still receive security codes by text message, you are using a system built on a protocol designed in 1975 that contains no security mechanisms. The SS7 signaling protocol that carries your text messages allows sophisticated attackers to intercept them at the network level. And that is the sophisticated version. The simple version is SIM swapping, where a criminal convinces your phone carrier to transfer your number to their device. The FBI documented 982 SIM swap complaints with nearly 26 million dollars in losses in 2024 alone. In March 2025, a California arbitrator ordered T-Mobile to pay 33 million dollars after a single SIM swap enabled the theft of 38 million dollars in cryptocurrency. In January 2024, a 26 year old used a fake ID at an AT&T store to SIM swap access to the Securities and Exchange Commission's official social media account. He posted a false announcement about Bitcoin that caused 230 million dollars in market liquidations.

The solution exists and it works. When Google required all 85,000 of its employees to use physical security keys in early 2017, the company experienced zero successful phishing attacks on employee accounts. Not one. The best entry level hardware key is the Yubico Security Key C NFC at 29 dollars or the Google Titan USB C with NFC at 35 dollars. Always buy two, one to use daily and one as a backup stored somewhere safe. Your total investment is 58 to 70 dollars for the strongest account protection available anywhere.

If hardware keys feel like too much, passkeys are the next best option and they are spreading fast. Over one billion people have activated at least one passkey, and 15 billion online accounts support them. Google reports 800 million accounts using passkeys. Amazon has 175 million passkey enabled customers. Each passkey is cryptographically bound to a specific website, which means a fake site will never match. You physically cannot be phished. NIST officially recognized synced passkeys in its July 2025 digital identity guidelines update and made phishing resistant authentication the new baseline.

For accounts that do not support hardware keys or passkeys, use an authenticator app. Google Authenticator is free and now offers cloud backup. Bitwarden and 1Password integrate time based codes directly into your password vault. The hierarchy from strongest to weakest is hardware security keys, then passkeys, then authenticator apps, then text message codes, then password alone. Move up the ladder as far as you reasonably can on every account that matters to you.

Privacy Tools That Actually Work: Search, Email, Messaging, and More

You do not need to become a technologist to protect your digital life. The ecosystem of privacy respecting tools has matured to the point where switching is painless and the products are genuinely good.

For search, Brave Search is the best free option, powered by its own web crawler with zero tracking. DuckDuckGo remains the easiest no fuss default. Kagi costs ten dollars per month for unlimited searches and is funded entirely by subscribers, with no advertising or tracking whatsoever.

For email, Proton Mail, based in Switzerland, offers end to end encryption with an architecture that prevents even Proton from reading your messages under legal compulsion. The free tier includes one gigabyte of storage. Proton Unlimited at 9.99 dollars per month bundles VPN, cloud storage, calendar, and a password manager. Tuta, based in Germany, encrypts more data than any competitor, including subject lines and address books, and now offers post quantum encryption by default starting at 3.60 euros per month.

For messaging, Signal is the gold standard. End to end encrypted by default for all communications, collecting virtually no metadata, and fully open source. When U.S. defense officials were found using Signal for sensitive communications in 2025, the headlines focused on the impropriety of using any messaging app for classified material. The underlying point was clear. Signal's encryption is trusted at the highest levels. Get more tips about how to setup and use Signal in "Everybody Has Something To Hide" by Guy Kawasaki and Madisun Nuismer (I was a consulting expert on the book).

WhatsApp uses the same Signal Protocol for message content, and Meta, which owns WhatsApp, collects extensive metadata including your contacts, timestamps, device information, and IP addresses, all shared across Meta's advertising ecosystem. Telegram does not encrypt regular chats or group chats at all. After CEO Pavel Durov's arrest in France in August 2024, Telegram began sharing user IP addresses and phone numbers with authorities.

For passwords, use a dedicated password manager. 1Password costs 4.99 dollars per month for families of five and includes a Travel Mode for border crossings. Bitwarden is fully open source with unlimited passwords and cross device sync on its free tier, and its premium tier costs just ten dollars per year. Avoid LastPass. The 2022 breach resulted in stolen encrypted vaults that criminals are still cracking. Five million dollars in cryptocurrency was stolen from LastPass users in December 2024 alone.

For your VPN, Mullvad costs 5.50 dollars per month, requires no email address to sign up, and accepts cash payment for maximum anonymity. Proton VPN provides the best free tier with unlimited bandwidth. A VPN protects your traffic from your internet service provider and on public Wi Fi. It does not prevent tracking through cookies, fingerprinting, or logged in accounts.

Teach Your Family: Every Generation Needs This Conversation

Privacy is a family project. Ninety five percent of teens have smartphone access, with nearly half online almost constantly. At the other end, FTC data shows fraud losses for adults over sixty reached 2.4 billion dollars in 2024, a fourfold increase since 2020, with real losses estimated between 10.1 and 81.5 billion dollars because of massive underreporting.

For elementary age children, five to ten years old, the American Academy of Pediatrics recommends shared family devices, co watching, and introducing the concept of private information, meaning name, address, and school. At this age, children generally perceive parental monitoring as a safety measure and respond well to simple rules.

For middle schoolers, eleven to thirteen, the conversation shifts to social media privacy settings, the permanence of anything posted online, and recognizing impersonation. Ask your kids directly whether they have reviewed the privacy settings on their accounts. Ask if they know how to block someone who makes them uncomfortable. Sit with them and look at the settings together.

For high schoolers, fourteen to eighteen, the conversation goes deeper into data collection, algorithmic targeting, and how a digital footprint affects college admissions and future employment.

Research strongly favors transparency over covert surveillance when it comes to monitoring your kids. A Digital Wellness Lab study found that children's reviews of parental control apps were seventy six percent one

star ratings, with the negative reviews focused on privacy invasion. A University of Wisconsin study found teens who had a voice in setting digital rules reported higher trust and better mental health outcomes. The best approach is scaffolding, meaning tighter controls for younger children and progressively more freedom for teens, with monitoring framed as temporary training wheels.

For elderly family members, start with the basics. Help them set up strong, unique passwords. AARP found that sixty four percent of Americans do not use distinct passwords across their accounts. Enable two factor authentication on their email and banking. Teach them one rule that will prevent the majority of scams. If anyone contacts them out of the blue with urgency, demanding immediate action or money, the contact is almost certainly fraudulent. AARP's Fraud Watch Network Helpline at 877-908-3360 fielded nearly 100,000 reports in 2025. AI powered voice cloning scams, where criminals replicate a loved one's voice to demand emergency money, are surging. Social media became the top contact method for scammers targeting seniors by dollar amount, with 561 million dollars in losses.

When the Breach Letter Arrives: Your Response Plan

If you have not received a data breach notification letter in the past year, you are in the minority. A 2025 poll found eighty percent of consumers received at least one breach notice, with forty percent receiving three to five. The record 3,322 data compromises in 2025 generated approximately 278.8 million victim notices. Nearly half of the people who received those notices did nothing. Forty eight percent cited breach fatigue. Thirty six percent did not trust the notice was real.

Do not ignore breach letters. Read them carefully. A legitimate breach notification should tell you what happened, what data was exposed, what the company is offering, and what steps you should take. Be aware that only thirty percent of 2025 breach notifications disclosed how the attack happened, down from nearly one hundred percent in 2020. Companies increasingly use vague language like "potentially subject to unauthorized access" when your data was actually stolen and posted on the dark web.

Your response depends on what was exposed. If your Social Security number was compromised, and it was involved in two thirds of 2025 breach reports, immediately freeze your credit at all three bureaus, place a fraud alert, create a my Social Security account at ssa.gov, file IRS Form 14039 to prevent tax fraud, and accept any free monitoring offered. If health data was breached, request your medical records and check for inaccuracies. Medical identity theft is dangerous because wrong information in your file affects your treatment. Health records sell for more than 250 dollars each on the dark web, compared to five to ten dollars for credit card numbers. If passwords were exposed, change them everywhere you reused them, enable two factor authentication, and check haveibeenpwned.com. If biometric data was compromised, the damage is permanent. You cannot change your fingerprints or your face.

File your complaints where they count. Visit IdentityTheft.gov to generate a personalized recovery plan and an official FTC Identity Theft Report, which is a legal document useful for disputing fraudulent accounts. Report fraud at ReportFraud.ftc.gov, where your report enters the Consumer Sentinel Network accessible by more than 3,000 law enforcement agencies. File with your state Attorney General. File with the Consumer Financial Protection Bureau at consumerfinance.gov for financial data issues. Individual complaints rarely yield direct payouts for you personally, and they drive enforcement priorities, fund investigations, and build class action cases.

Stay Informed: Privacy Law Is Moving Fast

Twenty state privacy laws are now on the books. New regulations take effect every few months. No federal privacy law exists. Staying current requires intentional effort, and several organizations make it straightforward.

The Electronic Frontier Foundation at eff.org provides the most accessible consumer resources, including the free Surveillance Self Defense guide and the Privacy Badger browser extension. The Electronic Privacy Information Center at epic.org produces rigorous policy analysis and graded all nineteen state privacy laws in January 2025, finding nearly half received failing grades. The ACLU's Speech, Privacy, and Technology Project drives landmark litigation on warrantless searches and facial recognition bans. Access Now operates a 24/7 Digital Security Helpline for at risk individuals. The Future of Privacy Forum publishes essential state law analyses and tracks AI legislation across all fifty states.

For ongoing news, subscribe to the IAPP Daily Dashboard, a free daily email that is the gold standard for privacy news. Follow Krebs on Security at krebsonsecurity.com for investigative cybercrime reporting. Follow The Markup at themarkup.org for data driven tech accountability journalism. For tracking state legislation, the IAPP US State Privacy Legislation Tracker is indispensable.

In 2026, the pace is accelerating. Three new state privacy laws took effect January 1 in Indiana, Kentucky, and Rhode Island. California's automated decision making technology regulations are now effective. Connecticut's age appropriate design code takes effect mid 2026. Oregon now bans the sale of data on consumers under sixteen and bans the sale of precise geolocation data. Internationally, the EU AI Act's major provisions for high risk AI systems, including AI used in hiring and financial decisions, take effect August 2, 2026, with reach that extends to U.S. companies serving European markets.

Privacy Is Not About Having Something to Hide

Every conversation about privacy eventually runs into the same objection. I have nothing to hide. As if privacy existed only to shelter secrets. As if the only people who need curtains are the ones doing something wrong behind them.

Privacy scholar Daniel Solove dismantled this argument methodically. The "nothing to hide" claim works only by narrowing privacy to concealing wrongdoing. It ignores the real harms, which include aggregation, exclusion, secondary use, distortion, and bureaucratic indifference. Solove invoked Kafka, not Orwell, to describe the true danger. The threat is not a visible tyrant watching your every move. The threat is a suffocating powerlessness created by institutions using your personal data while excluding you from any knowledge of or participation in the process.

Arguing that you do not care about privacy because you have nothing to hide is no different from saying you do not care about free speech because you have nothing to say. If you tell surveillance advocates to email you all their passwords, not a single person will take you up on that offer.

Real people who had nothing to hide have been destroyed by the misuse of their data. At least seven confirmed wrongful arrests based on facial recognition have been documented in this country, virtually all of them involving Black Americans. Robert Williams of Detroit was detained for thirty hours over a blurry

surveillance image matched to his expired driver's license photo. Porcha Woodruff, eight months pregnant, was arrested at home in front of her children for a carjacking she did not commit, based on an eight year old mugshot. Angela Lipps, a fifty year old Tennessee grandmother, was arrested at gunpoint while babysitting and jailed for six months based on a facial recognition match, despite bank records confirming she was 1,200 miles away at the time. NIST research shows that African American and Asian faces are up to one hundred times more likely to be misidentified by these systems.

The chilling effect is not theoretical. Jon Penney's landmark study found a twenty percent decline in Wikipedia page views of terrorism related articles after the Snowden revelations, and that decline was not temporary. PEN America surveyed more than 520 American writers and found one in six avoided writing or speaking on topics they believed would trigger surveillance. Self censorship in the land of the First Amendment was approaching levels found in authoritarian countries. Elizabeth Stoycheff's research demonstrated that surveillance knowledge causes people who hold minority viewpoints to silence themselves, homogenizing public discourse and creating the illusion of consensus.

Eighty one percent of Americans are concerned about how companies use their data. Seventy three percent believe they have little or no control over what happens to it. Seventy two percent want more government regulation, and that number is bipartisan. Only three percent of Americans say they understand how current privacy laws actually work. That gap between caring and acting is not hypocrisy. Privacy policies are not written to inform you. They are written to wear you down until you click agree.

Privacy is not a luxury. Article 12 of the United Nations Universal Declaration of Human Rights enshrines it as fundamental. The Supreme Court's 2018 decision in Carpenter v. United States recognized that aggregated digital data transforms formerly innocent information into categorically sensitive surveillance. Cardinal Richelieu said centuries ago that if given six lines written by the most honest man, he would find something in them to have him hanged. Bruce Schneier observed that if marijuana laws and anti homosexuality statutes had been perfectly enforced through surveillance, society would never have reached the point of reconsidering them. Privacy is the precondition for the social change that makes societies safer, freer, and more just. It protects the space where people think freely, disagree openly, make mistakes privately, and grow into who they are meant to become.

The Tools Exist. The Gap Is You.

Everything in this chapter is something you can do today. A thirty minute phone audit. A credit freeze across seven bureaus. A passkey setup. A DROP submission if you live in California. A browser switch. A password manager installation. A conversation with your kids about what they share online and a conversation with your parents about the phone call that sounds urgent and is almost certainly a scam.

The protective tools available to everyday Americans have never been more accessible. Credit freezes are free and take minutes. Passkeys have been adopted by a billion people. Hardware security keys costing 29 dollars eliminated phishing at the largest technology company on Earth. California's DROP platform enables one click deletion from more than 500 data brokers. The question is no longer whether the tools exist.

The question is whether you will use them before the next breach letter shows up in your mailbox. Before someone opens a credit card in your child's name. Before a data broker sells your location history to a company you have never heard of for a purpose you would never approve.

Privacy is not a spectator sport. Nobody is going to protect your data for you. The companies collecting it have no financial incentive to stop. Congress has failed to act for decades. The only person who will lock your front door is you.

Start now. Pick one section of this chapter, any section, and do it before you put this book down. Freeze your credit. Audit your phone. Switch your browser. Set up a passkey. Tell someone you love what you learned. Then come back tomorrow and do the next one.

Your privacy is worth fighting for. Not because you have something to hide. Because you have everything to protect.

About The Author

Mitch is an award-winning trial lawyer, private mediator, legal tech advocate, author, and keynote speaker who has spent the last four decades fighting for justice, guiding clients through high-stakes cases, and breaking down complex issues for audiences everywhere. He has taken more than 70 jury trials to verdict, winning 97 percent of them, and has earned honors like California Litigation Lawyer of the Year and Orange County Trial Lawyer of the Year.

Since founding his firm in 1986, Mitch has helped individuals and businesses navigate catastrophic injury cases, business disputes, and mediation, earning his firm a spot in the Bar Register of Preeminent Lawyers, a distinction given to fewer than five percent of law firms nationwide. He's been recognized year after year as a Southern California Super Lawyer and holds the highest "AV" peer recognition for ethics and ability. He has also served the legal community as a Judge Pro Tem, mediator for the Orange County Superior Court, and fee mediator for the Orange County Bar Association.

Mitch is a leading voice on legal/tech including AI, Web3, and the metaverse, helping professionals understand and apply these innovations to law and business. Including this book, he has authored 14 books, including "Leading with AI: Seven Steps to Transform Your Business and Empower Your People" "Artificial Intelligence in Law" "The Web3, Metaverse and AI Handbook," "Mastering the Art of Negotiation," and "The Metaverse Business Blueprint."

Mitch was also a consulting expert for Shame Nation (with foreword by Monica Lewinsky) and contributed to the California State Bar's authoritative guide "Effective Introduction of Evidence in California," Chapter 54 on electronic and social media evidence.

As a speaker, Mitch has keynoted on stages across the country and has shared his thoughts, approaches and techniques twice on stage at Tony Robbins' Business Mastery before audiences of thousands. Media outlets frequently call on Mitch for perspective on high-profile cases and emerging trends at the intersection of law and technology.

Mitch became a lawyer to help people, and that mission still drives him. His philosophy blends empathy, innovation, and integrity, whether he is trying a case, mediating a dispute, or mentoring the next generation of lawyers. Away from the courtroom, Mitch is an avid runner and paddle boarder.

Need an experienced trial lawyer, mediator, tech-savvy consultant, or engaging speaker? Mitch is always open to new conversations and collaborations—reach out. (mitch-jackson.com)

Other Books

📚 Leading with AI: Seven Steps to Transform Your Business and Empower Your People

📚 AI in Law- Revolutionizing Your Legal Practice with Innovative Strategies and Tools

📚 The Metaverse Business Blueprint: How to Build, Operate, and Grow in Spatial Commerce

📚 Mastering The Art of Negotiation- Insider Secrets for Business Owners, Entrepreneurs, and Professionals

📚 How to Create AI, Web3, and Metaverse Branding and Licensing Opportunities

📚 Power Moves- Battle-Tested Strategies From The Business Trenches

📚 The Mediator's Handbook: Turning Conflict into Collaboration

📚 Legal Tips for Creators

📚 From Courtroom to Boardroom: A Trial Lawyer's Guide to Winning Negotiations!

📚 The Web3, Metaverse, and AI Handbook

📚 From AI to Blockchain: 14 Technology Trends Every Lawyer Must Know!

📚 The Ultimate Guide to Social Media for Business Owners, Professionals, and Entrepreneurs

📚 Little Heroes- Big Tips for Bright Futures (for the little entrepreneurs in the world)

Recommendations

Let me be real with you for a second.

What you're about to read are words that mean the world to me, kind words from past clients, friends, and leaders in my community. They've been generous enough to share them over the years, and I'm deeply grateful for every single one.

Now, these aren't specifically about this book. But here's why they matter.

They give you a window into the person behind these pages. And I think that's important, because when you're investing your time in someone's ideas, you deserve to know who you're dealing with.

Everything I do comes back to the people in my life — my family, my friends, my clients, the professionals I've had the privilege of working alongside, both offline and online. Those relationships aren't background noise. They're the foundation.

So before you dive in, know this: I don't take your time, your attention, or your trust lightly. I never have. And I never will.

"Mitch is a strong legal and moral guide to those of us seeking to preserve and protect democracy and the rule of law. His insights and training are so valuable in that battle and his brilliant accessible written explanations of the law and the Constitution provide such an important platform in the work being done on Facebook and elsewhere. I look forward to his detailed posts daily and follow him closely. His opinions and observations are among my most valued."

Ty Cobb | former White House legal counsel | former Assistant U.S. Attorney for the District of Maryland

---

"It felt inevitable that Mitch and I connected, we share a deep commitment to privacy and security. His newsletter consistently delivers clear, practical insight, and his one-to-one advice has proven thoughtful, accurate, and immediately useful to me. Mitch has a rare ability to explain complex issues with clarity and precision, making difficult topics understandable and actionable. I highly recommend him for both his expertise and his exceptional communication skills. I hope his ideas are collected in a book someday so more people can benefit from Mitch's remarkable level of expertise."

Guy Kawasaki | Chief evangelist, Canva. Host, Remarkable People podcast (2.9M followers on LinkedIn)

"I first met Jon Mitchell Mitch Jackson on Spreecast.com, a video content based social networking site. I was hosting a web show that dealt with current events and politics. Mr. Jackson became a regular viewer and was kind and gracious enough to contribute his expert opinion on subjects pertaining to law and legislative affairs. There were countless times when I went to Mr. Jackson for assistance when I did not fully understand the nuances or implications of a new kind of legislation or court ruling. His advice was invaluable. Most of all, he was always incredibly approachable and willing to help. As a college student, I had previously felt

intimidated by lawyers and found myself reticent to approach those in the legal profession. Yet, since first connecting with Mr. Jackson, he has become a go to source for advice and guidance. He has also become a fantastic mentor. Mr. Jackson's welcoming demeanor and enthusiasm for helping others makes him a fantastic resource for anyone seeking legal assistance and advice. I would sincerely recommend him to any and all of my peers, and anyone seeking legal advice."

Dasha Burns | White House Bureau Chief for Politico

"Mitch Jackson is a definitive expert on constitutional law, privacy, and the rule of law. His prolific commentary is concise and to the point, making daily reading not just essential but enjoyable. His tone is authoritative yet down-to-earth, and educational without being condescending. Mitch's thoughtful approach to complex legal problems, combined with an unmatched ability to distill facts for public understanding and consumption, makes him one of the foremost thought leaders of our time."

Marisa Cianciarulo | Dean and Professor of Law, Western State College of Law at Westcliff University

---

_"Being a truth-teller comes with many risks and tends to make enemies in high places. Too many prefer the lies that make them feel good, that affirms what they want to be true. Although a mirror is only a reflection which should trigger corrective action not being smashed; often the truth teller acting as a mirror is attacked to distract from the reflection. Mitch and I share commitment to TRUTH not bound by political affiliation or any personal bias but to integrity in love of country, people and law. May history give the respect and recognition due in measuring Mitch's blatant commitment to hold the powerful accountable.

'The true measure of a man is not found in times of comfort and convenience but times of controversy and challenge.' -Dr. Martin Luther King, Jr."_

Ruth Torres | Change Agent & People Champion

"There are a lot of so-called "experts" and lawyers on LinkedIn, but Mitch stands out far above the pack because of his commitment to education in the areas of politics and privacy. Not only is he a great writer who can quickly get to the essential part of any discussion, but his writing is compelling because it is recognizable as expert and informed without being inflammatory. Every post is exceptionally well crafted, professional, educational and expands the conversation for those that care to learn. Regardless of one's politics, there is always something illuminating in his posts. LinkedIn is better because of Mitch. And I thank him for his contributions to education and truth telling."

Michael Mindlin | Managing Principal at MINDLINconsulting

_"Amidst a sea of content filled with flotsam and jetsam, the writings of Mitch Jackson are the lighthouse guiding his readers through the storm. That may read like metaphoric hyperbole, but it also happens to be true.

Mitch's articles and posts are the first writings I turn to because I know whether he's writing about the law, technology, or politics - including where the three intersect - the time I invest in reading him will inform and illuminate.

But it's his compassion and support that I find most valuable. When I engage in the public dialogue on this platform, I sometimes find myself under attack - and my attackers are almost always from the same "demographic" to which Mitch belongs. Mitch, however, has chosen to speak out at a time when it is safer, and at least in the short-term possibly more economically advantageous, to keep his opinions to himself. The courage he displays helps me maintain mine.

Though we've never met in person, I consider Mitch a comrade-in-arms, and a friend."_

Lori Block | (Mostly-retired) Mission-driven professional. Passionate traveler. Swimming addict. Volunteer focused on helping the next generations thrive. Outspoken force for good for widowed people.

"Mitch has evolved as a true leader in the legal community since our days in law school so long ago. I honor his commitment to our democracy and privacy. He is so courageous to speak out at a time when there is fear and confusion. He shares his wisdom and insights eloquently and for this I am grateful to share his writings. I'm thankful for his great work."

Mari Frank | Attorney (California) /mediator (Certified California and Florida) / Privacy Expert/Author/ Podcast Host

---

"Mitch Jackson is an invaluable source of legal information and education regarding the preservation of American democracy. Applying his formidable experience and knowledge of American jurisprudence, he has become a leading voice in preserving and defending American rule of law and the US constitution."

Virg Bernero | 51st Mayor of Lansing, Michigan

_"I met Mitch when we were in elementary school. Mitch was always bright and entrepreneurial. I was not surprised when Mitch became a successful trial attorney.

Recently, Mitch has been writing and posting commentaries on the actions of the current administration and other topics. His posts are thorough, clearly written, and factual. Mitch's writing directly addresses the threats to the rule of law and constitutional order that our country is facing, and the extraordinary corruption of the current administration.

Mitch posts his commentaries very publicly. Not many attorneys or business owners would risk alienating colleagues, clients, potential clients, and powerful interests. Mitch openly stated that the threats to our democracy are too important for hesitation. He does not c_are what anyone thinks. As a fellow attorney and citizen I applaud Mitch's courage."

Eric Morton, Esq. | Clear Sky Law Group

"I've had the privilege of following Mitch's work for some time, and it's been both inspiring and enlightening. His consistent focus on political integrity, data privacy, and transparency in public affairs demonstrates not only a deep knowledge of the issues but a genuine commitment to informing and empowering others. Mitch approaches complex topics with clarity and conviction, fostering civil dialogue in areas that too often divide us. His ability to connect the dots between policy, technology, and individual rights is exceptional, a rare blend of intellect and empathy. Mitch's insights have challenged and encouraged many of us to think more critically about the values that shape our democracy. I'm confident his upcoming work will further advance that important conversation."

Peter G. Goral | Social Media Strategist @ ArtEnvy Inc. | Marketing Solutions

---

_"Mitch Jackson is the real deal. In a world full of people who talk about privacy and digital rights, Mitch actually educates — consistently, generously, and without the fluff. He takes complicated legal and political issues and makes them matter to real people. That's hard to do, and he does it better than almost anyone I've seen.

30+ years as a trial lawyer, and he still shows up every day to teach and add value. Where does he found enough hours in the day?? If you're not already following his posts and his work, you're missing out."_

David Culbertson | Agile SEO | Board Member, Columbus Speech & Hearing

---

"Mitch is an authoritative expert on critical issues facing our world. He is a voice of reason, logic, professionalism, and practicality [in these crazy times] far more than most! I have greatly appreciated his courage and conviction, with his many writings and contributions, he cares ~ and it shows."

Graham Truax | EiR & Mentor | Startup, SME & AI Strategist | Technologist, Pragmatist, Provocateur & Adventurer

---

_"Mitch Jackson has proven himself to be a distinguished legal, political, and privacy expert whose knowledge, integrity, and commitment to educating the public about these issues truly sets him apart from other self-proclaimed subject matter experts. He prioritizes that which is in the best interests of his clients and the public. I am proud to be associated with him.

In an era where the intersection of law, politics, and personal privacy grows more complex by the day, few individuals possess both the depth of expertise and the ability to clearly communicate it to the public. As a practicing attorney and educator, he has an exceptional command of constitutional principles, regulatory frameworks, emerging privacy standards, and the evolving political landscape. More importantly, he has a rare

talent for translating complicated legal and policy issues into practical, understandable guidance that empowers everyday citizens. He performs a valuable service to our nationwide community on a daily basis by posting his carefully researched and thoughtful, fact-based ideas and opinions.

What distinguishes him from other content creators and opinion leaders is not only his educational background, legal experience, and technical proficiency, but his unwavering commitment to educating others. His thoughtful commentary on pressing issues consistently elevates the conversation. His work clearly reflects intellectual rigor, balanced judgment, and a deep respect for civil liberties and democratic principles.

Mitch is an invaluable resource and a trusted voice for the righteous. I eagerly anticipate the day when Mitch gathers his experiences, reflections, and writings into a publication to share with a wider audience."_

Gerard McAleer | Partner at MIS McAleer Integrity Services

_"I have had the privilege of connecting with Mr. Jackson and following his commentary, as well as engaging in discussions regarding the ongoing and deeply concerning climate in our country. Many of us are striving to understand current developments from legal and political perspectives, and to assess how they may affect our lives and the lives of those we care about.

In an era defined by an overwhelming volume of information—much of it inconsistent or unreliable—the value of credible, trustworthy sources cannot be overstated. It is essential to rely on individuals who are not only well-informed in their respective areas of law and government but who also communicate with integrity and clarity.

Mr. Jackson consistently provides thoughtful, accurate, and well-reasoned analysis of legal, governmental, and political issues. He has taken the time to address many of my own questions, offering insight that brings clarity to complex matters. I consider myself fortunate to know him and to benefit from his expertise."_

Marc Romano | Managing Partner, Founder, First Principle Group(FPG), Dad, Porsche 911 Enthusiast

---

"Mitch is a master connector. He's humanized his law practice with online content and through social networking. In fact he does such a great job that I've written about him in my books and discussed his ideas in my many speaking engagements around the world."

David Meerman Scott |Author of 12 books including "New Rules of Marketing & PR" and WSJ bestseller FANOCRACY | marketing & business growth speaker | advisor to emerging companies]

"Mitch is a rare breed of early adopters who can bring what's next from the edge back to the center to help everyone understand what's coming and what to do about it."

Brian Solis [Digital analyst, anthropologist, and futurist. Solis studies the effects of disruptive technology on business and society. He is an avid keynote speaker and award-winning best-selling author who is globally recognized as one of the most prominent thought leaders in digital transformation]

"I first met Mitch in Orange County at a LinkedOC event. Since then we've stayed connected on Twitter, Spreecasts and enjoyed a few podcasts together. I've watched Mitch's use of social media and he does a great job of connecting and engaging others at a very human level on the various digital platforms."

Gary Vaynerchuk [Co-founder and CEO of VaynerMedia, NY Times bestselling author and internationally acclaimed digital media marketing expert and speaker]

"Mitch has done nothing but good for the social community and is someone who is trusted and highly regarded by myself and many others in this space. He provides incredibly valuable and consistently worthwhile content to many around the world and is a true educator and trailblazer. Most importantly, he's there for you. Mitch has personally provided invaluable advice and guidance in the past, and I'm lucky enough not just to call him a great lawyer, but my friend."

Alex Pettitt [Award winning broadcaster, brand, media and biz expert, and top livestreaming personality on Periscope and other platforms]

"Do you have a Web3 or AI tech dispute? Mitch Jackson's Zoom mediation service is just what the doctor ordered. Beyond their vast expertise, what truly distinguishes Mitch and his team is their undeniable approachability and desire to help. Entrusting your dispute to Mitch is the smartest decision you could make for peace of mind."

Tom Martin- CEO of LawDroid

"Being truly human and connecting in today's tech age isn't easy, but if anyone exemplifies how best to engage people in the new digital ecosystem it is Mitch Jackson. If you have the chance to learn or work with Mitch, consider yourself lucky. The ROI of the value provided is undoubtedly going to be worth it."

Shama Hyder [Founder & CEO @ Zen Media | Keynote Speaker | Henry Crown Fellow at the Aspen Institute]

"Mitch is an amazing social networker and an all-around likable guy. I've watched his spreecasts and have been really impressed with his guests and the content. He's had so many notable people join him including Seth Godin, Leigh Steinberg and Chris Brogan. It's not at all surprising that influential people from many walks of life want to talk to Mitch because he asks great questions, he's extremely smart, and most of all, he's a super nice guy."

Jeff Fluhr, Partner at Craft Ventures; former Co-Founder and CEO of StubHub

---

"Mitch Jackson is the real deal. Rarely have I seen anyone combine high tech with high touch in such a powerful, effective and uplifting way. He's as authentic as they come and is absolutely focused on providing exceptional value to the lives of everyone he touches!"

Bob Burg [International bestselling author, speaker and coauthor of "The Go-Giver" and author of "Adversaries Into Allies: Win People Over Without Manipulation or Coercion"]

"Many leaders know how to talk. Mitch shows us how to actually share a message. His insight, knowledge, and incomparable touch make him the consummate communicator."

Sally Hogshead [Hall of Fame speaker, best-selling author, and the world's leading expert on fascination]

"Mitch's 30+ years of legal prowess, crowned with awards like 'California Litigation Lawyer of the Year,' make him an unparalleled mediator in Web3 and tech sectors. His Zoom and Metaverse venues offer a seamless and cost-effective way to resolve disputes. Mitch's services are easy to navigate, professional, and incredibly approachable. If you're a young entrepreneur hesitant about mediation, consider Mitch an invaluable resource for quick, fair, and convenient resolution."

Robert Hanna- KC Partners Founder & CEO; Legally Speaking Podcast Host

---

_"Mitch Jackson is an amazing lawyer and writer. He consistently writes timely, pertinent, and well researched articles on the current laws and regulations that pertain to our current political administration.

He writes clearly, concisely, and in terms that people not in the field of law can understand. I've learned so much from his LinkedIn and his Substack articles, he always manages to allay my fears when it comes to everything happening in our current political climate. Mitch has very high moral standards, speaks the truth, gives great advice, educates people, and somehow, beyond my understanding, manages to have a sense of humor, remains calm, all with a positive attitude. Mitch always gives me a sense of hope and inner peace. Love him!!!"_

Judith Marie C. | M.S. ZOOLOGY

"Mitch is the one that you want to have in your corner when it comes to navigating complex legal matters. With a passion for justice and a friendly and personable approach, he and his team will do everything to help you resolve the matter amicably and favorably!"

Francesca Witzburg- Founder and Managing Partner at ESCA.legal

---

"Mitch Jackson has more than 30 years' experience in civil disputes. The expertise he's acquired is perfectly applied to disagreements and disputes in the web3, DAO, and cryptocurrency spaces. Regardless of the industry, arguments and disagreements remain the same. Mitch is talented in managing conflicts, remaining neutral, and getting to the heart of the dispute so that it may be solved and the parties can move onto something more productive."

Nick Rishwain, Legal Technologist and Voice; Expert and Co-founder, CougarDAO, LLC

"Mitch Jackson's Zoom mediation services are excellent and I highly recommend them given his breadth of experience and expertise. Not only is Mitch incredibly friendly and approachable, but his entire team goes above and beyond to create an environment that is supportive for all parties involved. Their commitment to fostering understanding and achieving amicable resolutions is remarkable."

Colin Levy- Lawyer and Legal Technologist; Author of "The Legal Tech Ecosystem: Innovation, Advancement & the Future of Law Practice"

"I can't think of a more experienced lawyer living at the cutting edge of technology. While Mitch's expertise is impressive, it's who he is as a person that makes him truly remarkable. He is exceptionally generous, approachable, and personable. I am genuinely grateful for the insight and value he contributes."

Gyi Tsakalakis, Esq. [Co-Founder of AttorneySync]

"Mitch Jackson's decades of experience as a civil litigator sets him apart from other mediators in this space. Leveraging Zoom and the Metaverse saves time, cuts costs and reduces the inevitable stress of being in conflict. Mitch and his team are incredibly approachable, knowledgeable and helpful. Having Mitch and his team on your team will prove invaluable."

Bradley A. Friedman, JD

"After 30 years in FinTech, I can attest that no one is better suited to handle private mediation of the legal issues around Web3, AI and the Metaverse than Mitch Jackson. As a 'California Litigation Lawyer of the Year,' who runs a blog about the legal aspects of technology, Mitch is the premier resource for entrepreneurs seeking mediation involving blockchain-related tech. A growing number of companies are wrestling with legal issues in the technology space. This is especially true with AI. The value of Mitch's Zoom mediation service is that he is approachable, insightful and effective. Because the best mediators can transform conflict into collaboration."

Marc Angelos- Founder, Anvictus Communication

"Mitch Jackson leverages online channels like video and social networks to reach out and connect with people both on a personal level as well as a professional level. His efforts have taken him from being successful in his offline world to finding a whole new level of influence online, as well. In a very short time, Mitch has been able to reach out and connect with a lot of successful online influencers, and has been able to translate this into mutual value. Beyond all this, he's a great guy and doing yeoman work. I recommend him without hesitation."

Chris Brogan [CEO Owner Media Group; New York Times bestselling author of 9 books and listed by Forbes as one of the Must Follow Marketing Minds of 2014 while also recognizing Chris' website as one of the 100 best websites for entrepreneurs]

"Mitch Jackson is the anomaly. His approach is open and empathetic, yet determined at every turn to bring a conclusion to the case. I always felt educated about the status and that decisions were being made together. Having a guide like Mitch through the legal system isn't just necessary, it's critical."

Bryan Kramer [TED Talk & Keynote Speaker, CEO PureMatter]

"Mitch Jackson is most definitely a giver as he is extremely generous with his sound and insightful advice regarding all matters human interaction. To me, it is no surprise that Mitch is having a significant impact on people way beyond his courtrooms as he aptly translates the life lessons learned in such a high-pressure communications context to valuable communications tips to people from all walks of life including my grateful students. Mitch's interest in people is sincere and he is an extremely empathetic listener which allows him to find the perfect blend of professional and human elements of communication whether it be on or offline."

Niklas Myhr [The Social Media Professor | Chapman University]

"Mitch is a lawyer of tomorrow, today. He's the kind of lawyer and businessman who can make rain shine. Totally client focused with an aptitude to make you feel like the most special and important person in the world. Mitch reaches out and touches you where it matters most – in your mind and heart. He builds a relationship with you fast, to last; seemingly effortlessly – it's his human nature and star quality. He's a rainmaker lawyer (of the truly naked kind), meaning he's not afraid to be transparent, 'say it as it is' and do the extraordinary in order to get things done in a top quality fashion… and all for your benefit. I feel blessed that our paths crossed and entwined. You will too. There's a reason he's Top Gun. Enough said."

Chrissie Lightfoot- Author of "The Entrepreneur Lawyer"- Legal Futurist; International Speaker; Personal Brand and Digital Media Strategist

---

"Mitch Jackson is hands down one of my favorite people in the industry. I first met Mitch in 2014 when he attended my event Social Media Day San Diego. I knew Mitch from afar and had always respected his approach and expertise in the digital space. Since then, we've developed a solid friendship, and I consider him one of the most respected thought-leaders in our industry. So much so, Mitch is one of the people I turn to speak at my events on digital marketing, strategy, and of course, anything legal with digital marketing."

Tyler Anderson [Founder and chief strategy officer of Casual Fridays, a leading digital & social media marketing agency trusted by some of the biggest brands in the hospitality, tourism, non-profit, education, and entertainment industries]

"Mitch Jackson is exactly who you want having your back. While elevating my platform and speaking career, legal issues and needs naturally happen. When a potential issue was unfolding, Mitch shared specific ideas and actions with me on how to handle it. We got everything completely resolved in a matter of a few days, and he helped me alleviate a lot of stress over the holidays. I am extremely grateful for his above-and-beyond mindset and invaluable insight. Mitch is the best in the business, and I would HIGHLY recommend working with him!"

Brandon Farbstein [Gen Z Empowerment Speaker and Influencer]

"Mitch Jackson doesn't just advise, he connects. As a tech leader in VR and live video I've shifted any of my legal needs to his firm because he understands the landscape and the language. That knowledge of what's going on in tech saves me and my team valuable time… and comes with the added benefit of having an external friend and consultant. Mitch Jackson is seen by many on my team as, quite simply, another member."

Ryan A Bell [Media at NASA JPL; Emmy winner]

"Mitch Jackson is a social leader and professional that I use as a benchmark for executives whom I coach on personal branding and how to engage and build relationships on Social Media. Although we never talked about his law practice, I would recommend him to friends, family and business partners because of his authenticity, leadership and all around passion for connecting people and social good!"

Brian Fanzo [Keynote Speaker | Leading Digital/web3 and Social Business Change]

"Mitch Jackson is one of the most unique Human Beings I know and the fact that he is a Lawyer makes him even more amazing. He lives every day to help the people around him become better, smarter, faster by inspiring and educating people about how they can grow. It's an honor to know him and I am proud to call him my friend."

Jon Ferrara [American entrepreneur and the founder of Nimble. He is also best known as the co-founder of GoldMine Software Corp, one of the original contact management software companies]

"Mitch Jackson brings a rare combination of intelligence, clarity of communication, and strategy when it comes to helping people leverage technology and social media to further their business goals. I highly recommend that you pay attention to what he has to share."

Chris Lema- CEO, MotivationsAI

"Mitch is living proof that you can be professional and personable at the same time in business. He is one of the best communicators I know and proves this in the way he teaches others how to be effective in communicating. Whether it be speaking, writing or using video, Mitch demonstrates what he teaches."

Tim McDonald [Previous Director of Community at The Huffington Post; Community Engagement Strategist]

"Mitch and his team are expert communicators who understand the fast-moving targets of digital and social and weave in the very much needed human and relationship aspect of business. A lot of people can talk theory or great ideas, Mitch actually executes, usually with amazing results. It's my pleasure to write a few words of recommendation."

Bryan Elliott [Executive producer, writer and host of The GoodBrain Digital Studios, a full-service production company focused on great storytelling]

"I met Mitch through Chris Brogan, a business expert and friend I trust. When it comes to reporting on the state of our country, legally and morally, Mitch is an expert. His consistent effort in speaking truth to power (in the form of writing thoughtful commentary) is off-the-charts. We've personally discussed many of the most important issues facing our country, and like me, Mitch is focused on keeping America a democracy. Mitch is one of the good guys. Plus, I hear he's an excellent lawyer!"

Joel Libava | The Go-To Authority on Franchise Buying | Strategic Advisor, Author, and Industry Expert

---

_"I have had the privilege of being connected with Mitch here on LinkedIn and consistently engaging with his work around political accountability, privacy rights, and the protection of democratic values. His posts and articles are not only timely, but thoughtful, well-researched, and grounded in real legal and policy insight.

Mitch has a rare ability to break down complex political and privacy issues into clear, accessible discussions that educate rather than inflame. In an online environment that often rewards noise over nuance, he brings clarity, integrity, and substance. He doesn't just comment on current events, he provides meaningful analysis and practical context that adds real value to the conversation.

What stands out most is his commitment to principled dialogue and civic responsibility. His work consistently reflects a deep dedication to protecting privacy, strengthening democratic institutions, and encouraging informed engagement.

I highly recommend Mitch to anyone seeking insight, leadership, and thoughtful perspective at the intersection of law, politics, and privacy. His voice is both necessary and impactful. I am thankful for Mitch and at this moment in time, this country needs more people like Mitch!"_

Mauricio Pardo | Bilingual Sr Lead IM Engineer | US and International

---

More Client | Lawyer | Mediation | Technology recommendations at
https://mitch-jackson.com/recommendations/

Resources

Uncensored Objection Substack Newsletter (8,100 subs)

A newsletter and community for people who refuse to look away while truth is distorted, the law is twisted, and democratic norms are quietly eroded. This is a place for readers who believe the Constitution still matters, facts still count, and silence is not an option. Together, we cut through political gaslighting, reject performative outrage, and focus on clear thinking, legal reality, and accountability.

AI Legal/Biz Tech LinkedIn Newsletter (8,400 subs)

Look, if you're not staying ahead of AI, law, business, and tech right now, you're already falling behind — and I'm not going to let that happen. Join my newsletter and get the cutting-edge updates, expert tips, and exclusive interviews you need to thrive.

Smartphone Lock Down LinkedIn Newsletter (3,700 subs)

Your phone is the most powerful surveillance device ever built, and you paid for it yourself. Smartphone Lock Down is where I share 40 years of legal/tech experience and tell you exactly what your device is exposing, who is watching, and what you can do to protect yourself right now.

Stay connected online (platforms and socials)

An easy to follow and access list of my platforms and spaces.