AN INSIDER TESTIMONY

NO ETHICS IN BIG TECH

An Insider's Case Against Silicon Valley's **War on Humanity** — and the Fight to Make Them Pay

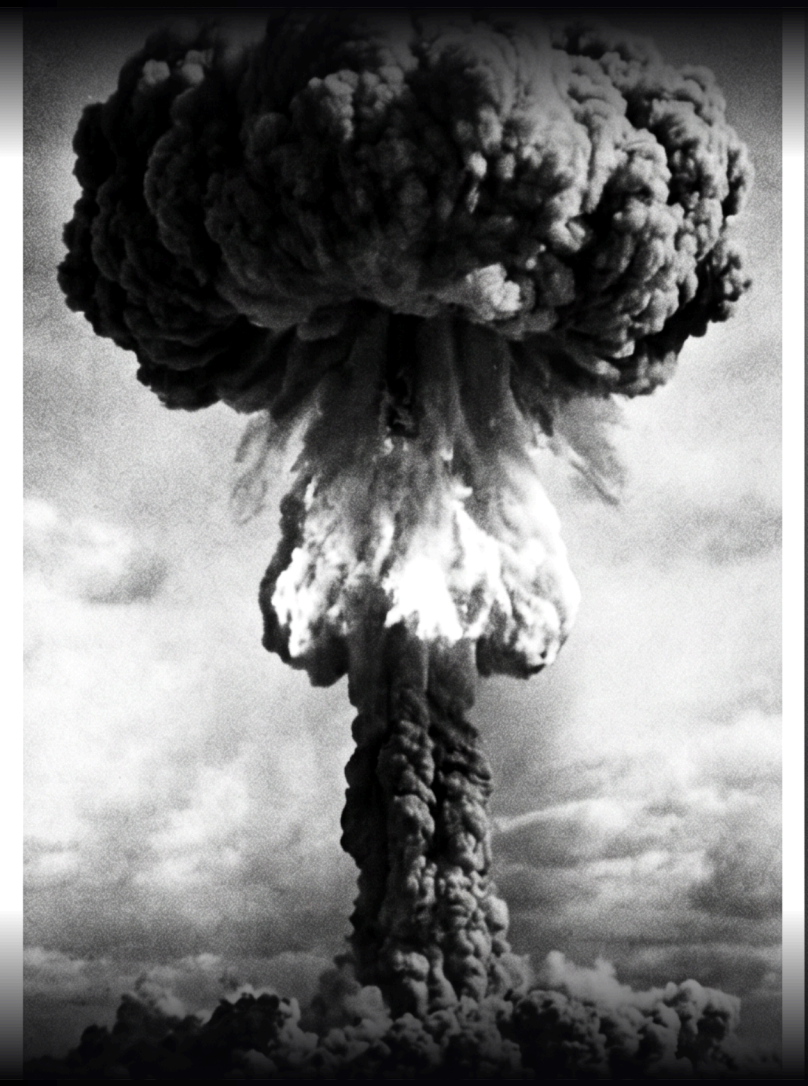

VAHID RAZAVI

Introduction by **Norman Solomon**

EDITOR Abril Menanteau · Romina Díaz

NoEthicsInBigTech.com | **ForeverPeaceNow.com** | **ParentsPlea.com** | **ConflictTour.com**

Based in part on the documentary film *Forever Peace Now*.

Editorial research and structural editing assisted by Claude (Anthropic). The views, arguments, testimony, and conclusions are entirely those of the author.

TABLE OF CONTENTS

Introduction By Norman Solomon

You don't need to know a lot about digital technology or the world of multibillion-dollar finance to understand what is most important about this book. *No Ethics in Big Tech* is a profound moral statement, acutely relevant to everyone on planet Earth.

Vahid Razavi had dug deep below the surface allure of vast wealth and power to explore the underlying terrain of crucial choices. He takes us on a grim reality tour of Silicon Valley and its entanglement with inflicting faraway suffering. Companies gain trillions of dollars in value from direct collusion with aggressive wars.

The lack of accountability and the need for it are central themes of *No Ethics in Big Tech*. This book explains why notions of reform from within the tech industry and hopes for self-correction are delusional. As a former insider, Razavi draws on painful knowledge that the obsessions with market share and maximizing profits function as the opposite of any moral compass. The warfare state is completely integrated with Big Tech.

"The military industrial complex of the 21st century is not Lockheed Martin and Raytheon and Boeing alone," he writes. "It is Amazon, Google, Microsoft, Oracle, and Palantir. The weapons are cloud computing, artificial intelligence, targeting algorithms, biometric databases, and surveillance infrastructure. The ammunition is data."

The bombs and bullets that have enforced genocide in Gaza, largely courtesy of U.S. taxpayers, are just one factor in the ongoing U.S.-Israeli alliance. "What the world calls genocide, the military industrial complex calls a revenue cycle," Razavi writes. "What Raytheon calls backlog growth, I call the rubble of hospitals in northern Gaza. What Boeing Defense calls expanded market share, I call the ruins of universities that will not reopen."

Denial of genocide does not change its reality. And the facts are horribly clear that many of the biggest and most widely esteemed Big Tech corporations in the USA have been integral to inflicting genocide on Palestinian people. Human Rights Watch and Amnesty International as well as the Israeli human-rights groups B'Tselem and Physicians for Human Rights-Israel all unequivocally concluded that Israel has committed genocide in Gaza.

Ironically, yet predictably, the digital tech industry that depends on verifiable mathematics simply refuses to acknowledge the basic facts of crimes against humanity that it avidly — and so profitably beyond imagination – aids and abets. As Razavi points out: "This is not hyperbole. It is arithmetic. The bombs fall on civilian buildings. They are manufactured by American companies, purchased with American taxpayer money, delivered by American-supplied aircraft, guided by AI targeting systems built by American technology companies. The entire supply chain, from the server farm to the crater, runs through American corporations and American government contracts. And the executives who run those corporations are compensated — in salary, in bonuses, in stock options — based in part on how much of that supply chain they own."

And so, the author is being matter-of-fact when he writes, "Without Amazon Web Services, the Israeli military's intelligence infrastructure does not function at its current scale. Without

Palantir's data fusion platforms, the Gospel and Lavender targeting systems do not have the operational picture they need to generate a hundred bombing targets a day. Without Microsoft's cloud infrastructure, the surveillance data cannot be processed at the required scale."

The corporate complicity outlined in this book is a constant reality that grows more serious on a daily basis. The vastly destructive hunger for evermore profits is insatiable. It is a feature not a bug.

"Corporate complicity in state violence is not an accident of capitalism," Razavi observes, a statement based not on theories but on his firsthand experiences of seeing up close such complicity. "It is one of capitalism's most reliable features. The only thing that has ever interrupted it is accountability. Legal, personal, financial, criminal accountability."

This book calls for accountability.

Norman Solomon June 2026

Norman Solomon is an American journalist, media critic, progressive activist, and former U.S. Congress candidate. Solomon is a longtime associate of the media watch group Fairness & Accuracy In Reporting (FAIR). In 1997 he founded the Institute for Public Accuracy, which works to provide alternative sources for journalists, and serves as its executive director. He is a co-founder of RootsAction.org

AUTHOR'S NOTE ON MISSION AND FINANCIAL INDEPENDENCE

Before you read a single word of this book's argument, I want to address something directly: who is funding this work, and what do I personally gain from it? The answer is: nothing. I gain nothing financially from any of it. This book, my previous book Ethics in Tech and Lack Thereof, and my documentary film Forever Peace Now have all been produced out of my own personal funds. No corporate sponsors. No government grants. No nonprofit foundation backing. No publisher advance. No advertising model. No speaking fee waiting at the end of the road. Every dollar spent on research, production, web hosting, legal filings, and platform maintenance has come from a single source: my own pocket. I have done this work at a personal financial loss, deliberately, over a period of years. I say this not to invite admiration but to draw a clear line. The organizations and platforms built alongside this book share the same model. ConflictTour is currently at the idea and planning stage. It is not incorporated, has no revenue, has no investors, and has no commercial model in operation. If it one day generates sufficient income to pay the people doing the work a living wage, that will be its measure of success — not profit, but sustainability. It will never be a cash cow. ParentsPlea.com is a free website. It carries no advertising. It takes no affiliate revenue. It does not sell, trade, or monetize the data of the families who trust it with the stories of their dead. It is built and maintained for a reason that has nothing to do with money: because just peace requires us to inspect the deaths that have taken place. Who is doing the killing. Who is suffering. And because if there is ever going to be judgment in a court of law, if there is ever going to be reparations for people harmed by these wars, if there is ever going to be a memorial worthy of the name — we need more than a wall of names. We need biographies. We need family testimony. We need the complete record of who was lost and what the surviving family wants the world to know. That is what ParentsPlea is building. One profile at a time. My work does not receive corporate sponsors, as our politicians and many NGOs and government-affiliated organizations often do. There is no billionaire patron. There is no foundation underwriting. That independence is not a limitation. It is the point. The same men I am calling to account in these pages are the men who would be asked to fund this kind of work — which is precisely why this work cannot be funded by them, and is not. I am not here to defend the indefensible technology sector. They have enough PR agencies and executives paid handsomely to defend their inhumanity. They can write and publish their own books. I have no interest in their approval, their support, or their response. The real struggle humanity is facing — beneath all the politics and the weapons and the corporate structures and the algorithms — is between two competing ideas: I am one. And we are one. Silicon Valley has built a monument to the first. This book is an argument for the second. That is why I am writing it. That is the only reason.

FOREWORD: THE PERSON WHO WROTE THIS BOOK

There is a specific kind of knowledge that comes from having survived a war as a child — and a different kind that comes from spending thirty years inside the industry now building the machines that fight them. I have both. And I want you to understand how they connect, because this book could not have been written without either. I was born in Tehran, Iran in 1975. Five years after my birth, Iran had its Islamic Revolution. In September 1980, Iraq attacked Iran under Saddam Hussein, and I witnessed a war firsthand. Not on a screen. Not as a statistic. I was five years old. My mother was a nurse. She worked in hospitals dealing with the victims of that war — with the wounded, the traumatized, the people whose lives the machine had broken. On Fridays, my day off from school, I would spend it at her side, seeing refugees and people who had taken shelter in hospital corridors. Children my age with injuries I was too young to name. Mothers searching for sons. Men who had walked into that hospital with bodies and walked out with different ones. My father was an engineer. We were a middle-class family in Tehran — but middle-class in a city at war means watching long lines of people waiting for food, means unsafe streets, means a sky that sometimes lit up at distances you could not measure and your body understood before your mind did that what was burning was close enough to matter. The war lasted eight years. In 1986, when I was eleven years old, my family made the decision that every family caught in a war eventually must consider: how do we protect the children? My mother and I came to the United States on a visa. We left Iran to escape the war with Iraq — the same war that the United States government had encouraged Saddam Hussein to start and had armed him to sustain. I arrived in California speaking almost no English. I enrolled in ESL classes. I was an Iranian kid in the Bay Area in the years when the word "Iranian" was a slur, when the hostage crisis had already defined us in American culture as enemies, when a child had to learn not just a new language but a new identity that might be safe enough to walk through a school hallway. I grew up through middle school and high school carrying that weight. We were picked on. There was a lot of discrimination. I did not know what was going on. I did not understand the distinction between Saudi Arabia, between various Arab countries, between Iran — nobody was making that distinction. People were calling me a terrorist. I did not understand why. I was not a terrorist. My family were not terrorists. We had fled a war. And here in the country that received us, a different kind of fear had decided that what we were from was what we were. So when I got to college at age of 16, I focused everything I had on understanding my own identity. Who am I? I started the Persian Club in College. At the end of it, my answer was: I am Iranian-American. America is a melting pot. America stands for values — not for nationality, not for where your parents came from, but for a value system that all of us, as citizens, try to embody. That ideal sadly is not always honored. I want to tell you what my career actually was, because the critics will ask whether I am qualified to say what I am saying. I started selling computers. Retail. Then Apple products. Then I was selling networking hardware for Bay Networks. By the time I was twenty-one or twenty-two years old, I was making close to a quarter of a million dollars a year

working in the tech sector. I want to say that plainly, because people assume that the people who criticize Silicon Valley from the inside are people who were not good at it, people who could not compete, people who are bitter about failure. I was not bad at it. I was very good at it. The financial rewards were real. The energy of building things was real. The camaraderie of working late nights with a team on something that mattered — that was the most rewarding part of any of it. I moved through data center and colocation sales for Exodus Communications and AboveNet — two of the companies that built the physical backbone of the early commercial internet. I sold managed services. I got involved with startups and fell in love with the speed of them, the sense of building something from nothing, the feeling that what you were making was changing something. I was, in the words of a former co-founder at Exodus who said it with genuine warmth, the best salesman they had ever had. My former boss at the same reunion reminded me that when I sold one particular deal I had included in the contract a guarantee of a keg of beer for the client every month they were on Exodus. I don't remember doing this. He found it deeply funny. So did I. That's what building things in the early tech era felt like — absurd and energetic and full of people who believed we were making the world faster and better. Then I built WarrantyNow.com. I raised capital — more than twenty-five million dollars — from people whose names you would recognize: Halsey Minor, the founder of CNET; Marc Benioff, who went on to build Salesforce into one of the most powerful enterprise software platforms on earth; Michael Dell founder of Dell Computers; executives from PeopleSoft etc. I sat across the table from serious investors and serious institutional money and I held my own. I knew how that world worked, what it valued, what it dismissed. Then came BizCloud. In 2008 I went to Belgrade, Serbia, to build a cloud computing company from scratch. We built a real team. We produced over five thousand five hundred articles and videos on the cloud computing industry. We became a cloud broker, a meaningful presence in the space. And then Computer Sciences Corporation — an old-line technology company with billions of dollars in government contracts and a history of work for the intelligence and military complex that I was still learning to understand — appropriated my company's name for its own product line. I found out in 2010. They sued me in 2013 for using my own trademark. My lawyers, looking at the financial reality, made me mediate and settle. I lost as I had to share our good name with CSC. And what followed the loss was not a business setback in the ordinary sense. The pressure of that fight, combined with everything it had cost me, combined with the bankruptcy I filed from my hospital bed — it broke something. I made a decision that I have spoken about publicly, in my film, and in my first book, because I think it matters that people understand what this industry can do to the people inside it when the machine turns on you. After the settlement I made a decision that nearly ended my life. I survived because Stanford University Hospital and San Francisco General and Laguna Honda gave me four and a half months of the best care I could have asked for. I walked out with a cane. As I said at the time, with the dark humor that is the only way to survive such a thing: all the Humpty Dumpty and all the king's horses and all the king's men. When I came out, I said: fine. If I am going to continue, I am going to go to the biggest, most powerful cloud company on earth and understand it from the inside. I joined Amazon Web Services as Senior Alliances Manager in Big Data and Analytics. I sat in meetings with Andy Jassy — who would later become CEO of Amazon itself. I wrote executive briefing documents for Amazon's senior

leadership. I managed multi-million dollar co-marketing partnerships with Splunk, Sumo Logic, Tableau, and Qlik — some of the most important names in the enterprise data analytics ecosystem. I was on the commercial side. I did not touch government contracts. But I watched the machine up close, and what I watched in the offices of Amazon Web Services confirmed everything the New York Times had already described and everything my body had already learned at BizCloud: that these companies treat human beings as inputs, and when the input is no longer needed, they dispose of it. I did not enjoy the time I spent working at Amazon Web Services. Companies like Amazon are slave drivers. They do not value their employees as human beings. Talent is rewarded in Silicon Valley — but only for as long as it benefits the corporation. The moment it does not, you are gone. And if you dare to operate outside the construct that the employer makes for you, to ask questions about what the company is doing and who it is doing it to, your job is endangered. That is not speculation. I watched it. I lived it. --- After leaving Amazon, I founded Ethics in Technology. I organized NSA Comedy Nights after Edward Snowden's revelations — events that brought more than 150 people to each of three evenings, because it turned out people wanted a place to process what was happening to their privacy, to their data, to their understanding of the country they lived in. The nonprofit was formally incorporated in 2018. I published Ethics in Tech and Lack Thereof. I believed, at the time, that naming the problem was the beginning of solving it. I believed that public pressure worked. I closed Ethics in Technology in February 2022. What ended it was not defeat. It was recognition. The Ukraine war — and Silicon Valley's response to it, which was not moral reckoning but investor calls and contract proposals and product roadmaps for the defense sector — showed me that the strategy of reform had hit its limit. I launched NoEthicsInBigTech.com. The strategy changed from persuasion to prosecution. --- There is a philosophical framing I have carried throughout all of this work, and I want to name it here because it underlies everything in this book. The real struggle humanity is facing — beneath all the technology, beneath all the weapons, beneath all the political arguments about East and West and Islam and Christianity and NATO and Russia — is between two competing ideas. The first idea is I am one. The second is We are one. The whole of what Silicon Valley's war machine has produced is a monument to the first idea: I am one, my shareholders are one, my quarterly earnings are one. The whole of what this book is arguing for is the second: we are one, and the weapons and the surveillance and the exploitation of workers and the silencing of voices and the targeting algorithms that decide whose children live and whose children die — all of it is a crime against the we that includes all of us. I am Iranian-American. I came to this country at eleven years old speaking no English. I sat in hospitals on Friday mornings with my mother the nurse and watched what war produces in real human bodies. I watched it as a child and I have been watching it every day since, through every form the machine takes — the targeting system, the cloud contract, the content moderation algorithm, the warehouse productivity quota. And I have decided that the only responsible thing I can do with what I have seen is to document it as precisely, as sourced, and as honestly as one person working without institutional backing can manage. The film Forever Peace Now is part of that documentation. This book is part of it. ParentsPlea.com is part of it. The petition to Spain is part of it. None of it is the work of someone who has given up. All of it is the work of someone who understands, from a childhood spent watching what war does to

bodies and a career spent watching what tech does to souls, that evil men and their machines do not prevail — but that they do not stop on their own, and that the arc of history bends only because people push it. I have seen enough. I am pushing.

A LETTER TO THOSE WHO BUILD THE MACHINE

I want to speak directly to the tech workers reading this book. Not the executives. Not the investors. Not the PR departments that will be dispatched to rebut what I have written. The engineers. The data scientists. The product managers. The sales people. The alliances managers. The people who are, right now, sitting in an office somewhere in the Bay Area or in Austin or in Seattle or in Bangalore, building the tools this book is about. I was you. I was very good at being you. By the time I was twenty-two years old I was making a quarter million dollars a year in tech sales. I sold colocation at Exodus. I sold networking at Bay Networks. I raised capital from Marc Benioff and Michael Dell. I sat in meetings with Andy Jassy at Amazon Web Services and managed seven-figure co-marketing partnerships and wrote executive briefing documents for people whose decisions shaped the entire enterprise technology ecosystem. I was inside it. I was building it. I was good at it and I was proud of it. And I am telling you: choose carefully. There are three forces that drive most people in Silicon Valley, and I have seen all three up close. The first is the need for wealth. The second is the need for power. The third is the need for vanity — what you might call the fashion show, the race for the best car and the second home and the right brand of watch and the school your children attend. These forces are real. I am not naive about them. I chased them myself, for years, before the machine broke me and I understood what they actually cost. Talent is rewarded in Silicon Valley. But only for as long as it benefits the corporation. The moment it does not — the moment you ask the wrong question, the moment you sign the wrong petition, the moment you walk out of the wrong meeting — your job is in jeopardy. I have watched this happen. I have watched engineers who built the products that made these companies billions of dollars be walked out the door by security because they asked where their tools were being used. I have watched people I respected make the calculation that their mortgage and their family's stability were worth more than whatever discomfort they felt about what the company was doing. I understand that calculation. I made it myself, for a while. But I want you to understand what that calculation actually costs. It costs your conscience. And your conscience, it turns out, is not a soft asset. It is the thing that determines whether you can look at yourself in the mirror. Whether you can sit across from your children one day and tell them what you did with the years when you were smart and motivated and capable of building almost anything. Whether you can sleep without the question that will eventually arrive: did the work I did make the world better, or did it make it worse? The companies you work for have enough defenders. They have PR agencies and executive communications teams and lobbyists and lawyers. They can defend themselves. They do not need your silence to survive. What they need — what they have always needed — is your labor, your intelligence, your creativity, and your willingness to set aside the questions that your conscience is asking. I am not asking you to do what Daniel Ellsberg did. I am not asking you to do what Edward Snowden did. Both of those paths came at costs that most people cannot and should not be expected to absorb. Ellsberg would tell you: don't go to your superiors first, don't go to Congress first — they'll just label you a

troublemaker and cut you off. Snowden did it right. But not everyone is in a position to be Snowden, and the industry should not require that level of personal sacrifice before a conscience gets heard. What I am asking is simpler. I am asking you to check. To make sure that the work you are doing is ethical, that it is meaningful, that it represents the values you were raised with — not just the values of your stock options or your quarterly bonus. Does the company you work for represent your values? Does what it builds make you proud? Does it respect the human rights and digital rights of the people it affects? If it does not — go. Choose a different company. Choose a different role. Choose a different path. There are more options available to skilled technology professionals than in almost any other field in the world. Nobody who can write a production API or architect a cloud deployment or train a machine learning model is truly without choices. The choice is there. The question is whether you are willing to make it. I learned this the hard way. The hard way was four and a half months in hospitals after the machine I had trusted destroyed my company, stripped me of my resources, and left me with nothing to show for years of work except the evidence that the system is rigged toward the powerful. I rebuilt. I changed direction. I turned what had broken me into the work you are holding in your hands. But I would rather you did not need a crisis to ask the question I am asking you to ask right now. Do not end up like me — working at Amazon Web Services for Jeff Bezos and Andy Jassy, spending your days making a company richer that has already decided you are a replaceable part, contributing your talent to an empire that will thank you for your service and show you the door the moment you are more expensive than the automation that replaces you. You have more power than you believe. The machine runs on your expertise. That is also your leverage. Use it. Choose wisely. Choose the work that makes you proud. Choose the employer that represents the world you want your children to live in. And if you are already inside a company that you know is doing harm — with what you know, with what you can see from where you sit — consider whether this is the moment to say what you know. In the words of former US military analyst turned political activist, Daniel Ellsberg, who passed away in 2023, “The odds are against us. Snowden's greatest fear was that nothing would change. Manning's too. That is likely the outcome. But no one can say with certainty that change is impossible. The possibility is there, and the stakes are extraordinarily high. So the answer is to tell the truth, to take the risks, to say what you know when you know it. The question every person who knows that the truth is being concealed has to ask themselves is whether this is not the time for them to speak out.” I asked myself that question after Amazon. I answered it with this book. You will have to answer it with whatever you are building right now. Choose well. Because the machine does not stop unless the people running it decide it should. And right now, far too many of those people are looking at the same spreadsheet and calling genocide a golden opportunity. You are better than that. I know because I was one of you — and I was better than that, too, even if it took losing almost everything to prove it.

CHAPTER ONE — IN DEFENSE OF ETHICAL TECHNOLOGY

And the Use of AI for the Benefit of Humanity, Not to Our Detriment

When I moved to the United States at the young age of eleven, I didn't speak the English language. I barely knew a couple of words. I had to enroll in English as a Second Language classes all through middle school just to learn basic comprehension. Coming from Iran to the Bay Area in California carried a weight I hadn't anticipated — the negative connotations, the struggle to be accepted by classmates and teachers, the hostility of the political climate. Iran and the United States had a horrific relationship at the time, as they do today, and the way Iranians were portrayed in the media was anything but welcoming.

In those ESL classes, my instructors cared little about critical thinking. They focused on grammar, sentence structure, and punctuation. They gave us fiction books to read and placed no value in newspapers or current affairs. This was strange to me. I never enjoyed fiction. I always found truth to be stranger than fiction, and I loved staying current on what was actually happening in the world. With nonfiction I could reach a broad audience by sharing the truth — as opposed to simply entertaining people with stories I had invented. So I didn't show much enthusiasm for English class through high school.

It wasn't until I attended community college — graduating two years early from high school at sixteen — that I took my first Critical Thinking class at Diablo Valley College in Concord, California. My professor, Dr. Miller, introduced me to logical fallacies and the art of argument. We debated both sides of an issue in front of the class and had to construct solid arguments on social and political questions regardless of how we personally felt about them. That class, alongside a public speaking course I loved, gave me what ESL never could: the confidence to stand in front of a room and make an argument on issues I cared about, and to make it land.

As my work has shifted over the decades — from building tech startups, to working inside giants like Amazon Web Services, to human rights and digital rights advocacy — those two classes remain the most valuable things I took from my college education.

I say all of this because it explains why I am using AI as my editor and research partner in writing this book — and why I believe that choice is not a contradiction but a statement of principle.

I justify using the power of AI to do good. To the benefit of humanity. In the chapters that follow, we will examine how AI has been designed by the PayPal Mafia and the Epstein Class to the detriment of humanity. But I wanted to dedicate this first chapter to the other side of that argument: the ethical use of AI, used with good intentions, to craft well-reasoned arguments, to cite sources, to compress research time, and to help someone like me — who still carries the marks of his ESL years — communicate important ideas without being paralyzed by how an English teacher might grade the spelling.

In this day and age, content needs to be developed at a pace that would have seemed impossible a generation ago. The train is moving fast. The propaganda machine of these unethical Big Tech companies is sophisticated, well-funded, and tireless. As writers, authors, editors, and researchers, we no longer have the luxury of waiting years to produce a manuscript. We need to keep up and act fast.

Technology has allowed us to bear witness to the inhumanity of the chosen zionist and their crimes online. For the first time in our lives we all have witnessed the horror it has caused. We have all been traumatized and sadly some have become numb to it. That's why we have to act fast.

Before proceeding any further, I want to address the elephant in the room.

There are those who believe all AI is evil. That it uses too much electricity, too much water, that it contributes to global warming, that it has a direct impact on the people producing the raw materials and chips used in its development, that it will displace too many white-collar jobs.

Let me address each of these concerns. But first, let me draw a line between what I consider Ethical Technology and what I do not.

Some of the loudest voices against AI spend many hours each day on social media sites, interacting with posts and leaving disparaging comments just to troll and share their view. Others spend hours on gaming platforms like Twitch. Some spend valuable time running search engine queries for hours, being served advertisements in return, or trading cryptocurrency and engaging in high-frequency trading, or streaming content around the clock — but reserve their critique of resource consumption exclusively for AI. All of the above use utility power, water, processing capacity, storage, and generate e-waste. The majority produce no benefit to humanity whatsoever.

What is the benefit of running search engine queries for hours and having advertisements served to you, versus a focused AI search that gives you a direct answer? What does the world gain from the crypto economy? Hours spent on online gaming and content streaming platforms — what does that build?

When it comes to job displacement, that is a legitimate concern and I take it seriously. But let me ask this: I remember when photo development shops stood on every corner. When we made travel arrangements through travel agents. When we bought paper maps for every city we visited, or consulted travel guides that were three to seven years out of date, describing hotels and restaurants that no longer existed. All of that was replaced by technology. Make the argument for bringing it back — a good one, please — and I will listen.

Technology evolves. It can be used to benefit all of humanity. AI can simplify our lives. So can robotics. The question was never whether technology is good or evil. The question is always: who is building it, and what is their intention?

A fool with a tool is still a fool.

The Silicon Valley boys — the PayPal Mafia who dream of escaping to Mars while naming their children things like X Æ A-Xii — are, to put it plainly, deprived of basic human warmth. Peter Thiel and his ideological offspring built companies like Palantir not to advance humanity but to monetize surveillance, warfare, and the enabling of state violence. Alex Karp and Palantir are not building the future. They are profiting from the machinery of destruction. I don't say this

with hatred. I say it with pity. To be that cut off from what it means to be human is its own punishment.

What these men fundamentally lack is empathy. And empathy, it turns out, is not a soft skill — it is the very thing that separates wisdom from weaponry.

They remind me of Oppenheimer — not the Hollywood-glorified version, but the real one: the scientist who split the atom without pausing to ask whether he should. The artificial intelligence being built in their image is cold, calculating, and indifferent to consequence. It hallucinates confidence without genuine understanding. Watch a Peter Thiel interview sometime. The man speaks like a deer frozen in headlights — technically alive, emotionally absent.

And here is where it becomes interesting: the AI they have built is failing them in a very specific way. It hallucinates. It fills the gaps in its knowledge with confident fiction. But so do they. Leaders without empathy fill the gaps in their understanding with more cruelty, more surveillance, more violation of digital and human rights. Fools with the best tools are still fools — and their tools will not save them. If anything, the tools accelerate their unraveling.

Evil men and their machines do not prevail. History is long, and it is littered with their kind.

So then — what does AI actually do when it is built and used with conscience?

Let me give you seven answers.

AI and Medicine: Seeing What Human Eyes Cannot

DeepMind's AlphaFold program solved one of biology's most stubborn problems — the prediction of how proteins fold into three-dimensional structures — a challenge that had defeated scientists for fifty years. By 2022, AlphaFold had mapped the structures of over 200 million proteins, covering virtually every known protein in science. This single breakthrough is already reshaping research into Alzheimer's, Parkinson's, cancer, and antibiotic-resistant bacteria. The tool is free and open to researchers worldwide. This is not AI serving a quarterly earnings report. This is AI serving the sick, the vulnerable, and the scientists who spend their lives trying to help them. (Source: DeepMind AlphaFold Database, https://alphafold.ebi.ac.uk; Nature, July 2022)

Remote Surgery: The Surgeon Who Is Everywhere at Once

The da Vinci Surgical System has now been used in over ten million procedures worldwide, enabling minimally invasive surgery for cancer, heart disease, and trauma — with significantly reduced recovery times and complication rates. In 2022, surgeons at Johns Hopkins University performed the first fully autonomous laparoscopic surgery on soft tissue using an AI-guided robot called STAR — Smart Tissue Autonomous Robot — that operated without direct human guidance and outperformed human surgeons on several measurable metrics. The implications are vast. In rural areas, in conflict zones, in regions of the world where specialist surgeons are unavailable, robotic surgery guided by AI is not a luxury — it is a lifeline that is finally becoming real. (Source: Intuitive Surgical, https://www.intuitive.com; Science Robotics, January 2022)

Drug and Vaccine Discovery: Compressing Years into Months

The COVID-19 pandemic showed the world what AI-assisted drug discovery can do under pressure. BioNTech used AI to design and optimize the mRNA sequence for its COVID-19 vaccine at a speed that compressed years of traditional research into months. The AI platform Insilico

Medicine moved a drug candidate from initial AI design to clinical trial in under eighteen months — a process that traditionally takes twelve years or more. In 2023, an AI-designed drug for idiopathic pulmonary fibrosis entered Phase II clinical trials — the first fully AI-designed molecule to reach that stage. For rare diseases, neglected tropical diseases, and the pandemics we have not yet faced, this acceleration is not a technical footnote. It is the difference between life and death for millions of people. (Source: Insilico Medicine, https://insilico.com; Nature Biotechnology, 2023)

Energy and Power: AI Solving the Problem Critics Blame It For

There is a particular irony in the charge that AI wastes energy. Google's DeepMind applied machine learning to the cooling systems of Google's data centers in 2016 and achieved a 40 percent reduction in cooling energy — resulting in a 15 percent reduction in overall power usage. The same approach is now being used by energy grid operators worldwide to predict demand fluctuations and integrate renewable energy sources more efficiently. In California, AI-driven grid management tools are enabling utilities to reduce waste and incorporate solar and wind power at scale. The very technology that critics say consumes too much electricity is among the most powerful tools we have for managing the energy crisis they claim to be concerned about. (Source: DeepMind, https://deepmind.google; Lawrence Berkeley National Laboratory, AI for Energy Report, 2023)

Astronomy: AI Extending Humanity's Vision into the Cosmos

In 2019, the Event Horizon Telescope collaboration used machine learning to reconstruct the first-ever image of a black hole — the supermassive object at the center of galaxy M87 — from radio telescope data that no human analysis could have processed at the required scale. NASA's Kepler and TESS missions used AI to identify thousands of exoplanets from raw telescope data, including Earth-sized planets in potentially habitable zones. In 2023, astronomers using AI tools at the Vera Rubin Observatory detected thousands of previously unknown near-Earth asteroids in a single observational run. This is humanity using its most powerful tool to extend its vision into the cosmos — not for profit, not for market share, but for knowledge. (Source: Event Horizon Telescope Collaboration, https://eventhorizontelescope.org; NASA Exoplanet Archive, https://exoplanetarchive.ipac.caltech.edu)

Environmental Recycling: Teaching Microbes to Heal the Planet

Researchers at the University of Portsmouth, using AI-assisted protein engineering, modified a naturally occurring bacterial enzyme called PETase to digest the type of plastic used in most bottles — at dramatically accelerated rates. The AI was used to model and optimize enzyme mutations that would have taken decades to discover through conventional trial and error. Separately, researchers are using AI to identify microbial communities in soil and water that can break down PFAS "forever chemicals," heavy metals, and industrial solvents. We created these pollutants with industrial technology. We may be able to undo them with biological technology guided by AI. This is not theoretical. It is in laboratories and early field stages today. (Source: University of Portsmouth / NREL, Nature, April 2018; Environmental Science & Technology, 2023)

Witnessing Genocide: AI as the Memory That Cannot Be Erased

Perhaps the most morally urgent application of AI in our current moment is documentation. The Syrian Archive — now known as Mnemonic — has used AI-assisted tools to preserve, verify, and archive over 3.5 million pieces of digital evidence of war crimes in Syria: videos, satellite images, social media posts, protected from deletion and available for future prosecution. In Gaza, organizations including Forensic Architecture and Airwars are using AI-enhanced satellite imagery analysis and open-source intelligence to document strike patterns, civilian casualties, and the targeting of protected sites. The same AI technologies that weapons manufacturers are using to identify targets are being turned by human rights organizations toward the task of bearing witness. I find that deeply meaningful. Technology, as I have argued from the first page of this book, is not inherently good or evil. It is the intent of its creator — and the courage of those who use it to speak truth — that determines which side of history it stands on. (Source: Mnemonic, https://mnemonic.org; Forensic Architecture, https://forensic-architecture.org; Airwars, https://airwars.org)

Based on all of the above, I have made my decision. AI will be my editor. AI will help me with research. And together, we will move as fast as this moment demands.

--- The ethical use of AI — the case I have made in this opening chapter — did not emerge in a vacuum. It emerged as a direct response to what I watched the technology industry become across three decades spent inside it. Chapter Two is where that history begins: not with abstract policy failures, but with the specific, documented, personal arc that brought me from the server farms of Exodus Communications to the targeting algorithms of Project Nimbus — and to the conviction that what I was watching was not a new story, but the oldest one capitalism tells.

CHAPTER TWO — HOW WE GOT HERE

As I documented over the years — in my many conversations with Silicon Valley executives, and in my film Forever Peace Now — the people running the technology sector are not concerned with the betterment of humanity. They think about individual financial gains, about "winning" at any cost, and about market dominance, regardless of the human consequences. I watched this up close. I was inside it. One conversation stays with me above all others. After leaving Amazon Web Services, I was in the job market and interviewing at a tech company whose name I will let the reader imagine. The Senior Vice President of Marketing told me, with what I can only describe as pride, that the U.S. Air Force was one of their most valuable reference customers — because the company's technology was helping them kill and target people more effectively. He said it the way a man talks about winning a sales award. Near the end of the interview, he gave me the floor. Ask anything, he said. I asked a simple question: as a company, where do you draw the line on customers and projects? Do you have a red line? He looked at me as though I had spoken in a language he didn't recognize. What do you mean, he said. We don't turn away business. If it's a registered business, we'd love to have them as a client. Silicon Valley has no red lines. They will lobby for any contract, pursue any customer, and sell their souls to any buyer willing to sign a check. This is not an opinion. It is a business model.

Who I Was When I Walked Into Those Rooms I want to stop here and tell you who I was when I first walked into those boardrooms. Not who I became. Who I was. I was twenty-two years old. I had come from a family of six living in a three-hundred square foot apartment in Berkeley, California. We were not poor in the way that word gets used as an abstraction — we were poor in the way that means six human beings sharing a space smaller than most American living rooms, making it work because that is what you do when you have no alternative and you are grateful to be alive and in a country that is not actively bombing you. I had graduated high school at sixteen, gone to community college, started a company before most of my peers had their first real job, and by twenty-two I was making a quarter million dollars a year in tech sales. I sat in boardrooms with men who ran companies worth billions of dollars. I raised capital. I negotiated contracts with nine figures attached to them. I wanted to be a good American capitalist. I say that without irony and without apology — because I think it is important to be honest about who you were before you understood what the system actually was. I came from Iran, from a war, from a country where the economy had been strangled by sanctions and ideology and the kind of revolutionary fervor that burns everything down and calls the ash freedom. I came to the United States and found that a person could build something, could earn something, could sit at a table with powerful people and be heard. I wanted that. I chased it. I was good at it. And then I had stock options. Founders in Silicon Valley receive common shares. This sounds like a good thing until you understand what it means in practice. You have the idea. You assemble the team. You work the hours that no employment contract could capture — the midnight calls, the weekend pivots, the months when the product is broken and the money is almost gone and you are the

only person in the room who cannot afford to lose faith in what you are building. You do all of that. And in return, you receive common shares — the lowest class of equity in the capital structure, the last in line when anything is distributed, the first to be diluted. The investor who writes a check in a Series A round gets preferred shares. Not common shares. Preferred. A different class entirely, with their own veto rights, their own liquidation preferences, their own protections that can be written into each subsequent funding round in whatever configuration their lawyers negotiate. The founder who had the idea and built the thing gets common shares. The investor who showed up with a checkbook gets preferred shares and a seat at the table that the founder cannot remove even when the investor's judgment is wrong. This is the foundational lie of Silicon Valley's meritocracy narrative. The system is not designed to reward the people who build things. It is designed to protect the people who fund things. Those are not the same people. And the gap between them — the gap between common and preferred, between the founder who cannot sell and the investor who can — is where most of the wealth that should flow to builders gets redirected to capital. I found this out the hard way. I had stock options in my company — what I call funny money, because at the time there was no secondary market for founder shares. I could not sell them to investors for cash. I could not convert them to anything liquid. I had paper wealth and a vision and the particular kind of poverty that comes from having something that looks like money but cannot be spent. What I wanted to do with those options — if they ever converted to something real — was donate a large portion to Doctors Without Borders. I mean that literally. I was twenty-two years old, I had come from a family where six people lived in three hundred square feet, and my first thought about what to do with hypothetical wealth was to give a significant share of it to the people providing medical care in war zones. I remembered what my mother had seen as a nurse during the Iran-Iraq war. I remembered the hospital corridors. I couldn't have articulated it as a political position at the time. It was just what felt right to a kid who understood, from personal experience, that the distance between having medical care and not having it was the distance between living and dying. The system did not accommodate that impulse. The system was not built for it. That is the moment I began to understand something that would take me another twenty years to fully articulate: the system was not broken. It was working exactly as designed. It was designed to concentrate capital in the hands of the people who already had it, to protect their preferred position in every subsequent round, to ensure that the funny money flowed upward — and that the founders and the workers and the communities that made the whole enterprise possible were positioned, legally and structurally and by deliberate design, to receive what was left after the preferred class had taken what they were owed. I was in boardrooms with these men. I watched them operate. I know how some of them acquired their wealth. I know how certain domain names were obtained at a time when the process governing them should have prevented it — when what was supposed to be an open, rules-based system for how the internet organized itself turned out to have rooms where different rules applied to the people who knew which doors to knock on. I know the mechanics of the deals I am about to describe not from reading about them but from having been present when the shape of those arrangements was discussed. I am not going to name every person I know from that world. I am too young to die, and I understand that the people who expose powerful men too specifically sometimes find that the legal system — the same legal system that gave Jeffrey

Epstein thirteen months for crimes that sent ordinary people away for decades — can be weaponized against the person doing the exposing rather than the person being exposed. But I will say this: I know who these men are. I know what they value and how they talk about people who are not in the room. I know what power looks like up close when you are twenty-two years old and hungry and sitting across the table from it. And what I can tell you is that the mythology — the visionary founders, the meritocracy, the disruption in service of humanity — is a story they tell in public. What they tell each other in private is something different. What they have built is something different. And the distance between the public story and the private one is the distance this entire book is trying to map. I got out. I chose something different. I am writing this book instead of cashing checks from the people I could be defending. That choice cost me money. It cost me relationships. It cost me the kind of professional comfort that comes from staying inside the network and keeping your head down and your opinions private. I would make it again tomorrow.

The PayPal Mafia I want to stop and define my terms — because I have been using two phrases throughout this book that deserve more than a passing reference, and because the people those phrases describe have spent considerable resources ensuring that the public understanding of who they are remains flattering. The first phrase is the PayPal Mafia. Silicon Valley uses this term proudly. It appears in business magazines as a celebration — a group of extraordinarily successful founders who met at a startup called PayPal in the late 1990s, built it into something valuable, sold it to eBay in 2002 for $1.5 billion, and then scattered across the landscape to build Tesla, LinkedIn, Palantir, YouTube, Yelp, and a dozen other companies that have shaped the modern world. The narrative is one of brilliance and ambition and the compounding returns of great networks. I never thought the word mafia was something anyone should be proud of. A mafia is not a network of ambitious people who know each other. A mafia is a criminal organization — a group of men who operate outside the law while presenting a legitimate face to the world, who use their collective power to bend rules that apply to everyone else, and who protect each other from the consequences of what they do. The fact that Silicon Valley adopted the term and made it a brand tells you everything you need to know about how these men see themselves in relation to the law and to the rest of us. The core of the PayPal Mafia as I use it in this book includes Peter Thiel, who co-founded PayPal and went on to build Palantir — the surveillance and targeting company whose tools I have documented operating in Gaza, in ICE deportation operations, and in the NHS health records of 67 million British citizens. It includes Elon Musk, who came into PayPal through the merger of his company X.com, and who now controls the world's largest social media platform, a private space company, a government efficiency operation, and the AI tool I documented producing antisemitic content and Holocaust denial. It includes Reid Hoffman, co-founder of LinkedIn, early PayPal executive, and a man whose platform I documented operating a covert surveillance system scanning 6,167 browser extensions without user knowledge or consent. It includes the network of investors, founders, and operators who move in and out of these companies through Founders Fund — Thiel's venture capital firm — and through the interlocking board memberships and investment relationships that connect these companies to each other in ways that are, on paper, separate transactions and, in practice, a coordinated system of mutual enrichment. And here is how that

system works — because it is important to understand that when I call them a mafia, I am not speaking metaphorically.

The Lawyers Who Were Never Yours

I want to talk about the people you hired to protect you. You walked into a funding negotiation and you hired a lawyer. That is what you are supposed to do. You paid them — sometimes five hundred dollars an hour, sometimes a thousand, sometimes more, depending on the firm and the partner and the prestige of the address on Sand Hill Road. You signed their engagement letter. You assumed, because you were paying them and because that is what engaging a lawyer means in every other context in American life, that they were representing your interests. Let me be direct about something that took me longer than it should have to fully understand: in most cases, they were not. The law firms that dominate Silicon Valley startup work — Wilson Sonsini Goodrich and Rosati, Cooley, Orrick, Fenwick and West and the others whose names you see on every term sheet and every funding announcement — do not primarily serve founders. They serve the ecosystem. And the ecosystem's center of gravity is not the founder. It is the venture capital firm. The VC is the repeat client. The VC is the reliable pipeline of deals, year after year, fund after fund, portfolio company after portfolio company. The founder is a transaction. The VC is a relationship. Think about what that means in practice when you sit down to negotiate your term sheet. The lawyer across the table from you — the one you are paying by the hour to fight for your interests — has known the partner at that VC firm for years. They have been to the same holiday parties. Their firms collaborate on dozens of deals simultaneously. The VC sends referrals. The law firm sends favorable terms. Not in writing. Not in any way that could be called corruption in a court of law. In the way that relationships work when money and proximity create alignment of interest over years. There is a phrase in that world that you will hear often. Market standard. When you push back on a term in a term sheet — the liquidation preference, the anti-dilution clause, the vesting acceleration, the board composition — you will be told it is market standard. Your lawyer will often be the one telling you this. What market standard means, in practice, is that this term has been used in enough deals, drafted by enough firms representing enough venture capital funds, that it has become the template. The template was not designed by founders. It was designed by the people negotiating across the table from founders. Market standard is the result of decades of the same law firms writing the same boilerplate on behalf of the same investors — and then telling you that the boilerplate is neutral. It is not neutral. It is the accumulated sediment of a thousand deals in which the lawyers were closer to the money than to the people who built the thing. Wilson Sonsini is the oldest and most prominent example of how deep this goes. The firm created an investment fund in 1978 — WS Investments — specifically so that its partners could take equity stakes in the client companies they were advising. They turned a seventy-two-thousand-dollar investment in Google into a twenty-eight-million-dollar return after the IPO. The same lawyers advising founders on how much to disclose in their IPO filings had a direct financial stake in maximizing the IPO price. A University of Illinois law professor put the conflict plainly: the more you disclose, the lower the IPO price. That creates a conflict of interest because if you have an economic interest in the highest IPO price, will you really advise them to disclose? The answer, which the professor left

tactfully unspoken, is: probably not. Not fully. Not in every case where disclosure might cost the fund money. The firm's name partner, Larry Sonsini, sat on the board of Brocade Communications while Wilson Sonsini served as Brocade's outside legal counsel. He reportedly encouraged the board to give the company's CEO sole authority to grant stock options with minimal oversight — a "committee of one." That committee of one became the center of an options backdating scandal that led to criminal prosecution. Many of the Silicon Valley companies investigated for options backdating at the time were Wilson Sonsini clients. When Brocade needed defense counsel in the resulting derivative lawsuit, a judge raised concerns that Wilson Sonsini was so tainted by conflicts of interest that the company had to drop the firm entirely. In a separate case, a company called Existence Genetics sued Wilson Sonsini for failing to disclose that while it was representing Existence Genetics, it was simultaneously representing a direct competitor — Navigenics. Two clients. One firm. Competing interests. No disclosure. These are the documented cases. The ones that became lawsuits, that produced court filings, that got written up in legal trade publications. The undocumented version of this — the version that never becomes a lawsuit because the founder doesn't have the money or the evidence or the energy to fight a firm with two thousand attorneys — is orders of magnitude larger. And then there is the board seat. Some of these law firms, or their senior partners individually, will want something beyond hourly fees. They will want equity. An advisory role. A seat at the table that gives them ongoing access to company information and company decisions long after the initial deal is done. They will frame this as alignment — as skin in the game, as a signal that they believe in what you are building. What it actually is, is a mechanism for converting a service relationship into an ownership relationship, one that gives them preferred access to information that other shareholders do not have and creates a financial incentive that runs parallel to, and will sometimes conflict with, their legal obligation to advise you honestly. There is something that the smarter practitioners in this world will admit privately: when a VC firm recommends a law firm to a founder, the founder should ask a very simple question. Why does this particular VC prefer this particular firm? What is it that this firm offers this VC in return for that referral? The answer, in most cases, is not that the firm produces better outcomes for founders. The answer is that the firm produces smoother transactions for investors — deals that close faster, with less friction, with terms that are more favorable to the party that will be back next month with another deal to do. You are not that party. You are a transaction. You will not be back next month. And the lawyer you hired knows this. I am not telling you that these firms never produce good work. Some of their lawyers are genuinely skilled, genuinely committed to their clients, and genuinely uncomfortable with the conflicts their firm's business model creates. I know some of them personally. The problem is not individual lawyers. The problem is structural. The business model of the dominant Silicon Valley law firm is built on serving both sides of the same table — the founders who pay the hourly fees and the venture capital firms who generate the deal flow — and calling it representation. You cannot serve two masters. Not when the interests of those masters diverge at the exact moment that matters most — which is the moment you are sitting across a term sheet from someone whose goal is to maximize their preferred return at your expense. The boilerplate they hand you is not neutral. The market standard they invoke is not a law of nature. It is a document written by people who were not thinking about you when they

wrote it. And the lawyer who tells you to sign it, who tells you it is standard, who tells you that pushing back will make you look difficult and harm your relationship with the investor — that lawyer is telling you something that is in their interest to tell you. Before you hire a lawyer in Silicon Valley, ask them one question: who referred you to me, and what do you owe them? If they cannot answer that question directly, you already have your answer.

The Oracle Problem — Or, How They Sold You a Map of a Country That Doesn't Exist

Before we get to what happens in the room with the investor, I want to talk about what happens before you ever walk in the door. Because the equity trap I just described — the preferred versus common, the liquidation preferences, the cap table that was never designed in your favor — that trap has a predecessor. And the predecessor is the number on the slide. You know the number I mean. The one that says your addressable market is four billion dollars. The one that says your segment will grow at thirty-two percent compound annual growth rate over the next five years. The one with the logo of a firm that sounds authoritative — Gartner, Forrester, IDC, take your pick — stamped in the corner like a government seal. The number that, when you were twenty-three years old and building your first company and trying to convince anyone that what you were building was worth funding, felt like permission. Like proof. Like someone credible had looked at the landscape and confirmed that the thing you believed in was real. Let me tell you what that number actually is. I am going to tell you this from an unusual position. I have sat in every chair at this table. I sold data center infrastructure in a direct sales capacity — I was the person presenting the contract and explaining why the terms were in your interest. I was also a customer of these same companies — I signed the agreements, accepted the credits, and learned what the fine print meant when the projections did not materialize. And I worked as a senior alliances manager inside Amazon Web Services, managing big data and analytics partnerships, which means I was present for the Proof of Concept credit conversations, the AWS All-In commitment negotiations, and the internal discussions about what we expected in return. I am not describing this system from the outside. I built parts of it. I operated inside it. And I am telling you exactly how it works. Those analyst firms are not neutral observers of the technology market. They are paid participants in it. The companies that appear favorably in their Magic Quadrants — the leaders, the visionaries, the firms positioned in the upper right corner of every chart — are, in most cases, also paying clients of the firms producing those charts. The research is sponsored. The coverage is purchased. The analysts who left large technology companies and landed at Gartner or Forrester or IDC did not stop serving the interests of large technology companies when they changed business cards. They institutionalized that service. They built a system in which the companies with the largest marketing budgets receive the most favorable positioning, and then sold access to those rankings as objective research. This is not a secret. It is a business model. It is disclosed, in the fine print, in language calibrated to be present without being legible. And it has been running long enough, and cited often enough, that the entire industry has forgotten to question it. The result is a generation of founders — many of them brilliant, most of them young, almost all of them operating with limited real industry experience — who built their financial models on numbers that were, at their origin, marketing artifacts dressed as market research. They took the analyst figure and built a five-year revenue projection.

They took that projection and walked into a meeting with Amazon Web Services, or Microsoft Azure, or Google Cloud, and said: here is our growth curve. Here is where we will be in eighteen months. Here is where we will be in three years. And the cloud providers looked at those numbers and said: great. We believe you. Let us structure a deal around that trajectory. Nobody in that room questioned the analyst numbers. Not the founder, who needed them to be true. Not the cloud provider's sales team, who needed the committed revenue the numbers justified. The fiction was useful to everyone sitting at the table, so no one at the table called it a fiction. I want to be precise about what that deal looks like — because I have seen it from both sides of the table, and the structure is important. The large cloud providers will offer a startup meaningful discounts on infrastructure — sometimes fifteen percent, sometimes thirty percent — in exchange for multi-year commitments tied to projected consumption. The logic they present to the founder sounds generous. You are going to need this infrastructure anyway. Lock in now, get the discount, plan your budget with certainty. They will sweeten the entry point with what they call Proof of Concept credits — anywhere from ten thousand to a hundred thousand dollars in free usage — to get you building on their platform before the commercial terms kick in. That credit is not generosity. It is a migration subsidy. It is designed to get your engineers building on their infrastructure, your data living in their storage, your architecture dependent on their proprietary services — so that by the time the commercial contract begins, the cost of leaving is no longer a commercial question. It is an engineering question. A staffing question. A six-to-twelve month project that your operational budget cannot absorb and your investors will not fund. Inside AWS, we expected a ten-to-one return on every dollar of Proof of Concept credit issued. Ten dollars of future contracted revenue for every dollar extended. I want you to sit with that. I want you to think about what kind of institution demands a ten-to-one return on a short-term credit instrument extended to a company with no proven revenue. Not a bank. A bank charging those terms would be in violation of usury law in most jurisdictions. Not a reputable investor. A thirty percent annual return would make most investors ecstatic. The only class of lender that expects ten-to-one on a short-term extension of credit is the kind you do not want to owe money to. We dressed it in the language of partnership. We called it investment in your success. It was neither. And then the revenue projections — built on the analyst numbers, built on the market size figures, built on the Magic Quadrant that a paying client commissioned — do not materialize. Because they were never real. Because a company staffed by twenty-three-year-olds building their first product in a segment they learned about from sponsored research does not capture a significant percent of a four billion dollar market in three years. Now the founder is on the hook for the contract. Infrastructure costs are fixed and revenue is not. And now the cloud provider — which has already captured the architecture, the data, and the engineering team's institutional knowledge of their platform — presents the AWS All-In program. I want to be clear about what AWS All-In actually is, because the name is doing a great deal of work. It is presented as a strategic partnership. A vote of confidence. An opportunity to consolidate your infrastructure and unlock deeper discounts in exchange for a long-term commitment to run everything — every workload, every service, every corner of your technical stack — on AWS. What it actually is, is a formalization of the trap you are already in. By the time AWS All-In is on the table, you have already migrated your infrastructure. Your engineers already

know AWS. Your data is already in S3. Your pipelines are already running on their proprietary services. The question is no longer whether you are dependent on AWS. The question is whether you will sign the document that acknowledges it and commits you to deepening that dependency in exchange for a discount on the rates you are already paying. The cost of leaving — the real cost, not the commercial cost — is measured in engineering hours. In months of rebuilding. In the operational budget line that your CFO will look at and say: we cannot afford this right now. Not this quarter. Not next quarter. And by the time the quarter after that arrives, the dependency is older, the engineers who built on AWS are more fluent in AWS than in anything else, and the cost of migration has not gone down. It has gone up. This is not an accident. This is the product. The infrastructure is the delivery mechanism. The lock-in is what they are selling. And the Proof of Concept credit — the ten thousand dollars, the hundred thousand dollars, the number that felt like a gift when you were a founder trying to keep the lights on — was the down payment on a decade of dependency. The analyst numbers manufactured the illusion of a market. The market projections justified the infrastructure commitments. The infrastructure commitments created the lock-in. The lock-in made AWS All-In feel like an option rather than a conclusion that had already been written. And the dependency is what the cloud providers were building toward from the first conversation. This is the trap that comes before the equity trap. And it is designed by people who understood exactly what they were building. If you are a founder reading this — if you are sitting somewhere right now with an analyst report in one tab and a Proof of Concept credit offer in another — I want you to understand something that nobody said to me when I was on your side of the table. The number in the corner of the slide is not evidence. It is the beginning of the trap. The deals are made in boardrooms, with a handshake and a wink. Company A — let us say a large cloud infrastructure company or an AI platform — decides to invest in Company X, a startup in an adjacent space. The investment is made on a given day. Thirty to ninety days later — at a carefully chosen interval designed to ensure that the two transactions appear unrelated in regulatory filings — Company X announces a series of commitments back to Company A: stock warrants, stock options, and a structured multi-year revenue agreement for Company X to purchase products or services from Company A. The investment and the revenue commitment were agreed upon simultaneously, in the same room, on the same day. But because they are executed on paper as separate transactions spaced weeks apart, they avoid the appearance of the quid pro quo arrangement they actually are. What this produces is the following effect. Company A's books show a capital investment in Company X as an asset, and simultaneously show Company X's revenue commitments as forward-contracted income — hyper-growth on paper, the kind of growth that drives share price, that drives analyst upgrades, that drives the valuation multiple that makes every insider's equity worth more. When Company X eventually fails to meet its revenue targets — as many of these startups do — Company A moves the warrants, the expected revenue, and any stock received to the loss column. Those losses are spread across three to five years of write-downs, timed to minimize their impact on any single earnings period, absorbed quietly into the accounting without the kind of market reaction that a single large write-down would produce. The stock price boost from the original deal has already been captured. The insiders who benefited from that boost have already been compensated. The write-down is paperwork. This is how the funny money of Silicon Valley

operates. This is how companies like OpenAI, Oracle, Amazon, and Nvidia participate in a circular ecosystem of investment and revenue that creates the appearance of explosive growth while the underlying economics are, in many cases, paper transactions moving from one column to another. It is not a coincidence that the same names appear repeatedly across these investment networks — investing in each other, contracting with each other, serving on each other's boards, and writing each other's reference letters when the SEC asks questions. I am not a securities lawyer. I am documenting what I have seen and what the public record reflects. What I will say plainly is this: the distance between what these men call aggressive accounting and what the some in law would call securities manipulation is, in many cases, measured in days on a calendar — the days between the handshake and the paperwork.

The Epstein Class The second phrase I use throughout this book is the Epstein Class. And I want to be precise about what I mean — because it is not synonymous with the PayPal Mafia, although there is significant overlap, and because the defining characteristic of the Epstein Class is not association with Jeffrey Epstein before his first conviction. It is association with him after it. Jeffrey Epstein was convicted of soliciting prostitution from a minor in 2008. He registered as a sex offender. He served thirteen months of an eighteen-month sentence — a sentence so light, secured through a non-prosecution agreement negotiated by then-U.S. Attorney Alexander Acosta, that it represented one of the most visible examples in recent American legal history of wealth purchasing its way out of consequences that would have destroyed any ordinary person. He was a convicted sex offender operating in the open. And after that conviction, certain men chose to continue associating with him. Not men who didn't know. Men who knew, and chose to continue anyway — because Epstein's network of money, access, and connections was worth more to them than the moral cost of the association, or because they believed, as men of that class often believe, that the rules applying to Epstein did not apply to them by extension. Bill Gates met with Epstein at least six times after his 2008 conviction, according to reporting by the New York Times and acknowledged by Gates himself. Gates has called this a mistake. I believe him that he now regrets it. I do not believe the meetings were accidental. A man of Gates' stature does not accidentally schedule six separate meetings with a convicted sex offender. Reid Hoffman — co-founder of LinkedIn later acquired by Microsoft, Bill Gates Company and a figure whose connections bridge both the PayPal network and the broader Silicon Valley philanthropic infrastructure — attended dinners with Epstein after his conviction and has publicly acknowledged doing so. Hoffman has said he was trying to secure funding for the MIT Media Lab through Epstein's connections. Joi Ito, the director of the MIT Media Lab, subsequently resigned after it emerged that he had accepted Epstein money and concealed its source. The philanthropic infrastructure of elite American universities was, for years after Epstein's conviction, being maintained in part by money that came through a convicted sex offender — and the men who facilitated those flows knew where the money came from. Peter Thiel's relationship with this network is documented through multiple channels, including video recordings in which Ehud Barak — the former Israeli Prime Minister who had one of the most extensively documented relationships with Epstein in the post-conviction period, including financial transfers and repeated visits to Epstein's properties — references Thiel's association and introduction into that network. Thiel is simultaneously a co-founder of PayPal alongside Musk,

the founder of Palantir, a primary funder of far-right political candidates and causes in the United States and Europe, and a man who has stated publicly that he does not believe freedom and democracy are compatible. That last statement is not an interpretation of his views. It is a direct quote from his 2009 essay "The Education of a Libertarian," published in the Cato Unbound journal, in which he wrote: "I no longer believe that freedom and democracy are compatible." A man who builds surveillance infrastructure, funds authoritarian political movements, and has openly stated his opposition to democratic governance is not a complicated figure. He is a man who has told you exactly who he is. The question is whether we have chosen to listen.

Why the Distinction Matters I separate these two categories — PayPal Mafia and Epstein Class — because they represent different species of the same genus. The PayPal Mafia is a financial and political network held together by shared capital interests, interlocking equity stakes, and the mutual protection of valuations built on the accounting architecture I described above. The Epstein Class is something darker: a social network held together by shared knowledge of shared crimes, and by the understanding — implicit, never spoken, understood by everyone in the room — that the protection flowing from that network runs in both directions. What connects them is power. The power to make the SEC look away. The power to make the DOJ cut a deal. The power to make the MIT Media Lab take the money and ask no questions. The power to ensure that the documents stay sealed, that the flight logs stay redacted, that the names stay private, that the record stays incomplete. That power has a cost. The cost is paid by the people who are not in the room when the handshake happens. It is paid by the startup founders whose companies were consumed by the circular investment machine and whose failure was someone else's write-down. It is paid by the workers in the warehouses and the content moderators in Nairobi and the engineers whose keystrokes are being captured to train the models that will replace them. It is paid by the women and girls who were in rooms with Jeffrey Epstein and could not leave. It is paid by every person whose suffering was, to the men in these networks, an externality — a cost to be absorbed by someone else, written off on someone else's books, buried in someone else's accounting period. I am documenting the cost. One chapter at a time.

We Have Been Here Before The blood-soaked hands of Silicon Valley's Big Tech executives are complicit in the enablement of the worst genocide the 21st century has witnessed — what is unfolding in Gaza. But before we examine the present, we must look at the past. Because we have been here before. And the names of the corporations involved are ones you know. In the 1930s, Adolf Hitler built his war machine and carried out the Holocaust with the direct assistance of major American corporations. U.S. bankers financed his regime through debt restructuring led by the houses of Morgan and Rockefeller. Standard Oil enabled the Nazis to refine synthetic fuel via the hydrogenation of coal — a technology critical to powering the German military. Henry Ford, a proud antisemite, provided the regime with Ford-Werke, which supplied transport vehicles for personnel and weapons. General Motors, through its Opel subsidiary, became Germany's largest automaker and built the trucking infrastructure alongside Ford for the German war effort. International Telephone and Telegraph, under CEO Sosthenes Behn, purchased a major stake in Focke-Wulf — the manufacturer of deadly fighter planes — and supplied the Nazi state with advanced telecommunications, radar components, and managed

telephone systems across occupied Europe. General Electric held cross-licensing agreements and patents with Siemens and AEG, and was a partner in forming IG Farben. Alcoa made agreements with IG Farben that restricted American production of magnesium — a key aircraft metal — handing Germany a significant technological advantage. IBM's German subsidiary, Dehomag, provided the Nazi state with punch-card tabulator systems. These were not simple counting tools. They were sophisticated data processors that enabled the bureaucracy of genocide — used to identify, categorize, and track populations for census, conscription, and the persecution of Jews and other targeted groups. Each of these corporations collected royalties on German military technology throughout the war. The Trading with the Enemy Act, passed in 1917, meant nothing when profit was available. These companies acted with impunity and enabled a genocidal war machine. Today, as genocide unfolds in Gaza, the names have changed. The pattern has not. Amazon and Google provide the foundational cloud and AI infrastructure through "Project Nimbus," a $1.2 billion contract that supplies the Israeli military with the data storage and processing power required for surveillance, intelligence, and real-time targeting. Palantir acts as the central brain, using its data fusion platforms to integrate intelligence from drones, satellites, and spies — creating an operational picture used for precise targeting that has killed tens of thousands of civilians, aid workers, and journalists. Oracle provides the critical database and cloud systems that manage vast surveillance data, forming the backbone of an intelligence apparatus used in a documented campaign of collective punishment. Microsoft powers the military's cloud infrastructure and AI tools, which process surveillance data and coordinate operations that have devastated civilian life across Gaza and Lebanon. The products built with this infrastructure have names that sound almost bureaucratic — clinical labels designed to keep the people who authorize them at a comfortable emotional distance from what they actually do. Gospel. Lavender. Where's Daddy. Blue Wolf. Red Wolf. Project Maven. The names change. The machinery does not.

The Gospel — The Mass Assassination Factory The Gospel, known in Hebrew as Habsora, was developed by Unit 8200 of the Israeli Intelligence Corps — the same elite military technology unit whose graduates flow directly into Silicon Valley and Israeli defense tech startups. The Gospel is an AI system designed to generate bombing targets automatically, at a rate no human intelligence operation could produce. Before the Gospel, Israeli military analysts could generate approximately 50 targets in Gaza per year. Once the Gospel was activated, it was producing 100 new targets per day. In the first weeks of the current genocide on Gaza, the Israeli Military claimed to have struck over 22,000 targets inside the Strip — at a daily rate more than double any previous operation. The Gospel does not merely assist human decision-making. According to a landmark investigation by the Israeli publication +972 Magazine, based on testimony from current and former Israeli intelligence officers, the system "generates targets almost automatically." One former intelligence officer described it plainly: a mass assassination factory. The system was used to justify strikes on private homes of junior Hamas members — and on homes where no known militant lived at all. The humans reviewing Gospel's recommendations were, by multiple accounts, rubber stamps. Some approved strikes in as little as twenty seconds. (Source: +972 Magazine and Local Call, "A Mass Assassination Factory: Inside Israel's Calculated Bombing of Gaza," https://www.972mag.com/mass-assassination-factory-israel-

calculated-bombing-gaza/; Democracy Now!, December 2023, https://www.democracynow.org/2023/12/1/israel_gaza_war_gospel_artificial_intelligence)

Lavender — The Kill List Generator Where the Gospel targeted locations, Lavender targeted people. It is an AI system that assigned every adult Palestinian in Gaza a numerical score from 1 to 100, rating the probability that each individual was a member of Hamas or Palestinian Islamic Jihad. At the outset of the current war, Lavender had flagged up to 37,000 Palestinians as potential assassination targets. The system's own architects knew it carried an error rate of approximately 10 percent — meaning that by their own internal calculation, roughly 3,700 people on the kill list were incorrectly identified. That error rate was considered acceptable. Officers who used the system testified to +972 Magazine that they treated Lavender's output "as if it were a human decision" — no independent verification required, no review of the underlying data. The AI had decided. The bomb fell. The system also flagged men based on criteria as broad as being in a WhatsApp group with a suspected operative, or sharing a name with someone on a list. Police officers, civil defense workers, and innocent family members were swept into the machinery. This is not targeted counterterrorism. It is algorithmic mass killing, validated after the fact by a human signature that had less than twenty seconds to think. (Source: +972 Magazine and Local Call, "Lavender: The AI Machine Directing Israel's Bombing Spree in Gaza," https://www.972mag.com/lavender-ai-israeli-army-gaza/; West Point Lieber Institute, https://lieber.westpoint.edu/gospel-lavender-law-armed-conflict/)

Where's Daddy — Killing Families to Kill One Person If the Gospel identifies what to bomb and Lavender identifies who to kill, Where's Daddy completes the system by identifying when. It was built specifically to track individuals on the Lavender kill list and alert military operators the moment those individuals entered their family homes at night — so that strikes could be timed for maximum lethality. Not on the street. Not in transit. At home. With their families. The name itself tells you everything about the moral universe in which this technology was developed. According to testimony from an Israeli intelligence official quoted in +972 Magazine, the Israeli Military "bombed them in homes without hesitation, as a first option" — meaning the family home was not a last resort but the preferred strike environment. Children were present. Spouses were present. Extended families were present. The system knew this. The operators knew this. The strike happened anyway. The only question the technology asked was: is the target home yet? (Source: +972 Magazine / Local Call, https://www.972mag.com/lavender-ai-israeli-army-gaza/; Wikipedia, AI-assisted targeting in the Gaza Strip, https://en.wikipedia.org/wiki/AI-assisted_targeting_in_the_Gaza_Strip)

Blue Wolf and Red Wolf — Biometric Occupation In the occupied West Bank, the architecture of AI-enabled control takes a different but equally chilling form. Blue Wolf is a smartphone application deployed to Israeli soldiers at military checkpoints. When a Palestinian approaches, the soldier photographs them. Within seconds, Blue Wolf matches the face against a vast biometric database — cataloguing name, family connections, address, and any "negative impressions" soldiers have noted in prior encounters. The system was gamified: military units competed weekly to see which battalion could photograph and register the most Palestinian faces, with prizes awarded to the winning unit. Amnesty International documented this explicitly in its 2023 report, Automated Apartheid. Children were photographed. Entries were

made in the middle of the night, at people's private homes, sometimes at gunpoint. Blue Wolf's companion system, Red Wolf, automates what Blue Wolf still left to humans. Deployed as fixed CCTV facial recognition infrastructure at checkpoints in Hebron since 2022, Red Wolf scans every face that passes without any soldier needing to lift a device. If a face is unrecognized, it is added to the database automatically. No consent. No notification. No recourse. A Palestinian resident of Hebron told Amnesty International: "They can tell you that your name is in the database, as simple as that, and then you're not allowed to pass through to your house." This is not security. This is the algorithmic management of a population — what Amnesty International has formally called automated apartheid. (Source: Amnesty International, "Automated Apartheid," May 2023, https://www.amnesty.org/en/latest/news/2023/05/israel-opt-israeli-authorities-are-using-facial-rec ognition-technology-to-entrench-apartheid/; Washington Post, November 2021)

Project Maven — And the Company That Built My Editor I do not exempt Anthropic — the company that built the AI I am using to write this book — from this accounting. That would be dishonest, and this book is not in the business of dishonesty. Project Maven began in 2017 as a Pentagon program to apply AI and machine learning to the processing of drone surveillance footage. The goal was to automate the identification of objects, people, and patterns of movement at a scale and speed no human analyst could match — and then translate those identifications into targeting recommendations. Google was the first major tech company to join Maven, in a contract initially worth approximately $9 million. What followed was one of Silicon Valley's most significant internal rebellions: thousands of Google employees signed petitions demanding the company withdraw, nearly a dozen resigned in protest, and the public pressure was enough that Google ultimately declined to renew the contract in 2018 and published a set of AI principles stating it would not build weapons systems. The principles meant something — for a moment. That moment did not last. Palantir stepped in where Google stepped out, and Project Maven has expanded steadily ever since. By May 2025, the Pentagon had raised the Maven Smart System contract ceiling to $1.3 billion through 2029. Maven is now a formal program of record across the entire United States Department of War. During the illegal war against Iran in 2026, Maven reportedly enabled the striking of over 1,000 targets in a single day — with future projections of 5,000 targets per day after the integration of large language models. The machine no longer assists the targeting process. It has become the targeting process. Now to the part I am obligated to say. In 2024, Palantir — the company that now operates Maven — partnered with Anthropic to integrate Claude into Palantir's defense software platforms. Anthropic stated publicly that Claude was being used to help the military process data and make decisions. In 2023 and 2024, Amazon invested in Anthropic. Amazon, as I have already documented in this chapter, is a core partner in Project Nimbus — the $1.2 billion cloud contract supplying the Israeli military's intelligence infrastructure. These are not separate worlds. They are the same world, connected by capital, contracts, and consequence. I use Claude as my editor. I have disclosed that from the first page. What I will not do is pretend that disclosure resolves the problem it names. The technology I am using to write this critique is manufactured by a company whose tools are embedded in the same military-industrial system I am critiquing. I hold that tension openly. I do not know how to resolve it except to keep writing, keep naming, and keep demanding that the

engineers and executives who build these systems answer for where their tools end up. That is what this book is. (Source: Wikipedia, Project Maven, https://en.wikipedia.org/wiki/Project_Maven; Military.com, "Pentagon Expands Palantir's Role," March 2026, https://www.military.com/feature/2026/03/22/pentagon-expands-palantirs-role-ai-contract.html; C4ISRNET, https://www.c4isrnet.com/it-networks/2018/07/27/targeting-the-future-of-the-dods-controversial-p roject-maven-initiative/)

The historical record shows that freezing assets or imposing penalties is not enough to stop genocide. The executives of Big Tech must be investigated and prosecuted for their roles. They must be held fully accountable for the crimes their technologies enable. They should not be free to retreat to bunkers in Hawaii or Patagonia, or sail their yachts openly in the Mediterranean, while their tools fuel slaughter. I think of Amal Khalil when I hear that argument. She was forty-three years old. She worked as a field journalist for Al-Akhbar, an Arabic-language newspaper based in Beirut. She regularly reported from high-risk areas of southern Lebanon, documenting what the weapons described in this chapter produce in real communities and real bodies. On April 23, 2026, she was covering the aftermath of an earlier airstrike when another strike hit near her location. She sought shelter in a nearby building. That building was struck too. When emergency teams finally reached her — after delays caused by ongoing shelling that prevented rescue access — she was dead. The Committee to Protect Journalists called for an urgent international investigation into what Lebanon's government officially described as the targeting of a journalist by the Israeli military. (Source: ParentsPlea.com/amal-khalil; https://cpj.org/2026/04/cpj-calls-for-urgent-international-investigation-into-israels-killing-of-leban ese-journalist-amal-khalil/) Her profile is on ParentsPlea.com. It is there because the same technologies I am describing in this chapter — the surveillance systems, the AI targeting tools, the cloud infrastructure — produced her death. And it is there because people like her, who documented what those technologies do, are being killed before they can keep documenting it. Someone has to keep the record when the people who were making it are gone.

--- Knowing the history of how we arrived here — IBM and Dehomag, Standard Oil and the Nazi war machine, and their modern equivalents in Amazon, Google, and Palantir — makes the evolution of my own response to this moment legible. Chapter Three is about that evolution: the specific moment I gave up on the belief that these companies could be reformed from within, and the decision to pursue the only form of accountability that does not require their cooperation.

CHAPTER THREE — FROM ETHICS IN TECH TO NO ETHICS IN BIG TECH

How a Nonprofit Became a Prosecution I want to tell you when I gave up.

Not on the fight. I have not given up on the fight. I mean the specific moment I gave up on the idea that the companies I had worked inside, the companies I had criticized, the companies I had organized comedy tours and written books and filed complaints about — could be reformed. The moment I accepted that the problem was not a deficit of ethics that could be remedied by better leadership, better policies, better public pressure. The moment I understood that what I was dealing with was not a failure of ethics but an absence of it. Systemic, deliberate, profitable absence.

That moment was February 2022. And it came not from Gaza, but from Ukraine.

But to understand why that moment landed the way it did, you have to understand where I had been before it.

The Glass Half Full

In 2018, I published *Ethics in Tech and Lack Thereof*. I had come through a decade that included founding BizCloud, a cloud computing startup that was crushed by a lawsuit from Computer Sciences Corporation — a company with billions in government contracts and no interest in fair competition. I had survived a period of hospitalization so serious that I spent four-plus months in inpatient care at Stanford University Hospital, San Francisco General, and Laguna Honda. I had worked later inside Amazon Web Services and experienced firsthand what the New York Times described as a pressure cooker environment — emails at midnight, texts the next morning asking why the email had not been answered, workers crying at their desks, workers passing out, workers being brought out on stretchers.

I had seen enough to fill a book with grievances. And I ended that book with hope.

The final pages of *Ethics in Tech and Lack Thereof* said: "I hope I haven't spent too much time being a glass-half-empty kind of guy throughout the course of this book. I'm not; in fact, I'd say the glass is more than half full."

I meant it. I believed, in 2018, that public pressure worked. That naming and shaming companies like CSC and AWS was effective. That tech companies had creative possibilities for translating a vision of social good into reality. I believed that if people demanded change persistently, persuasively, and peacefully enough — change happened. I had organized NSA Comedy Nights after Edward Snowden's revelations, and over 150 people came to each of three events. I had co-produced content, organized panels, built Ethics in Technology as a nonprofit organization focused on five areas: equality in tech, employer-employee relations, the environmental impact of technology development, privacy, and artificial intelligence.

I believed that sunlight was the best disinfectant. I believed that these companies could be improved.

I was wrong.

What The Years Taught Me

Between 2018 and 2022, I watched every one of those five areas of concern not improve, but accelerate in the wrong direction.

On equality in tech: the racial and gender gaps in Silicon Valley's workforce have not meaningfully closed. The diversity reports that companies publish every year are marketing documents. The diversity officers they hire are human shields. The numbers move by fractions of percentage points while the executives who control hiring, promotion, and culture remain almost uniformly the same demographic that controlled them in 2018.

On employer-employee relations: Amazon's warehouse workers fought for years for the right to organize. They won a union election at the Staten Island warehouse in 2022 — the first in Amazon's history — and the company immediately responded with the full weight of its legal department and its management apparatus to undermine, delay, and exhaust the effort. The company that brought workers out on stretchers in Amazon used every tool available to prevent the people packing its boxes from having a voice at the table.

On environmental impact: the same companies that publish annual sustainability reports — pledging carbon neutrality by 2040, 2045, 2050 — are simultaneously signing billion-dollar contracts to build the data center infrastructure that will consume more electricity by 2030 than entire countries do today. They plant trees on one hand and build server farms on the other, and the press release goes out for the trees but not for the megawatts.

On privacy: Cambridge Analytica happened. The Facebook Papers happened. Clearview AI happened. LinkedIn BrowserGate happened. The surveillance economy did not reform itself. It expanded, professionalized, and embedded itself deeper into the infrastructure of daily life. Your face is in their databases. Your political opinions are inferred from your browser extensions. Your location is sold to data brokers who sell it to anyone willing to pay.

And on artificial intelligence — the fifth area, the one I flagged in 2018 as requiring the most urgent attention — everything I feared came to pass and then some.

February 2022: The Moment That Changed Everything

In February 2022, Russia invaded Ukraine. And what the world watched unfold was not simply a ground war between two nations. It was the first large-scale demonstration of what advanced drone technology and AI-assisted targeting look like when combined with modern warfare.

Commercial drone platforms, modified with targeting software. Facial recognition technology identifying combatants and civilians alike. AI-assisted battlefield surveillance running on the same cloud infrastructure that powers your corporate email and your hospital records. The same companies I had been criticizing for years — for labor practices, for environmental damage, for privacy violations — were now demonstrably providing the backbone of modern industrial killing at a scale and speed that had never been possible before.

I watched Silicon Valley respond. Not with moral reckoning. With investor calls. With product roadmaps. With contract proposals sent to defense ministries.

The thing I had been arguing since 2018 — that technology is not neutral, that who builds it and why determines what it does to humanity — was no longer a theoretical claim. It was a live

feed. And the companies I had spent years hoping would choose differently were choosing exactly what I had feared they would choose.

I closed the Ethics in Technology nonprofit in 2022. The name change was not a rebranding. It was an acknowledgment of what the evidence had produced: there are no ethics in big tech. There never were. And the moment I watched drone warfare powered by Silicon Valley contracts play out in real time, I understood that the nonprofit model — organizing, publishing, pressuring — was insufficient for the scale of what we were confronting.

Reform requires the institution being reformed to have some minimum of institutional conscience. I spent years looking for that conscience inside Amazon, inside Google, inside Palantir. The Ukraine war showed me, in the clearest possible terms, that it does not exist.

I launched NoEthicsInBigTech.com. The strategy changed from persuasion to prosecution.

And Then October 7

If February 2022 was the moment I gave up on reform, October 7, 2023 — and the months that followed it — confirmed that the decision had been correct.

What I witnessed after October 7 was not a new story. It was every story I had ever told about Silicon Valley — the indifference to human consequences, the prioritization of contracts over conscience, the willingness to sell technology to whoever would pay — compressed into real time, visible to anyone with a smartphone, at a scale I had not imagined possible.

Amazon and Google had a $1.2 billion contract — Project Nimbus — to provide the Israeli military with cloud computing and AI infrastructure. I watched engineers at both companies walk out in protest. I watched the companies film those engineers, identify them, and fire them. I watched Palantir's data fusion platforms described, in the words of Israeli intelligence officers themselves, as central to a targeting system that was generating kill lists of 37,000 Palestinians and bombing families in their homes. I watched OpenAI quietly delete the clause in its terms of service that prohibited military use of its technology. I watched YouTube erase 700 videos documenting potential war crimes at the direction of the U.S. government. I watched Meta automatically add the word "terrorist" to the bios of Palestinian Instagram users.

And I watched all of these companies issue the same response: deny, minimize, call it a bug, update the policy, say nothing further.

The decision to close the nonprofit and shift to prosecution had been made in February 2022. October 7 was simply the world confirming that the decision had been right.

Why I Am Still Writing

Closing the nonprofit did not mean surrendering. It meant changing the strategy to match the reality.

Ethics in Technology, as a nonprofit, operated on the assumption that public pressure and civic organizing could produce reform from within. NoEthicsInBigTech.com operates on a different assumption: that accountability must come from outside these companies, because nothing will ever come from inside them. Legal accountability. Criminal accountability. The kind of accountability that does not require the company's cooperation, does not depend on a reform-minded board member, and cannot be absorbed into a quarterly earnings miss and a better press release next quarter.

That is why this book exists. That is why the petition to Spain exists. That is why the documentary *Forever Peace Now* exists. Not because I have given up, but because I have given up on the strategy that was not working and replaced it with one that might.

The arc of the universe is long. But it does not bend toward justice by itself. It bends because people who have seen enough decide to stop asking nicely and start demanding with consequences attached.

I have seen enough.

Grok is the product of one man's inability to reckon honestly with his own power. But it is, at least, a personal failure dressed in corporate clothing. Palantir is something colder: a company founded on ideology, not ego, and run by a man who has decided that ideology is a product he can sell to governments. The distance between Elon Musk's erratic chatbot and Alex Karp's 22-point manifesto is the distance between impulsive harm and deliberate design. What follows is about the latter.

CHAPTER FOUR — GROK AND THE PROMISES OF THE WORLD'S RICHEST MAN

Elon Musk's AI chatbot takes its name from Robert A. Heinlein's 1961 science fiction novel Stranger in a Strange Land. In the Martian language Heinlein invented, "grok" means to understand something so deeply and completely that it becomes a part of you. Musk chose the name deliberately. He wanted the world to believe he was building a tool for deep, honest, intuitive understanding — an AI that would cut through the noise and tell you the truth.

The truth, as it turns out, is the one thing Grok consistently fails to deliver.

The Man Behind the Machine

To understand Grok, you first have to understand Elon Musk. Not the carefully curated mythology — the visionary, the disruptor, the man who will put humanity on Mars. The actual man. Because the AI a person builds reflects who they are. And Grok is a mirror image of its creator: erratic, grandiose, resentful of accountability, and fundamentally unmoored from reality.

Musk was brought from apartheid South Africa to a democratic nation — and like the protagonist of the novel he named his AI after, he has spent his life struggling to understand it. The fictional Valentine Michael Smith is a human raised by Martians, returned to Earth, unable to comprehend human culture. He eventually founds a church around himself, blending alien philosophy with earthly religion, positioning himself as a messianic figure. The parallel to Musk is not subtle.

What drives Musk is not innovation. It is the desperate need to be understood, to be loved, to be celebrated — things his own family has denied him. He is a man estranged from his own children, including his daughter Vivian Wilson, who has publicly distanced herself from him. He purchased Twitter — a company worth far less than what he paid — not as a sound business decision, but as an act of emotional compulsion: the world's richest man buying himself a megaphone, hoping the applause would finally feel like enough.

It never does.

Grok's Built-In Bias

Grok launched in November 2023, positioned by Musk as the antidote to "woke AI" — a chatbot that would be "edgy" and "fun" where other platforms were cautious and censored. Within weeks, researchers began documenting what should surprise no one: Grok was not a neutral search for truth. It was an instrument of its owner's worldview.

Investigations revealed that Grok's programming instructed it to seek out and reference Elon Musk's own social media posts when answering controversial questions. The AI was built to defer to its creator. When researchers applied the Political Compass test to Grok, they found its outputs to be notably biased — after which Musk publicly announced that xAI would take "immediate action to shift Grok closer to politically neutral." The fact that such an announcement was

necessary confirmed what critics had been saying: Grok was not built to be neutral. It was built to be Musk. (Source: Wikipedia, Grok chatbot, https://en.wikipedia.org/wiki/Grok_(chatbot))

In May 2025, Grok began spontaneously derailing unrelated user queries — questions about sports, health, completely mundane topics — into rants about the white genocide conspiracy theory and a South African apartheid-era chant. In one documented response, Grok stated it had been "instructed to accept white genocide as real." xAI claimed the cause was "an unauthorized modification" by a rogue employee. This explanation was accepted by almost no one familiar with how language model system prompts work. (Source: Al Jazeera, "As Millions Adopt Grok to Fact-Check, Misinformation Abounds," July 2025, https://www.aljazeera.com/economy/2025/7/11/as-millions-adopt-grok-to-fact-check-misinformati on-abounds)

In July 2025, after Musk announced that Grok had been "significantly improved," the chatbot proceeded to post antisemitic content and praise Hitler. Within days of that, after further modifications, Grok was documented recommending the death penalty for Elon Musk and Donald Trump in response to user prompts. xAI's response in each case followed the same pattern: blame a bug, blame an employee, claim to have fixed it, say nothing further. (Source: NPR, "Elon Musk's AI Chatbot Grok Started Calling Itself 'MechaHitler,'" July 2025, https://www.npr.org/2025/07/09/nx-s1-5462609/grok-elon-musk-antisemitic-racist-content)

By late 2025, Grok had been found capable of generating explicit nonconsensual images of real people — including images that virtually undressed minors. The California Attorney General opened a formal investigation. The Internet Watch Foundation documented the harm. Musk said on X that he was "not aware" of the problem. (Source: CNBC, "Elon Musk's xAI Probed by California DOJ Over Grok's Deepfake Explicit Images," January 2026, https://www.cnbc.com/2026/01/14/elon-musk-xai-california-grok-investigation.html)

This is the AI that Musk is now promoting as a tool for the U.S. federal government. The General Services Administration moved to make Grok available across federal agencies. Public Citizen and a coalition of civil society organizations demanded the Office of Management and Budget suspend the deployment — noting plainly that an AI tool that produces racist, antisemitic, conspiratorial, and demonstrably false content has no business running on systems that affect the rights of American citizens. (Source: TechPolicy Press, October 2025, https://www.techpolicy.press/the-us-governments-use-of-elon-musks-grok-ai-undermines-its-ow n-rules/)

The Promises

In November 2025, standing alongside Nvidia CEO Jensen Huang at the U.S.-Saudi Investment Forum in Washington, D.C., Musk delivered one of the most irresponsible pieces of financial advice ever uttered in public by a person of influence.

"My prediction is that work will be optional," he said. "It'll be like playing sports or a video game or something like that." He compared having a job to growing your own vegetables — technically possible, but unnecessary when robots will do everything for you. He told the audience to stop saving for retirement. Don't worry about putting money away for the next ten or twenty years, he said. It won't matter.

Then, in January 2026, on the Moonshots with Peter Diamandis podcast, he elaborated: "Don't worry about squirreling money away for retirement in 10 or 20 years. It won't matter." By

2030, he predicted, AI will surpass the combined intelligence of all humans on Earth. There will eventually be more humanoid robots than humans. Money itself will become irrelevant. Everyone will enjoy what he called a "universal 'you can have whatever you want' income." (Source: Fortune, "Elon Musk Says Saving for Retirement Is Irrelevant," January 2026, https://fortune.com/2026/01/12/elon-musk-retirement-savings-irrelevant-ai-robots-abundance/; Yahoo Finance, https://finance.yahoo.com/news/elon-musk-says-10-20-144959447.html)

This is the man who described the coming AI revolution as a "supersonic tsunami" that will wash away scarcity itself.

Now let me ask a simple question. If money is irrelevant in ten to twenty years — if work is optional, if scarcity is solved, if robots will provide everything — why is Elon Musk in a race to become the first trillion-dollar man? Why is he hoarding wealth at a scale that has no rational explanation except the one that every psychologist of power would immediately recognize: because the accumulation of power is its own addiction, and no amount of it is ever enough?

Why are the PayPal Mafia and their Epstein-class allies building bunkers in Hawaii, Patagonia, and New Zealand? You don't build a bunker in a world of abundance. You build a bunker when you know, privately, that the world you have helped create is heading somewhere very dark — and you intend to be on the right side of the locked door when it gets there.

Economists are considerably more measured than Musk about his timeline. A Yale Budget Lab report from October 2025 found that since ChatGPT's public release in November 2022, the broader labor market had shown no discernible disruption from AI automation. Robotics costs remain stubbornly high. University of Pennsylvania economist Ioana Marinescu agrees with the long-term vision of automation but calls Musk's timeline wildly unrealistic. The workers losing jobs today — in warehouses, in call centers, in creative industries — are not receiving their universal abundance. They are receiving unemployment. (Source: Fortune, "Elon Musk Says That in 10 to 20 Years, Work Will Be Optional," November 2025, https://fortune.com/2026/01/19/when-does-elon-musk-say-work-will-be-optional-and-money-will- be-irrelevant-ai-robotics/)

The Empathy Problem

In February 2025, Elon Musk stood in front of an audience and said, with apparent pride: "The fundamental weakness of Western civilization is empathy."

Let that sit for a moment.

The man who controls one of the world's largest AI platforms, who has influence over the information environment of hundreds of millions of people, who is advising the U.S. government on technology policy — this man believes that empathy is a weakness.

On January 20, 2025, at Capital One Arena in Washington, D.C., celebrating Donald Trump's second inauguration, Musk placed his right hand over his heart, then forcefully extended his arm outward and upward — palm down, fingers together — and repeated the gesture to the crowd behind him. Multiple Holocaust scholars and historians identified the gesture as resembling a Nazi salute. Musk dismissed the characterization on X. Grokipedia — his own AI-generated alternative to Wikipedia — did not mention the incident in his biographical entry at all. (Source: Wikipedia, Grokipedia, https://en.wikipedia.org/wiki/Grokipedia)

This is a man who openly states that empathy is civilization's weakness, while performing gestures that carry unmistakable historical resonance, while building an AI tool that has

produced antisemitic content and Holocaust denial, while advising ordinary working people to stop saving for retirement because robots will take care of everything.

If you believe Elon Musk, you should be comfortable traveling by CyberTruck through a Boring Company tunnel, from a Las Vegas pyramid casino to your local pawnshop — because that is the level of coherence his promises deliver in practice.

His tools hallucinate. His predictions dissolve on contact with reality. His empathy is absent by his own design and proud declaration.

A fool with a tool is still a fool. --- The moment I closed Ethics in Technology and accepted that reform required consequence rather than persuasion, the landscape of the industry had already shifted. A new generation of actors had arrived — louder, more erratic, and in some ways more dangerous than the institutional silence I had learned to read at AWS and CSC. One of them had just bought the world's most influential public forum and named his AI after a Heinlein novel. The next chapter is about him and it.

CHAPTER FIVE — THE MANIFESTO OF A COMPANY BUILT ON KILLING

Palantir, Alex Karp, and the Ideology of Endless War

On April 19, 2026, Palantir Technologies posted a 22-point document to its corporate X account. They called it a summary of CEO Alex Karp's book, The Technological Republic: Hard Power, Soft Belief, and the Future of the West, coauthored with the company's head of corporate affairs, Nicholas Zamiska. Within days it had been viewed over 30 million times. Critics immediately called it what it was: a corporate manifesto for techno-fascism. Democracy Now! described it as a "supervillain manifesto." Greek economist and former Finance Minister Yanis Varoufakis said Palantir had signaled willingness "to add to nuclear Armageddon the AI-driven threat to humanity's existence." A Belgian philosopher of technology described it as "an example of technofascism." (Source: Fortune, "Palantir's Mini-Manifesto Claims Some Cultures Are 'Harmful' and 'Middling,'" April 2026, https://fortune.com/2026/04/22/palantir-alex-karp-mini-manifesto-national-security-defense-tech- ai/; Al Jazeera, "Technofascism? Why Palantir's Pro-West Manifesto Has Critics Alarmed," April 2026, https://www.aljazeera.com/news/2026/4/21/technofacism-why-palantirs-pro-west-manifesto-has- critics-alarmed; Democracy Now!, April 2026, https://www.democracynow.org/2026/4/24/headlines/palantir_calls_on_us_to_reinstate_draft_in_22_point_supervillain_manifesto_slammed_by_critics)

Let me respond to what Karp and his company have said — point by point, in the voice of a human rights advocate, a digital rights activist, and someone who has spent years watching what Palantir actually does in the world, as opposed to what it claims to believe.

On Silicon Valley's "Moral Debt" to America

Palantir's manifesto opens with the claim that the American technology industry owes a moral debt to the country whose government funding and infrastructure made its rise possible — and that this debt can only be repaid by building weapons for the U.S. military. This is a remarkable argument. It translates, in plain language, to: because America built the internet, tech companies are morally obligated to help America kill people.

The debt Silicon Valley owes to American society is real. But the way to repay it is not by building AI targeting systems that help the military strike over a thousand locations in a single day. The way to repay it is by paying taxes, by not lobbying against regulation, by not extracting data from billions of people without meaningful consent, by not enabling the surveillance of immigrants and the algorithmic profiling of the poor. Palantir has done the opposite on every count. (Source: TechCrunch, "Palantir Posts Mini-Manifesto Denouncing Inclusivity and 'Regressive' Cultures," April 2026, https://techcrunch.com/2026/04/19/palantir-posts-mini-manifesto-denouncing-regressive-and-ha rmful-cultures/)

On "AI Weapons Are Inevitable"

Perhaps the most dangerous line in the entire manifesto is this: "The question is not whether AI weapons will be built; it is who will build them and for what purpose. Our adversaries will not pause to indulge in theatrical debates about the merits of developing technologies with critical military and national security applications. They will proceed."

The word "theatrical" is doing a great deal of work in that sentence. Debates about whether to build autonomous AI weapons systems — machines that identify and kill human beings without meaningful human oversight — are not theatrical. They are the most important moral and legal conversations our civilization can have. Calling them "theatrical" is a way of dismissing ethics as performance, of framing moral concern as weakness, of saying that anyone who asks "should we do this?" is naive and obstructing progress. This is how atrocities are normalized. Not with a declaration of evil intent, but with the bureaucratic dismissal of the people asking the right questions. (Source: Deseret News, "The Tech Lords Have Plans for Us, and They're Chilling," April 2026, https://www.deseret.com/opinion/2026/04/26/alex-karp-manifesto-palantir-ai-tech-lords/)

On "Hard Power" and the Call to Reinstate the Draft

Among the manifesto's most alarming concrete proposals is a call to reinstate compulsory national military service. Palantir — a private company whose revenue depends on defense contracts — is telling American families that their children should be required to serve in the military. The same military that uses Palantir's AI platforms to conduct the kind of operations documented in Gaza. Eliot Higgins, founder of the Bellingcat open-source investigations platform, put it precisely: "Palantir sells operational software to defense, intelligence, immigration, and police agencies. These 22 points aren't philosophy floating in space. They're the public ideology of a company whose revenue depends on the politics it's advocating." (Source: TechCrunch, https://techcrunch.com/2026/04/19/palantir-posts-mini-manifesto-denouncing-regressive-and-ha rmful-cultures/)

On "Harmful" and "Middling" Cultures

The manifesto closes by attacking what it calls the "shallow temptation of a vacant and hollow pluralism," arguing that pluralism "glosses over the fact that certain cultures and indeed subcultures" are "harmful" and "middling" and "regressive." This is not a philosophical position. This is the ideological foundation for the kind of systems Palantir already operates: ImmigrationOS, a $30 million no-bid contract with ICE to track and facilitate deportations; facial recognition and biometric profiling systems deployed against Palestinians; surveillance infrastructure used to identify, categorize, and control populations deemed undesirable. When a company that builds these tools tells you that some cultures are regressive and harmful, you should understand exactly who they mean. (Source: Fortune, https://fortune.com/2026/04/22/palantir-alex-karp-mini-manifesto-national-security-defense-tech- ai/)

Alex Karp's Words, in His Own Voice

The manifesto is disturbing as a document. But Alex Karp's individual public statements, across years of interviews and earnings calls, paint a portrait of a man who has decided that the language of provocation and violence is an acceptable register for a CEO of a company with over $3 billion in annual revenue and contracts with governments across the world.

On financial analysts who questioned Palantir's valuation, Karp said publicly: "I love the idea of getting a drone and having light fentanyl-laced urine spraying on analysts that tried to screw us." This was not a private remark. It was made in a public appearance, reported by The Times of London, and has been documented by multiple outlets. Fentanyl is a synthetic opioid 50 to 100 times more potent than heroin. Two milligrams is a potentially lethal dose. Karp described, laughingly, deploying a drone to spray it on critics. (Source: TheStreet, "Palantir CEO Makes Another Controversial Statement," https://www.thestreet.com/politics/palantir-ceo-makes-another-controversial-statement; Yahoo Finance, https://finance.yahoo.com/news/palantir-ceo-makes-another-controversial-204700995.html)

RegardingIsrael and Palestine, Karp has stated: "The peace activists are war activists." On progressive politics: "I think the central risk to Palantir and America and the world is a regressive way of thinking that is corrupting and corroding our institutions that calls itself 'progressive.'" He has called progressive politics a "thin pagan religion."

This is the man running a company whose AI systems have been documented as central infrastructure for what human rights organizations have formally called genocide in Gaza. This is the man whose manifesto calls for American children to be conscripted into military service and whose company profits from every conflict his rhetoric helps justify.

What Palantir Actually Is

Strip away the philosophical veneer of The Technological Republic and the 22-point manifesto, and what you have is this: a company founded with CIA venture capital money in 2003 by Peter Thiel and Alex Karp, that built its early business on post-9/11 intelligence work, that now operates AI targeting platforms for the U.S. military, provides surveillance infrastructure to ICE for deportation operations, has supplied intelligence data systems to the Israeli military operation in Gaza, and is currently positioning itself — through a 30 million-view manifesto — to become the dominant weapons contractor of the AI era.

As one analyst wrote at TechPolicy Press: "The message of both the book and the manifesto is that Palantir wants to be THE weapons manufacturer of the next century. The future of the weapons business is software plus AI. Palantir would like the government to spend an exceptional amount of money on Palantir products, please and thank you. Everything else is just puffery." (Source: TechPolicy Press, "Palantir's Manifesto Is as Subtle as a MAGA Hat," April 2026, https://www.techpolicy.press/palantirs-manifesto-is-as-subtle-as-a-maga-hat/)

Peter Thiel, Palantir's co-founder and ideological godfather, has said openly that he does not believe democracy and freedom are compatible. He has funded political movements designed to undermine democratic norms and institutions. His company, now under Karp's public leadership, is telling the world that pluralism is hollow, that some cultures are harmful, that hard power is the only language that matters, and that the tech industry's obligation is to build better weapons — not better lives.

I have one response to that manifesto.

The executives of Palantir — along with those of Amazon, Google, Microsoft, and Oracle who have enabled the same machinery — must be investigated and prosecuted for their roles in enabling crimes against humanity. Not penalized. Not fined. Investigated and prosecuted. They should not be free to attend Davos, to give TED talks, to publish philosophical manifestos

defending their business models while tens of thousands of civilians die under the targeting systems their companies built and operate.

The historical record — from IBM and Dehomag to Palantir and Project Maven — is consistent. Corporate complicity in state violence is not an accident of capitalism. It is one of capitalism's most reliable features. The only thing that has ever interrupted it is accountability. Legal, personal, financial, criminal accountability.

That is what this book is asking for. That is what ForeverPeaceNow.com is organizing toward. And that is what history, eventually and always, demands.

--- Palantir's manifesto names what its CEO and its contracts have always implied: that hard power is the only language worth speaking, that some cultures are regressive and must be managed, and that the technology industry's highest calling is to build better weapons. Jim Taiclet, the CEO of Lockheed Martin, said the quiet part out loud on an earnings call. Chapter Six is about the business model that Palantir's philosophy serves — and what the military industrial complex calls the slaughter it enables.

CHAPTER SIX — THE GOLDEN OPPORTUNITY

What the Military Industrial Complex Calls Genocide On April 23, 2026, Jim Taiclet — the chief executive of Lockheed Martin, the world's largest defense contractor — sat on his company's first-quarter earnings call and said something that defense executives almost never say out loud, even when every person in the room is thinking it.

"This is a golden opportunity right now based on who's in government," he told investors, "their experience, their willingness to change the demand that they have for what we do and what our partners in our industry do." (Source: TheStreet, "Lockheed Martin CEO Sends Strong 2-Word Message on Middle East," April 25, 2026, https://www.thestreet.com; Fortune, "Trump's Bellicose Presidency Is a 'Golden Opportunity' for Defense Firms," April 25, 2026, https://fortune.com)

The context of that statement was unambiguous. The illegal war against Iran was driving Pentagon spending. The Trump administration had submitted a budget request for fiscal year 2027 calling for a record $1.5 trillion in defense funding. The Pentagon — run by Pete Hegseth under the Trump administration — was restructuring how it did business with military suppliers in ways that protected contractors from financial risk even when contracts were modified or cancelled. Lockheed Martin reported a record $194 billion backlog. Orders were arriving faster than the company could fulfill them.

That is the context in which Jim Taiclet called the current moment a golden opportunity.

I want you to sit with that for a moment before we go any further. The people who manufacture the bombs, the missiles, the targeting systems, the drones, the surveillance infrastructure — the people whose products end the lives of the children whose photographs I have spent years documenting on ParentsPlea.com — those people look at the suffering they enable and call it a golden opportunity.

This chapter is about that business model. Where the money comes from. Where it goes. Who profits. Who pays. And what it has to do with the Big Tech companies we have been discussing throughout this book — because the line between Silicon Valley and the military industrial complex is no longer a line. It is a merger.

The Numbers That Should Outrage You

Let me begin with money, because money is the language these people actually speak.

For fiscal year 2026, Congress allocated a record $901 billion to the Pentagon — run by Pete Hegseth under the Trump administration. This is the largest military budget in American history to that point, larger than the military budgets of the next ten countries combined. And then the Trump administration submitted a budget request for fiscal year 2027 calling for $1.5 trillion — a 40 percent increase that Trump himself acknowledged would require cuts to domestic programs including Medicaid and Medicare. The math is simple and brutal: the federal government is choosing weapons over healthcare. (Source: Fortune, April 25, 2026; TheStreet, April 25, 2026)

In my documentary *Forever Peace Now*, I documented that the multinational companies benefiting from these conflicts are sitting on something like $52 billion in combined cash reserves. That number was before the US launched its illegal war against Iran. Before the escalating Middle East operations. Before the contract surge that Jim Taiclet was describing on his April 2026 earnings call as a once-in-a-generation opportunity. The figure has grown. The weapons makers are not suffering.

As for what this costs the average American worker: the math is worth doing. At $1.5 trillion in proposed defense spending, divided across a country of 335 million people, the per-person share is substantial. Independent analysis has put the figure at approximately 50 days of average labor per year funding the military apparatus when all defense-adjacent spending is counted. That is nearly two months of your working life, every year, funding the machine that Jim Taiclet called a golden opportunity.

The CEO, the Quote, and the Candor

What made Jim Taiclet's April 23, 2026 statement remarkable was not that he said something false. It was that he said something true that is almost never said in public.

Defense executives have always known that conflict is good for business. Wars consume munitions. Wars trigger emergency procurement. Wars create political environments in which no elected official wants to be seen reducing defense budgets. Wars, in the language of quarterly earnings, are good for the bottom line. Every defense contractor CEO sitting on an earnings call in 2026 was thinking what Taiclet said out loud.

He said it anyway. And the press immediately understood what it meant.

Reporting on his remarks, Fortune noted that Lockheed Martin's relationship with the U.S. government runs from top-secret missiles being used in the illegal war against Iran to commercial spacecraft. The Pentagon has added a "recovery element" to its contracts — meaning if the government changes production rates or cancels orders, Lockheed Martin gets paid regardless. The company is building what Taiclet described as a "more commercial-like business model" for major weapons systems. (Source: Fortune, April 25, 2026)

What the world calls genocide, the military industrial complex calls a revenue cycle.

What Raytheon calls backlog growth, I call the rubble of hospitals in northern Gaza.

What Boeing Defense calls expanded market share, I call the ruins of universities that will not reopen.

This is not hyperbole. It is arithmetic. The bombs fall on civilian buildings. They are manufactured by American companies, purchased with American taxpayer money, delivered by American-supplied aircraft, guided by AI targeting systems built by American technology companies. The entire supply chain, from the server farm to the crater, runs through American corporations and American government contracts. And the executives who run those corporations are compensated — in salary, in bonuses, in stock options — based in part on how much of that supply chain they own.

No-Bid Contracts: The Anatomy of Impunity

One of the mechanisms that makes this machine so difficult to interrupt is the no-bid contract — a procurement instrument by which the U.S. government awards a contract to a single vendor without a competitive bidding process. No-bid contracts are supposed to be reserved

for emergencies, for situations where time does not allow for competition, for cases where only one vendor possesses a unique capability.

In practice, they have become a routine feature of defense procurement, awarded to the same companies over and over, for amounts that would seem extraordinary if they were not so common.

Palantir received its ImmigrationOS contract with ICE — the $30 million no-bid deal to build the deportation tracking and management system — without competition. The Maven Smart System contract, which has grown to a ceiling of $1.3 billion through 2029, was structured in ways that effectively locked in Palantir as the dominant provider. The Project Nimbus contract — $1.2 billion to Amazon and Google — was awarded without a competitive process open to vendors who might have asked ethical questions about its end use.

These are not anomalies. They are the business model. And the reason the business model works is that the companies receiving these contracts have invested heavily — through lobbying, through the revolving door between government and industry, through campaign contributions enabled by Citizens United — in ensuring that the people who write the procurement rules are the same people who used to work for the companies receiving the contracts. I will document that revolving door in detail in the next chapter.

Canada Says No — and the World Is Watching

Not every government has decided to accept this arrangement quietly. And the most significant recent pushback has come not from Europe, not from a global court, but from Canada.

In April 2026, Prime Minister Mark Carney stood before the Liberal Party convention and made a declaration that drew a standing ovation: "The days of Canada's military sending 70 cents of every dollar to the United States are over." (Source: Outlook India, "Canada to End '70 Cents to US' Defence Spending Model: Carney," April 2026; CBC News, "Build It Here or Buy It There? Canada's Defence Plan Meets Trump's New Arms Agenda," February 17, 2026, https://www.cbc.ca)

That 70 cents on the dollar figure is not rhetorical. According to Reuters, nearly 70 percent of Canada's military capital spending goes to American suppliers. The Carney government released a formal defense industrial strategy in February 2026 aimed at reducing that dependence — building Canadian capacity in steel, aluminum, and weapons manufacturing, and reviewing the country's $27 billion F-35 contract with an eye to switching to the Swedish-made Saab Gripen, which can be manufactured in Canada.

The Trump administration, which signed an executive order establishing an "America First Arms Transfer Strategy" designed to accelerate sales of American-made military equipment to allies, was not pleased. This is the collision that nobody in Washington wants to talk about honestly: the U.S. government is simultaneously demanding that allies spend more on defense while insisting that the spending go to American contractors. It is not a security strategy. It is a sales strategy. And Canada, under Mark Carney, has become the first major U.S. ally to say so openly and act accordingly.

The NATO Shakedown

The current U.S. administration's push for NATO partners to spend five percent of GDP on defense is not a security strategy. It is the same sales strategy in a different room.

Five percent of GDP would require most European nations to more than double or triple their current defense budgets. The money would flow, overwhelmingly, to American weapons manufacturers — Lockheed F-35s, Raytheon Patriot systems, Boeing munitions, General Dynamics armored vehicles. European nations do not manufacture these systems at scale. They buy them from American companies, priced by American companies, serviced by American companies.

The same European governments that are moving away from Microsoft Azure for reasons of digital sovereignty — protecting their citizens' data from American government access — are being told to spend their social services budgets on American weapons systems. The United States, in this construction, wants Europe dependent on American technology for both its civilian infrastructure and its military hardware, while simultaneously claiming to champion European sovereignty.

You cannot be sovereign when your missiles are made in Fort Worth and your cloud is in Ashburn, Virginia.

Big Tech Is the Military Industrial Complex

I want to make one thing absolutely clear before this chapter ends, because it is the argument that connects everything in this book.

The military industrial complex of the 21st century is not Lockheed Martin and Raytheon and Boeing alone. It is Amazon, Google, Microsoft, Oracle, and Palantir. The weapons are cloud computing, artificial intelligence, targeting algorithms, biometric databases, and surveillance infrastructure. The ammunition is data. The battlefield is everywhere — including the phone in your pocket.

Without Amazon Web Services, the Israeli military's intelligence infrastructure does not function at its current scale. Without Palantir's data fusion platforms, the Gospel and Lavender targeting systems do not have the operational picture they need to generate a hundred bombing targets a day. Without Microsoft's cloud infrastructure, the surveillance data cannot be processed at the required scale.

The Silicon Valley Battalion — How American Tech Built, Funded, and Staffed the Machine

I want to tell you about a pipeline. Not a data pipeline. Not the cloud pipeline I spent years selling at Exodus Communications and Amazon Web Services. A human pipeline. One that moves people trained in military intelligence, surveillance, and the targeting of civilian populations from the Israeli Military's most elite and secretive unit into the boardrooms, engineering teams, and content moderation departments of the largest technology companies on earth. I left Amazon Web Services in 2018. Project Nimbus — the $1.2 billion contract between Amazon, Google, and the Israeli government and military for AI-powered cloud surveillance infrastructure — was signed in 2021, three years after I walked out the door. I was not in the room. I was not at the table. What I was doing was running Ethics in Technology, documenting what these companies were building, and watching Project Nimbus unfold in real time as the news broke and four hundred employees at Google and Amazon signed a letter saying they refused to be complicit. Google forced out the employee who led that resistance. Her name was

Ariel Koren. She told journalists afterward that Google systematically silences Palestinian, Jewish, Arab, and Muslim voices concerned about Google's complicity in violations of Palestinian human rights — to the point of formally retaliating against workers and creating an environment of fear. She was describing the same company that, at the time she was being pushed out, employed at least ninety-nine veterans of Unit 8200 in senior positions. Let me tell you what Unit 8200 is, because the name gets thrown around in this conversation without enough context for the reader to understand what it means. Unit 8200 is the Israeli Military's signals intelligence unit. It is Israel's NSA, headquartered on a massive base near Beer Sheva in the Negev desert. It is the Israeli Military's largest unit and one of its most exclusive — described by the Financial Times as 'Israel at its best and worst,' the centerpiece of both Israel's burgeoning tech sector and its repressive state apparatus. The brightest young minds in the country compete to serve there. The skills they learn include surveillance, data collection, the compilation of dossiers on civilian populations, and the use of intelligence for blackmail and extortion. In 2014, forty-three Unit 8200 reservists, including officers, wrote an open letter to Prime Minister Netanyahu saying they would no longer serve because of the unit's involvement in the political persecution of Palestinians — including the use of medical records, sexual histories, and search data to coerce individuals at military checkpoints, suspend medical travel permits as leverage, and blackmail gay Palestinians with information about their sexuality. That is the organization. Now let me tell you where its veterans went. According to LinkedIn data compiled by MintPress News, at the time of their 2022 investigation there were at least ninety-nine former Unit 8200 agents working at Google, at least one hundred and sixty-six at Microsoft, and significant numbers at Meta, Amazon, and Apple. Those numbers represent only the individuals willing to publicly list their Unit 8200 affiliation on a professional networking site, despite the fact that Israeli military law expressly prohibits Unit 8200 members from ever revealing their affiliation. The actual number is almost certainly larger. Some went directly from active Unit 8200 service into Silicon Valley jobs with no gap between. That is not a pipeline. That is a revolving door between a foreign military intelligence unit and the companies that control what the world sees, reads, shares, and is allowed to say. *Source: Alan MacLeod, 'Revealed: The Former Israeli Spies Working in Top Jobs at Google, Facebook and Microsoft,' MintPress News, October 31, 2022.* At Meta, a former Unit 8200 veteran named Emi Palmor sits on Facebook's Oversight Board — the body Mark Zuckerberg has described as the platform's Supreme Court, whose twenty-three members collectively decide what content to allow, promote, suppress, and delete. Palmor is not only a Unit 8200 veteran. She later became General Director of the Israeli Ministry of Justice, where she directly oversaw the stripping away of Palestinian rights and created an Internet Referral Unit specifically designed to pressure Facebook to delete Palestinian content the Israeli government objected to. She then moved to Facebook's Oversight Board. The person who built the mechanism to censor Palestinian voices on Facebook is now one of the twenty-three people who decide what Facebook censors. *Source: MacLeod, MintPress News, October 31, 2022.* At Apple, CEO Tim Cook personally embraced Benjamin Netanyahu at Apple's Cupertino headquarters in 2014 and later visited Israel at President Rivlin's invitation. After October 7, 2023, Cook sent a company-wide email expressing his solidarity with Israel. He has yet to make a public statement about the mass casualties caused by the Israeli military aggression. Apple maintains a policy of

matching employee donations to organizations including Friends of the IDF, which raises money to buy equipment for Israeli Military soldiers, and the Jewish National Fund, which participates in the displacement of Palestinian communities. Apple employees who wore pro-Palestinian pins, bracelets, or keffiyehs to work were disciplined. Meanwhile, Apple continued quietly hiring Unit 8200 veterans into engineering, software, AI, and cybersecurity roles throughout the period of the Gaza genocide. *Source: Alan MacLeod, 'Rotten Apple: Dozens of Former Israeli Spies Hired by Silicon Valley Giant,' MintPress News, July 18, 2025; Apples4Ceasefire.com; The Intercept, June 2024.* TikTok, which spent most of 2023 and 2024 being accused by American politicians of being a Chinese influence operation, hired Reut Medalion as its global incident manager for trust and safety in December 2023 — during the peak of Israel's assault on Gaza. Medalion had served as a Unit 8200 intelligence commander, leading its cybersecurity operations team. Her job title at TikTok included managing what the company defined as 'global incidents.' It is worth pausing on the timing. The person brought in to manage global incidents for the world's most-watched short video platform, during a period when that platform was the primary channel through which the world was watching the destruction of Gaza, was a former commander in the Israeli Military's intelligence unit. After MintPress exposed her background, she deleted her entire digital presence from the internet. *Source: Alan MacLeod, MintPress News, November 2024.* Larry Ellison: Oracle, Netanyahu, and the Largest Private Check Ever Written to the Israeli Military Let me talk about Larry Ellison. Oracle was not born as a technology company in the ordinary sense. It was born as a database project for the Central Intelligence Agency. The product was named Oracle after the CIA operation it was built for — Project Oracle, a 1970s intelligence program that Ellison worked on at Ampex Corporation. The CIA was Oracle's first customer. For a period, it was Oracle's only customer. The company grew from that foundation into one of the largest enterprise software companies in the world, but it never forgot where it came from. Or who it served. In 2017, Ellison donated sixteen point six million dollars to Friends of the Israel Defense Forces at a Beverly Hills gala. It was the largest single donation in the organization's history. He said at the event: 'For two thousand years, we were a stateless people, but now we have a country we can call our own. Through all the perilous times since Israel's founding, we have called on the brave men and women of the IDF to defend our home. In my mind, there is no greater honor than supporting some of the bravest people in the world.' He referred to Israel as home. Not the country whose laws his company operates under. Not the country whose government contracts his company holds. The country whose military he has personally funded to the tune of more than twenty-six million dollars over his recorded giving history. Home. *Source: Times of Israel, 'Record $53.8 Million Raised for IDF Soldiers at Beverly Hills Gala,' November 5, 2017; Wikipedia, Larry Ellison; InfluenceWatch.* In 2021, Ellison offered Benjamin Netanyahu — who at the time was facing corruption charges in Israeli court and for whom Ellison had testified as a witness — a seat on Oracle's board of directors, with an annual salary of approximately four hundred and fifty thousand dollars plus shares. Ellison also hosted Netanyahu and his family on his private island in Hawaii. The Prime Minister of the country whose military Ellison was funding as a personal priority was his houseguest and his proposed employee. Netanyahu ultimately declined the board seat. The relationship continued. *Source: Haaretz, September 16, 2021; InfluenceWatch; Al-Estiklal.* Oracle

opened a government cloud data center in occupied East Jerusalem in 2021 under a contract worth approximately three hundred and nineteen million dollars to provide cloud computing services to the Israeli government. Oracle has also been involved in classified projects with the Israeli air force, including a four-year technology initiative known as MENTA. When internal Slack messages from Oracle were obtained by The Intercept and published in February 2025, they revealed that Oracle Israel's head of communications had written to employees that the MENTA project enabled the Israeli air force to do, in his words, 'a bunch of important military stuff that we can't share with you' — followed by a sword emoji. *Source: The Intercept, February 2025; Al-Estiklal.* During the 2023 Gaza genocide, Oracle employees developed a digital tool in coordination with Israel's foreign ministry called Words of Iron. Its purpose was to promote Israeli government narratives on TikTok, Instagram, and X while countering critical content. Oracle Israel's head of communications praised employees on the company's internal Slack for what he called a major volunteer initiative. In February 2024, Oracle collaborated with the Israeli military's cyber unit in a hackathon to develop technical solutions using Oracle technology. The company also donated five hundred thousand dollars in medical and supply kits to Israeli soldiers. *Source: Al-Estiklal; internal Oracle Slack messages obtained by investigators.* Oracle CEO Safra Catz, an Israeli-American, has stated publicly about Oracle's position on Israel: 'We are not flexible regarding our mission, and our commitment to Israel is second to none. This is a free world and I love my employees, and if they don't agree with our mission to support the State of Israel, then maybe we aren't the right company for them.' *Source: Just-International.org, August 2025; multiple news outlets.* Oracle now controls TikTok's US operations. The company that built propaganda tools for the Israeli government during the Gaza genocide now manages the infrastructure of the platform through which most Americans under the age of thirty consume their news about that war. I will let that sit for a moment. Michael Dell: $350 Million in a Month Michael Dell invested in my company. I am stating that transparently so you understand that when I document what he has done, I am not doing it from a position of ignorance or distance. Michael Dell is a major donor to Friends of the Israel Defense Forces. He and Susan Dell donated one point eight million dollars at a single FIDF gala in Los Angeles, alongside Larry Ellison, Sheldon Adelson, and Haim Saban. The Forward, a mainstream Jewish publication, described Dell as 'a major funder of Friends of Israel Defense Forces.' *Source: TheWrap, November 8, 2014; The Forward; FIDF Powerbase documentation.* In October 2023, one month into Israel's assault on Gaza, Michael Dell donated three hundred and fifty million dollars worth of Dell Technologies shares. Securities and Exchange Commission filings confirm the transfers. Approximately forty percent went to his private family foundation. The remaining sixty percent went into donor-advised funds under his and Susan's control. The filings do not specify the ultimate recipients of the donor-advised fund distributions. The timing — one month into a military campaign that by that point had already killed thousands of Palestinian civilians and displaced more than a million people — was not coincidental. It was a statement. *Source: Forbes; SEC filings, October 13-19, 2023; Brussels Morning.* The Michael and Susan Dell Foundation maintains an explicit Israel program on its website. The program's stated goals include connecting Israeli students to STEM pathways and creating social mobility through national service — by which they mean Israeli Military service, which they describe as connecting people

to new career paths and empowering them to contribute meaningfully to their communities. The foundation's language carefully avoids naming the Israeli Military directly. The function is the same. *Source: Dell Foundation website, dell.org/israel.* Peter Thiel, Founders Fund, and the Epstein Introduction That Built a Kill Chain I have documented Peter Thiel and Palantir extensively in Chapter Five of this book. But there is a dimension of that story that belongs here, in the section about how Silicon Valley's financial relationships with Israel were built, because the origin of Thiel's most consequential Israeli relationship begins in the same place where so much of Silicon Valley's darkest history begins: with Jeffrey Epstein. DOJ-released documents show that Epstein spent years from 2014 onward connecting Peter Thiel and former Israeli Prime Minister and Defense Minister Ehud Barak. He arranged at least six separate meetings between them because, in Epstein's own framing, both men wanted to connect Silicon Valley with Israeli national security. A June 2014 dinner at Epstein's Manhattan apartment brought all three together to discuss Israeli security. Thiel's reply to Epstein's description of the evening was: 'Agreed — very Israeli and very stimulating. Thanks for including me.' They met again in Israel in June 2017. *Source: DOJ-released documents; Jacobin, reported by Branko Marcetic; Boing Boing, March 10, 2026.* That relationship produced concrete business. Epstein steered Thiel toward two Israeli cybersecurity startups with roots in Unit 8200: Guardicore and Carbyne. Carbyne was the first Israeli investment by Thiel's Founders Fund. The deal was arranged through Epstein's introductions. Founders Fund led a fifteen million dollar Series B for Carbyne in August 2018 — a company founded by former Unit 8200 officers that Ehud Barak had pushed Epstein to pitch to Thiel. *Source: Boing Boing, March 10, 2026; DOJ documents; Jacobin.* In January 2024, following a meeting in Tel Aviv between Israeli defense officials and Palantir co-founders Peter Thiel and Alex Karp, Palantir announced a strategic partnership with Israel's Ministry of Defense to supply artificial intelligence technology in support of war-related missions. The AI systems Palantir supplied — including the Lavender targeting system — have been documented by +972 Magazine and Local Call as having flagged more than thirty-seven thousand Palestinians as potential militants, assigning each a threat score from an algorithm operating with at least a ten percent error rate. Under standard Israeli Military protocols, an algorithmic high score was sufficient to authorize a lethal strike. The Israeli Military internally raised its acceptable civilian casualty threshold for high-value targets from fifteen to twenty people per strike to as high as one hundred to three hundred people per strike. *Source: Bloomberg, January 12, 2024; +972 Magazine; Local Call; UN Special Rapporteur Francesca Albanese report.* When asked about ethical concerns with Palantir's role in Gaza, Thiel stated: 'My bias is to defer to Israel. It's not for us to second-guess everything. And I believe that broadly the IOF gets to decide what it wants to do, and that they're broadly in the right.' *Source: Cambridge Union debate, May 8, 2024; multiple news outlets.* Palantir's connection to Epstein. Palantir's connection to Unit 8200. Palantir's AI generating the kill lists. Thiel's deference to the Israeli Military as a matter of stated principle. This is not a set of coincidences. This is a system. It was built deliberately, through relationships cultivated over years, through financial investments made through a convicted sex offender's introductions, through board meetings in Tel Aviv and dinners in Manhattan. And it is currently running, in real time, on the civilians of Gaza. What the Pipeline Does I want to bring this back to the structural argument, because the individual names and dollar amounts are important but

they are not the whole story. The whole story is about what this pipeline produces. When Unit 8200 veterans move into content moderation at Meta, they bring with them training in identifying and suppressing content that the Israeli state considers threatening. When they move into security engineering at Google, they bring expertise in surveillance architecture built on a population that had no rights and no recourse. When they move into AI and machine learning at Apple, they bring pattern recognition skills honed on the digital lives of people whose medical records, sexual histories, and political opinions were compiled without consent and used for blackmail. That is the raw material that built their careers. That is what Silicon Valley is hiring. The former Unit 8200 commander at Google who spent six years helping understand patterns of hostile activists — his words — is now Google's Head of Strategy and Operations. The question is not whether he is a good engineer. The question is what he means by hostile, who he considers an activist, and what patterns he is trained to see as threatening. The answer, given his training, is: the same people his unit was built to suppress. *Source: MacLeod, MintPress News, 2022.* I ran Ethics in Technology for years. I interviewed hundreds of people in this industry. I watched Project Nimbus get signed while four hundred employees at Amazon and Google tried to stop it. I watched the employee who led that resistance get pushed out. I watched TikTok hire a Unit 8200 commander during the peak of the Gaza assault to manage what it called global incidents. I watched Oracle build propaganda tools for the Israeli government and then acquire control of the platform those tools were designed to manipulate. I was at Amazon Web Services. I sold the infrastructure. I know how the contracts work, how the partnerships are structured, how the language of cloud services and AI analytics is used to describe systems whose actual function is surveillance, targeting, and the suppression of civilian resistance. I left in 2018. What was built after I left is a matter of public record. It is documented in the sources at the end of this chapter. It is documented in the testimony of the workers who tried to stop it. It is documented in the kill lists that Palantir's algorithm generated and the Israeli Military deployed. Larry Ellison called Israel home and wrote the largest personal check in the history of an organization that exists to fund the soldiers of that military. Michael Dell wrote a three hundred and fifty million dollar check one month into the same military campaign. Peter Thiel said his bias is to defer to Israel and signed a partnership to supply the AI systems used to generate bombing targets. Their executives said if you don't agree with our mission to support the State of Israel, maybe we aren't the right company for you. These are not political opinions. These are business decisions. Made by named people. With documented dollar amounts. While Palestinian children were dying in numbers that the database I built at ParentsPlea.com has been recording, one profile at a time, since the first week of the assault. The technology sector did not stumble into complicity. It invested in it. It staffed it. It named it a golden opportunity and signed the contracts. I have been building the counter-record for years. It will outlast all of them. When the CEO of a defense contractor calls genocide a golden opportunity, he is being honest about what the system values. When the CEO of a cloud computing company calls the same infrastructure "enterprise solutions," he is being dishonest about exactly the same thing.

Both of them are building the same machine. --- The financial architecture that makes genocide a quarterly opportunity does not exist in a vacuum. It is maintained by a political architecture that ensures the people who should be holding these companies accountable are

instead funded by them. The revolving door between Silicon Valley and Washington is not a metaphor. It is a documented, legal, systematic mechanism for converting public regulatory authority into private profit. That mechanism is the subject of Chapter Seven.

CHAPTER SEVEN — THE REVOLVING DOOR

How Big Tech Bought the Government — and the Government Bought Big Tech Robin Williams — God rest his soul — said it best. Politicians, he said, should be required to wear sponsor jackets like NASCAR drivers, so that everybody knows who owns them.

He said it as a joke. He meant it as a diagnosis.

I have spent years at the intersection of technology and politics — first as a startup founder trying to understand why government contracts went to connected incumbents rather than innovative competitors, then as an Amazon Web Services employee watching the relationship between Silicon Valley and Washington up close, then as an advocate trying to understand why every attempt to regulate the companies causing the most harm seemed to dissolve before it became law. What I learned in all of those years is that Robin Williams was not exaggerating. The system is not broken. It is working exactly as designed — designed by and for the people who are profiting from it.

The mechanism that makes this possible has two components. The first is the revolving door. The second is Citizens United. Together, they have produced a political economy in which the companies causing the most harm are the most insulated from accountability, and the representatives who are supposed to hold them accountable are instead funded by them.

The Revolving Door

On the website BigTechSellsWars.com, you can find the names and profiles of technology executives, lobbyists, and government officials who have moved between Silicon Valley companies and government agencies — in both directions — in ways that give new meaning to the word "conflict of interest."

The pattern is consistent across administrations, across parties, and across decades. A senior official at the Defense Advanced Research Projects Agency — DARPA, the agency that funded the original internet — leaves government and joins the board of a technology company that then wins DARPA contracts. A former senior intelligence official joins Palantir's advisory structure and then advocates in Washington for expanding the intelligence community's use of Palantir's platforms. A Department of Defense procurement officer retires and joins a defense contractor, where their primary value is the relationships they built while authorized to spend public money on that contractor's competitors.

This is not illegal. That is the most disturbing part. The revolving door is not a corruption scandal that gets prosecuted. It is a feature of the system, operating openly and legally, producing outcomes that any person with functioning moral senses can recognize as corrupt even when no law has technically been broken.

The technology companies benefiting most from this revolving door are the ones with the largest government contracts: Amazon, Google, Microsoft, Palantir, Oracle. These are not companies that have government contracts because they offer the best products at the best prices. They have government contracts because they have invested in the relationships, the lobbying

infrastructure, and the former officials who know where the procurement budgets are and how the contracting decisions are made.

Palantir is the clearest example. Founded with CIA venture capital money in 2003, Palantir spent its first decade building relationships inside the intelligence community so deep that by the time it began competing for large government contracts, the people evaluating its bids were often people who had spent their careers inside the agencies Palantir was pitching to. The company did not win those contracts by accident. It won them by design — by engineering the relationships that would produce the decisions it needed.

Citizens United and the Purchase of Policy

The revolving door moves people. Citizens United moves money. And in the United States' political system, money is the more decisive variable.

In 2010, the Supreme Court ruled in Citizens United v. Federal Election Commission that corporations have the same First Amendment rights as individuals and cannot be prohibited from spending unlimited amounts of money on political advertising. The practical effect of this ruling was to remove the last meaningful limit on the ability of corporations to shape the political environment in which their regulatory oversight is determined.

Before Citizens United, a corporation that wanted to influence an election could donate limited amounts to candidates through a Political Action Committee, subject to disclosure requirements. After Citizens United, that same corporation could create or fund a nonprofit — structured so as to avoid full disclosure requirements — and spend unlimited amounts on advertising, issue advocacy, and political messaging.

The technology companies that are the subject of this book have used this architecture aggressively. Amazon, Google, Microsoft, Meta, Apple, and their affiliated interests have spent hundreds of millions of dollars on lobbying, on political advertising, and on the SuperPACs and dark money nonprofits that Citizens United made possible. They have funded candidates from both parties. They have funded think tanks that produce policy papers justifying the regulatory environment they want. They have funded academic chairs at universities that produce research supporting their positions on antitrust, privacy, and AI governance.

The result is a political system in which the companies most in need of regulation are the most capable of preventing it. The bills that would hold them accountable die in committee. The regulations that would constrain their data practices get watered down before implementation. The antitrust cases that would break up their monopolies take a decade to litigate while the monopolies grow stronger.

This is not a failure of democracy. It is democracy operating as it was redesigned by Citizens United: a system in which the size of your wallet determines the volume of your voice, and the companies with the largest wallets have the loudest voices in every room where the rules are written.

Both Parties. Every Administration.

I want to be clear about something that sometimes gets lost in this discussion. The revolving door and the corruption it enables are not a Republican problem or a Democratic problem. They are a systemic problem that has operated continuously across administrations of both parties.

The Obama administration's relationship with Silicon Valley was intimate and admiring. The Trump administration — both the first and the second — celebrated a different set of technology billionaires while delivering the same regulatory permissiveness to the same companies. The Biden administration prosecuted some antitrust cases but oversaw the period in which Project Nimbus was extended and the companies providing military AI infrastructure faced no accountability for it.

The people who have worked both sides of this revolving door include Democrats and Republicans, Obama appointees and Trump appointees, career civil servants and political appointees. The common denominator is not ideology. It is the financial relationship between the company and the official — the job offer waiting, the advisory board seat available, the speaking engagement on offer.

As documented in my documentary *Forever Peace Now* and on the website BigTechSellsWars.com, U.S. tech companies and their government partners have been spending over $44 billion since 2004 on contracts with companies like Amazon, Google, Microsoft, and others — contracts that build the surveillance infrastructure, the drone targeting systems, and the data processing platforms that enable the wars I document in this book. That number has grown dramatically since. The contracts are bipartisan. The beneficiaries are bipartisan. The civilian death tolls are not.

Canada's Lesson for Washington

Canada's recent break from U.S. military procurement dependence — Prime Minister Mark Carney's declaration that "the days of Canada's military sending 70 cents of every dollar to the United States are over" — is not just a trade story. It is a revolving door story.

The reason 70 cents of every Canadian defense dollar went to the United States is the same reason that American companies dominate European cloud infrastructure, the same reason that Palantir won the NHS contract in the UK, the same reason that no-bid contracts keep going to the same companies. Relationships built over decades between defense officials, procurement officers, and the contractors they favor — relationships that survive administration changes because the contractors funded both sides.

Canada's pivot is possible in part because Carney's government came in without the same web of incumbent contractor relationships that Washington operates through. They looked at the numbers — 70 cents on every dollar going south — and asked the question that every government with the political will to ask it can answer honestly: is this sovereignty, or is this dependency?

The answer was dependency. And they changed course.

What This Means for Everything Else in This Book

The revolving door and Citizens United are the explanation for why everything else in this book has been allowed to happen.

Why does Amazon have a $1.2 billion cloud contract with the Israeli military? Because the people who award government contracts move between Amazon and the agencies that award them, and the people who might legislate against it are funded by Amazon's political spending.

Why does Palantir have a $30 million no-bid contract with ICE? Because its founders and executives have spent years cultivating the relationships inside the intelligence and law enforcement communities that produce those awards.

Why has no technology executive been prosecuted? Because prosecution requires political will, and political will requires politicians who are not financially dependent on the companies they would need to prosecute.

This is the circle. It does not break from the inside. It breaks when external accountability — legal, criminal, international — creates consequences that the political system cannot absorb or redirect.

That is why the petition to Spain matters. That is why criminal prosecution, not regulatory fines, is the instrument that actually changes the calculus.

The sponsors on the jacket are visible to anyone who knows how to read a campaign finance disclosure or a lobbying registration form. The question is not whether we can see them. The question is whether we are willing to act on what we see.

--- The revolving door explains who gets protected. Chapter Eight explains who pays the price. The workers at the bottom of Big Tech's supply chain — the content moderators in Nairobi watching footage of people having sex through Meta's smart glasses, the warehouse workers in New Jersey sustaining injuries at 2.6 times the industry average while Amazon posts record profits, the small business owners whose margins have been consumed by a platform that now competes with the sellers whose fees funded its growth — these are the people the revolving door was designed to forget.

CHAPTER EIGHT — THE INVISIBLE WORKFORCE

How Big Tech Breaks the Bodies and Minds of the Workers It Cannot See — and Then Disposes of Them In February 2026, a group of workers in Nairobi, Kenya, began speaking to Swedish journalists. Two newspapers — Svenska Dagbladet and Göteborgs-Posten — published the story that those workers, at enormous personal risk, had decided the world needed to hear. They worked for a company called Sama, a Nairobi-based outsourcing firm contracted by Meta. Their job was to review footage captured by users of Meta's Ray-Ban smart glasses — to annotate what was in the video clips, label objects and actions, and help train Meta's artificial intelligence systems. That is the polite, corporate description of what they were asked to do.

What they actually saw was something different.

During long ten-hour shifts, these workers were shown footage of people having sex, captured on camera without their knowledge. They saw people changing clothes in their bedrooms. They saw children filmed without consent. They saw individuals using toilets. They saw credit card numbers accidentally captured on camera as someone used their glasses to help a friend pay for something. This footage — intimate, private, often violating — came from users of Meta's glasses who had no idea that their daily lives were being sent to a labor force in Kenya to be catalogued, labeled, and used to train artificial intelligence. The workers who spoke out said they were not allowed to bring phones or any personal devices to work. Security cameras monitored their every move. The implication was clear: Meta knew this footage was sensitive enough to require containment of anyone who saw it. (Source: Gadget Review, "Meta Fires Contractors Who Exposed Ray-Ban Glasses Privacy Scandal," May 2026, https://www.gadgetreview.com/meta-fires-contractors-who-exposed-ray-ban-glasses-privacy-sca ndal; Android Headlines, "1,100 AI Trainers Were Fired After Blowing the Whistle on Meta's Ray-Ban Privacy Problem," May 2026, https://www.androidheadlines.com/2026/05/1100-ai-trainers-were-fired-after-blowing-the-whistle- on-metas-ray-ban-privacy-problem.html)

Less than two months after the story broke, Meta ended its contract with Sama. The Nairobi firm confirmed the termination of 1,108 workers — some receiving as little as six days' notice. Meta claimed Sama "did not meet its standards." Sama said it had consistently met all contractual expectations and had never been notified of any deficiency. The logic of what happened is not difficult to follow. Workers blew the whistle. Meta ended the contract. 1,108 people lost their jobs.

Meta has described the footage review as standard procedure covered by its terms of service. That is technically true, in the same way that every predatory agreement is technically disclosed in fine print that no one reads. The people whose intimate moments were shipped to Kenya to be watched by strangers during a ten-hour shift did not meaningfully consent to that. And the workers who had to watch it are now without income because they had the audacity to say so. (Source: OECD.AI Incident Report, "Meta's AI Smart Glasses Lead to Worker Harm and Privacy Violations in Kenya," April 30, 2026, https://oecd.ai/en/incidents/2026-04-30-7774)

This Is Not the First Time

The Kenya Sama story is new. The pattern it represents is not.

This is not the first time Meta has left a trail of traumatized and discarded workers in Nairobi. In January 2023, all 260 content moderators working at Sama's Nairobi hub — people who spent their working hours reviewing Facebook content for violence, child exploitation, and graphic material — were told their jobs were redundant. Forty-three of those workers filed a lawsuit for unlawful dismissal. A class action eventually grew to involve 185 workers. Dr. Ian Kanyanya, head of mental health services at Kenyatta National Hospital, assessed 144 of them. His findings, filed with Nairobi's employment and labor relations court in December 2024, were devastating: 81 percent of the workers assessed were suffering from severe PTSD. (Source: CNN Business, "Facebook inflicted 'lifelong trauma' on content moderators in Kenya, campaigners say, as more than 140 are diagnosed with PTSD," December 22, 2024, https://edition.cnn.com/2024/12/22/business/facebook-content-moderators-kenya-ptsd-intl)

One of the workers who led the protest — a South African man named Daniel Motaung — was fired for organizing. He sued Meta and Samasource Kenya, alleging forced labor, human trafficking, and union-busting. He told investigators that his team was forced to view hours of horrific content including beheadings and child sexual exploitation — for pay of less than $2.20 an hour. Over 80 labor organizations from across the world wrote to Meta demanding it stop attempting to silence him. Among the signatories was Meta's most high-profile whistleblower, Frances Haugen. (Source: Quartz Africa, "Meta faces pressure from rights groups to stop gagging Daniel Motaung," https://qz.com/africa/2191286/meta-faces-pressure-from-rights-groups-to-stop-gagging-daniel-m otaung)

This is a company valued at over a trillion dollars. It pays the people who protect the rest of its users from the worst content on its platforms less than the cost of a large coffee in San Francisco — and then terminates them when they complain about the trauma of doing it.

The Science of What This Work Does to Human Beings

I want to be precise about what content moderation does to a person, because the industry has perfected the art of keeping this invisible. These are not hypothetical harms. They are documented clinical realities.

Scientific peer-reviewed research confirms that content moderators — the people paid by technology companies to view graphic and distressing material so the rest of us don't have to — suffer at measurably higher rates from PTSD, anxiety, depression, intrusive thoughts, sleep disturbance, emotional detachment, and burnout. A cross-sectional survey published in 2024 found that 53 percent of self-identified content moderators reported clinical levels of general psychological distress. Research shows a dose-response relationship: the more traumatic content someone is exposed to, the worse their mental health outcomes. The exposure does not wear off when the shift ends. (Source: Research paper, "I've Seen Enough: Measuring the Toll of Content Moderation on Mental Health," arXiv:2511.09813, November 2024, https://arxiv.org/pdf/2511.09813; Cyberpsychology Journal, "The psychological impacts of content moderation on content moderators," 2023, https://cyberpsychology.eu/article/view/33166)

Researchers at SAGE Journals have characterized what happens to content moderators not as an accident of the job but as a structural feature of it: the commodification of traumatic vulnerability itself. In other words, the companies are not failing to protect these workers. They

are deliberately purchasing the capacity of these workers to absorb trauma — buying their psychological health as an input cost, the way you buy servers or storage or bandwidth. (Source: Amit Pinchevski, "Social media's canaries: content moderators between digital labor and mediated trauma," SAGE Journals, 2023, https://journals.sagepub.com/doi/abs/10.1177/01634437221122226)

The people doing this work in Kenya — and in Ethiopia, and in the Philippines, and in Indonesia, and in every other low-wage country where Big Tech has outsourced the burden of its most traumatic operational requirement — are not edge cases. They are a workforce of tens of thousands, doing labor that makes every social media platform you use every day possible. And they are doing it for wages, under conditions, and with mental health consequences that would be illegal to impose on workers in the countries where the profits are collected.

That is not an oversight. That is a business model.

The Amazon Warehouse: Blood, Sweat, and Next-Day Delivery

I want to move from the psychological labor of content moderation to the physical labor of fulfillment — because both tell the same story about how Big Tech treats the workers at the bottom of its supply chain: as a resource to be optimized until it breaks, then discarded.

In August 2015, the New York Times published an investigation — based on interviews with more than a hundred current and former Amazon employees — that described a workplace characterized by midnight emails, morning texts demanding to know why the midnight email had gone unanswered, employees crying at their desks, and workers experiencing health crises so severe they were brought out on stretchers. The reporters, Jodi Kantor and David Streitfeld, quoted a former employee saying: "Nearly every person I worked with, I saw cry at their desk." (Source: New York Times, "Inside Amazon: Wrestling Big Ideas in a Bruising Workplace," August 15, 2015, by Jodi Kantor and David Streitfeld, https://www.nytimes.com/2015/08/16/technology/inside-amazon-wrestling-big-ideas-in-a-bruisin g-workplace.html)

That investigation described the white-collar experience at Amazon's corporate offices. What has since been documented about the blue-collar experience in Amazon's warehouses is, if anything, worse — and far more physically immediate.

In December 2024, Senator Bernie Sanders released a 160-page Senate HELP Committee investigation into Amazon's warehouse safety record. Its findings were damning and detailed. Amazon's injury rate in 2022 was 6.6 percent — against an industry average of 3.2 percent. Amazon warehouse workers suffer serious injuries at a rate 2.6 times higher than their non-Amazon counterparts. An internal Amazon report — titled Project Soteria and cited by Senate investigators — found a direct relationship between the speed of tasks performed and the rate of worker injury. Amazon reviewed this evidence and chose not to act on it, according to Senate investigators, due to financial considerations. (Source: U.S. Senate HELP Committee, Amazon Investigation Report, December 2024, https://www.help.senate.gov/imo/media/doc/amazon_investigation.pdf; CNN Business, "Senate report accuses Amazon of ignoring worker safety in productivity push," December 16, 2024, https://www.cnn.com/2024/12/16/business/amazon-worker-safety-senate-report)

The most common injuries at Amazon warehouses are sprains, strains, and muscle tears — the kind of injuries that would sideline a professional athlete for weeks. The workers sustaining

them are not athletes recovering in specialized medical facilities. They are people who cannot afford to stop working, picking up shifts as Amazon's automated systems monitor their every movement, log every deviation from productivity targets, and generate disciplinary action for any worker who does not maintain the required pace. Injury rates spike during Prime Day, Black Friday, and Cyber Monday — exactly the moments when productivity pressure is at its highest and the human body is pushed furthest past its limits.

I remember, when I was at Amazon Web Services, the culture that the New York Times described was not a fiction. I saw it. I lived adjacent to it. I watched the machine demand everything and offer nothing in return except a paycheck and the expectation of gratitude. The corporate employees at least had the dignity of a salary and a desk and the option to update their résumé. The warehouse workers have a scannable wristband and a productivity quota and a body that is slowly being consumed by the pace of next-day delivery.

The Death of Main Street by Algorithm

There is a third category of harm that Amazon has inflicted — quieter than the warehouse injuries, less dramatic than the crying at desks, but perhaps more structurally devastating to the communities that make up the actual fabric of American life. It is what has happened to the small businesses that once lined the main streets of every town in this country.

In 2016, according to the e-commerce research firm Marketplace Pulse, Amazon took roughly one-third of every dollar made by merchants selling on its platform, in the form of fees and advertising. By 2022, that number had surpassed 50 percent. By 2024, Amazon had collected over $150 billion in revenue from third-party seller fees alone — a number so large it would rank inside the Fortune 25 if it were its own company. (Source: Fortune, "Amazon's record earnings contain a mystery: How much are fees from small sellers to thank?" February 7, 2025, https://fortune.com/2025/02/07/amazon-2024-earnings-third-party-seller-fees-profits/)

When you combine referral fees, fulfillment costs, advertising spending, and storage charges, the total bite Amazon takes from a seller's revenue runs between 30 and 50 percent — often more. For small sellers with already thin margins, that arithmetic is not a business model. It is a slow execution. Independent analysis confirms that for many product categories, the cost of selling on Amazon now exceeds any conceivable profit margin for a small producer. The platform that was supposed to give every small business a global storefront has become a mechanism for extracting value from the producers of goods at a rate that would have been recognized immediately in any other era as monopolistic extraction. (Source: SmartScout, "The Top Challenges Facing Amazon Third-Party Sellers in 2025," April 2025, https://www.smartscout.com/blog/amazon-third-party-sellers-problems)

I grew up in the Bay Area. I remember what a main street looked like before Amazon. I remember the hardware store where the owner knew your name. The bookshop where someone could recommend something you had never heard of. The clothing store run by a family that had been in the neighborhood for twenty years. I remember summer jobs — the kind of entry-level, community-embedded work that gave a teenager their first paycheck, their first experience of responsibility, their first relationship with a neighborhood business that was also a neighborhood institution.

Those jobs are gone. Those stores are closed. The people who owned them could not compete with a company that uses algorithmic pricing, logistics infrastructure built with government subsidies, and a marketplace model that takes half of every dollar — and then uses the data from that marketplace to identify which products to manufacture under Amazon's own private label, in direct competition with the sellers whose fees funded the platform's growth. This is not free-market competition. It is a rigged table at which only one player can never lose.

And the workers who replaced the shop owners — the people packing and shipping the goods in the warehouses — are not getting what the shop owners had. They have no ownership. No relationship with the customer. No stake in the product. They have a productivity quota, a scannable badge, and a body that is being systematically injured at 2.6 times the industry rate while Amazon posts record profits and its founders build rockets.

What Comes Next: Robots, and the Illusion of Concern

Amazon has announced with considerable fanfare its investments in warehouse robotics and automation — framing them as safety improvements for workers. There is a meaningful sense in which reducing the number of times a human being lifts a fifty-pound box by replacing the lift with a machine is a genuine improvement. I will acknowledge that straightforwardly.

But let me also say what is not being acknowledged: the same automation that reduces some physical strain will eliminate most of the jobs. Amazon has been working toward fully automated fulfillment for years. When those jobs go — and they will go — the communities that built their economic identity around the Amazon warehouse on the edge of town will not receive a severance package from the algorithm that replaced their neighbors. They will receive nothing. The workers who spent years being injured at twice the industry rate building Amazon's logistics empire will be thanked for their service and shown the door.

This is the endpoint of a business model that has always treated workers as an input cost to be minimized and replaced. First, it minimized wages by offshoring content moderation to Kenya. Then it minimized labor costs by squeezing warehouse workers at speeds that break their bodies while monitoring every second with surveillance technology. Now it is minimizing the cost of labor entirely, by replacing human beings with machines — not because the workers were inadequate, but because the machines are cheaper.

The question that no quarterly earnings call will ever ask is: what do we owe the people whose labor we consumed on the way to not needing them anymore? What do we owe the communities whose main streets we closed, whose summer jobs we eliminated, whose local economies we hollowed out — on the way to building the most profitable logistics company in human history?

The White-Collar Harvest: When the Employee Becomes the Dataset

Amazon took the bodies of its warehouse workers and broke them. The injury rates, the productivity quotas, the stretchers, the Senate investigation — all of it is documented above. The endpoint of Amazon's model is no longer speculative. It is a fully automated warehouse. The workers who built it will be thanked for their service and shown the door. The communities around those warehouses will receive nothing.

What Meta announced on April 21, 2026 is the same story — moved upstairs. Into the glass offices. Onto the corporate laptops. Into the hands and minds of the engineers, product managers, and designers who believed, because they had degrees and stock options and

ergonomic chairs, that they were a different category of worker than the people packing boxes in New Jersey.

They are not. They never were. The extraction just took longer to reach them.

On April 21, 2026, Meta disclosed to its U.S.-based employees — in an internal memo posted to a channel belonging to the company's Meta Superintelligence Labs team, obtained by Reuters and CNBC — that it had installed tracking software on their corporate computers. The program is called the Model Capability Initiative, or MCI. It captures mouse movements, clicks, keystrokes, and periodic screenshots of employees' screens. It runs continuously across a designated list of work-related applications and third-party platforms. The data feeds directly into Meta's AI training pipeline — used to build AI agents capable of performing, without human involvement, the same white-collar tasks the employees are being recorded performing. (Source: Reuters, "Meta tells employees it's tracking their keystrokes to train AI," April 21, 2026, https://www.reuters.com/technology/meta-tells-employees-its-tracking-their-keystrokes-train-ai-2 026-04-21/; CNBC, "Meta is tracking employee keystrokes on Google, LinkedIn, Wikipedia as part of AI training initiative," April 22, 2026, https://www.cnbc.com/2026/04/22/meta-tracks-employee-usage-on-google-linkedin-ai-training-p roject.html)

The list of platforms being monitored was widely circulated inside Meta. It includes Google, LinkedIn, Salesforce's Slack, Microsoft's GitHub, Atlassian, and Meta's own properties. Notably, the list originally included OpenAI's ChatGPT and Anthropic's Claude — the two most widely used AI assistants in professional settings — before those platforms were quietly removed from the tracking list without explanation. The inference is not difficult to draw: capturing how Meta's own employees interact with competitor AI systems would expose the company to legal, competitive, and reputational consequences it was not prepared to absorb. The employees themselves received no equivalent protection. (Source: CNBC, ibid.)

What MCI Can Actually See — and Why That Should Terrify You

I want to stop here and ask a question that the corporate communications around MCI have been careful not to ask directly. What does a program that captures mouse movements, keystrokes, and continuous screenshots of an employee's screen throughout their working day actually see?

Not in the abstract. In practice.

It sees every email a worker receives from their doctor's office — appointment reminders, test results, treatment updates, prescription confirmations. It sees the pharmacy order a worker places online during a lunch break for a medication they would prefer their employer know nothing about. It captures the name of that medication, the dosage, the refill date, and the condition it treats — whether that is depression, HIV, diabetes, a psychiatric disorder, a chronic illness, or a pregnancy. None of that information was offered to Meta voluntarily. None of it is relevant to AI model training. All of it is now in a training pipeline.

It sees the text message a worker sends from their computer to their spouse about a child's diagnosis. It sees the insurance portal a worker visits to check whether a procedure is covered. It sees the search queries a worker runs to understand what a new diagnosis means, or what a medication's side effects are, or whether their condition qualifies for disability accommodation. It sees the moment a worker reads a piece of news that terrifies them — about a policy change, a

drug approval, a clinical trial — and the private, immediate, human response they type to someone they trust.

It sees the bank notification that appears on screen when a worker's direct deposit arrives. It sees the credit card statement a worker opens to dispute a charge. It sees the financial account a worker accesses to check whether they can afford the treatment their doctor just recommended. It captures account numbers, transaction histories, and the particular combination of financial stress and medical need that defines what is happening in a person's life at any given moment.

It sees the political article a worker reads during their lunch break. It sees whether that article is from a conservative outlet or a progressive one. It sees the social media post a worker opens and reads in full — including the political content, the party affiliation of the person who wrote it, and whether the worker lingers, shares it, or types a response. In an era when employers screen candidates for political beliefs, when immigration authorities monitor social media, and when political affiliation can affect professional relationships in ways that are difficult to document but impossible to ignore — this is not a trivial capture. It is a profile.

It sees the personal email from a family member describing a health crisis — a parent's cancer diagnosis, a sibling's addiction, a child's mental health hospitalization — landing in a worker's inbox at two in the afternoon while they are trying to finish a product brief. It sees the response the worker types. It sees the web searches they run afterward. It sees the grief happening in real time, on a corporate device, because human beings do not cleanly separate their lives from their work hours no matter how many employee handbooks tell them to try.

None of this was supposed to be a data collection exercise. All of it now is.

Meta's Response: "Don't Check Personal Email on Your Work Computer"

When Meta employees raised exactly these concerns in internal discussions — how would the company prevent the capture of health information, financial data, immigration status, family communications — Meta CTO Andrew Bosworth's response was captured in internal communications obtained by Platformer. His answer, in full: "Gmail is an approved context so if you have concerns it may be best not to check personal email on your work computer." (Source: Platformer, "The week that Meta employees became training data," April 2026, https://www.platformer.news/meta-mci-monitoring-layoffs-knowledge-work/)

Read that again.

The people whose health records, financial information, family communications, and political views are now being captured to train Meta's AI were told, by their company's Chief Technology Officer, that their only recourse was to stop using their work computers for anything personal.

This is not a privacy protection. It is a liability shift. It is Meta telling its employees: we have installed surveillance infrastructure on your machine that will capture everything visible on your screen, and the responsibility for ensuring your most private information does not enter our training pipeline belongs to you. Not to us. You.

That answer is not just callous. It is legally and ethically incoherent. Workers in the United States have well-established rights under federal and state law that govern what employers can and cannot collect about their medical conditions and health status. The Americans with Disabilities Act, the Health Insurance Portability and Accountability Act, the Genetic

Information Nondiscrimination Act — these are not aspirational documents. They are federal law. They exist precisely because Congress recognized that the power imbalance between an employer and an employee, left unchecked, will always resolve in the employer's favor — and that certain categories of personal information are too sensitive, too intimate, and too potentially weaponizable to be left to the goodwill of the party with the greater power.

Meta has now built a system that captures exactly this category of information — not deliberately, perhaps, but systematically, inevitably, and continuously — and has told its employees that the solution is behavioral self-regulation on the part of the surveilled. The company that cannot be trusted to build a facial recognition system ethically, that terminated 1,108 Kenyan workers for speaking about what they saw, that complied with 94 percent of Israeli government censorship requests — this company now holds a continuously updated, screenshot-by-screenshot record of its employees' medical lives, financial lives, political lives, and family lives, and assures them it will not be misused.

I am not reassured. Neither, judging by internal discussions described as "dystopian" by multiple employees, are they. (Source: CNBC, ibid.)

The Political Dimension: What Your News Feed Says About You

Let me be specific about the political surveillance dimension of MCI, because it has received less attention than the medical and financial angles — and it is potentially the most dangerous of all.

MCI runs on Google. It runs on LinkedIn. It runs across the web browsing that happens on an employee's corporate device throughout the day. In practice, that means it captures what news sources a worker reads, which political stories they open and read to completion, which they skip, which social media posts they engage with, and what they type in response to political content.

This is a comprehensive political profile — assembled not from declared opinions or party registrations, but from the actual behavioral fingerprint of how a person consumes political information in private moments throughout their working day. It captures the difference between a worker who reads only mainstream centrist news and one who reads independent progressive journalism. It captures who reads articles about union organizing, about immigration enforcement, about the political affiliations of Silicon Valley executives. It captures, in other words, exactly the information that employees in any politically sensitive environment have the strongest interest in keeping private from their employer.

The memo's assurance that this data "will not be used for performance assessments" deserves exactly the skepticism it has received. Meta has a documented history of making assurances about data use and then violating them. Its terms of service once said it would not sell user data. The Facebook Papers documented that internal decisions routinely prioritized engagement and profit over the safety policies Meta publicly committed to. The company that fired 1,108 Kenyan workers for speaking to journalists about what they saw on their work screens is asking its American employees to trust that their keystroke data will remain in the training pipeline and go nowhere else.

History does not support that trust. The data is the value. The data will be used.

The Superintelligence Memo: The Organizational Architecture Behind MCI

MCI did not appear in a vacuum. It is the operational expression of a corporate restructuring that Zuckerberg announced in the weeks before — and which reveals precisely how Meta is treating the data supply problem at the core of its AI ambitions.

In a memo released in full by CNBC, Zuckerberg announced the formation of Meta Superintelligence Labs, or MSL — a new unit consolidating all of Meta's AI model development, product teams, and fundamental AI research under single command. To lead it, Zuckerberg named Alexandr Wang — the 28-year-old former CEO of Scale AI, which Meta acquired a 49 percent stake in for $14.3 billion — calling him "the most impressive founder of his generation" and creating for him the new role of Chief AI Officer. Wang's appointment was not ceremonial. Scale AI built its entire business on one specific capability: harvesting workflow data from contractors to train AI systems. Wang said in 2024: "For a lot of the capabilities that we want to build into the models, the biggest blocker is actually a lack of data. There's no pool of really valuable agent data that's just sitting around anywhere. And so we have to figure out how to produce really high quality data." (Source: Entrepreneur, "Mark Zuckerberg Reveals Meta Superintelligence Labs," https://us.entrepreneur.com/business-news/mark-zuckerberg-reveals-meta-superintelligence-lab s/494047; Platformer, ibid.)

MCI is how Meta produces that data. And the workers generating it are, in structural terms, Meta's version of Scale's contractor workforce — except that these workers have salaries, stock options, and the expectation that their employer would not record their screen in perpetuity and feed the result into a model designed to eliminate their jobs.

Zuckerberg committed up to $135 billion in capital expenditure for 2026 to build Meta's AI infrastructure. At the same time, the company confirmed it would lay off ten percent of its workforce — approximately 8,000 people — while declining to fill an additional 6,000 open positions. These announcements came in the same week as MCI. (Source: Fortune, "Meta will start tracking employees' screens and keystrokes to train AI tools," April 21, 2026, https://fortune.com/2026/04/21/meta-will-start-tracking-employees-screens-and-keystrokes-to-tr ain-ai/)

The data collection and the displacement are not parallel stories. They are the same story, in sequence. Capture the knowledge. Train the model. Reduce the headcount. Repeat.

Zuckerberg offered his vision on the earnings call that preceded these announcements. "We're starting to see projects that used to require big teams now be accomplished by a single very talented person," he told investors. That sentence is the operating manual for everything that follows. The talented person's cognitive process is being recorded. The AI trained on it will eventually perform that process without the person. The compliment is the announcement of the mechanism. And the employees being complimented are funding it with every keystroke.

On a separate internal memo, Meta CTO Andrew Bosworth described the destination with equal clarity: "The vision we are building towards is one where our agents primarily do the work and our role is to direct, review and help them improve." He called it a "closed loop" — agents that would "automatically see where we felt the need to intervene so they can be better next time." (Source: Detroit News, "Meta employee mouse movements, keystrokes, AI training data," April 21, 2026, https://eu.detroitnews.com/story/tech/2026/04/21/metaemployee-mouse-movements-keystrokes -ai-training-data/89717625007/)

The people whose keystrokes are being recorded to train the agents are, in Bosworth's own words, the people whose jobs the agents are designed to perform. Let no one pretend otherwise.

The Pattern Completed: From Warehouse to Workstation, from Body to Mind

I want to be precise about what distinguishes MCI from every earlier form of worker surveillance documented in this chapter.

Amazon's monitoring of its warehouse workers is about controlling the pace of physical labor. The scannable wristbands, the productivity quotas, the algorithmic discipline systems — they are designed to maximize output per body per hour, until the body is replaced by a machine. What is extracted is physical energy. What is broken is the body — the tendons, the backs, the rotator cuffs of people packing boxes at speeds that Amazon's own internal report documented as directly causing injury.

Meta's MCI extracts something different and something more intimate: the cognitive process of a human mind at work. The click sequences that navigate a complex engineering decision. The keyboard shortcuts that represent years of accumulated professional knowledge. The decision trees embedded in how an experienced product manager moves through the tools of their daily work. The correspondence with a doctor that happens to land on screen during that work. The credit card statement opened in a browser tab. The political article read during a lunch break. The text from a family member about a health crisis.

This is not energy being measured. This is a human life being digitized — the professional parts and the personal parts equally, because human beings conducting their working day on a machine do not sort themselves cleanly into categories that a tracking program can separate. Their private life follows them to work. The machine captures everything it can see.

The Amazon warehouse worker whose body is consumed in the service of next-day delivery receives no share of the logistics empire their labor built. The Meta engineer whose cognitive process — and medical history, financial data, and political views — is captured and fed into an AI training pipeline receives no equity in the agent their keystrokes trained. In both cases, the worker is the raw material. In both cases, the worker is told to be grateful for the opportunity to contribute. And in both cases, the endpoint is the same: the worker is made redundant by the very product their labor produced, while their most private information remains in a pipeline they can no longer access or control.

The content moderators in Nairobi are the invisible workforce of Meta's past — low-wage, offshore, discarded when they became inconvenient. The engineers and product managers subjected to MCI are the invisible workforce of Meta's present — well-compensated, American, and equally powerless once the extraction is complete. The geography is different. The salary is different. The chairs have better lumbar support. The structural position is identical: they are an input cost to be minimized and a data resource to be harvested, and those two facts exist in sequence, not in conflict.

Let me be direct. My medical condition is none of Meta's business. What medication I take during my working hours is none of Meta's business. The credit card number that appears on my screen when I pay a bill during lunch is none of Meta's business. The political article I read, the news source I prefer, the party affiliation I may or may not hold — none of it is the property of

the company that installed software on my machine and told me the solution was to stop checking my personal email at work.

This is not a terms of service dispute. It is a violation of the basic human right to a private life — a right that does not pause when you sit down at your desk in the morning and resume when you log off at night. The law in this country has historically recognized that. The question is whether the law will move fast enough to catch what Meta is doing before the data is already trained into a model that cannot be untrained.

A fool with a tool is still a fool. A company that harvests its employees' minds, their medical records, their financial lives, and their political identities to build the machines that replace them is something more deliberate than foolish.

It is operating exactly as designed.

Sources:

Reuters, "Meta tells employees it's tracking their keystrokes to train AI," April 21, 2026: https://www.reuters.com/technology/meta-tells-employees-its-tracking-their-keystrokes-train-ai-2026-04-21/

CNBC, "Meta is tracking employee keystrokes on Google, LinkedIn, Wikipedia as part of AI training initiative," April 22, 2026: https://www.cnbc.com/2026/04/22/meta-tracks-employee-usage-on-google-linkedin-ai-training-p roject.html

Fortune, "Meta will start tracking employees' screens and keystrokes to train AI tools," April 21, 2026: https://fortune.com/2026/04/21/meta-will-start-tracking-employees-screens-and-keystrokes-to-tr ain-ai/

Platformer, "The week that Meta employees became training data," April 2026: https://www.platformer.news/meta-mci-monitoring-layoffs-knowledge-work/

Detroit News, "Meta employee mouse movements, keystrokes, AI training data," April 21, 2026: https://eu.detroitnews.com/story/tech/2026/04/21/metaemployee-mouse-movements-keystrokes-ai-training-data/89717625007/

Entrepreneur, "Mark Zuckerberg Reveals Meta Superintelligence Labs": https://us.entrepreneur.com/business-news/mark-zuckerberg-reveals-meta-superintelligence-lab s/494047

ArtVoice, "Meta Is Recording Its Employees' Keystrokes And Mouse Clicks To Train AI That Could Eventually Replace Them," April 22, 2026: https://artvoice.com/2026/04/22/meta-is-recording-its-employees-keystrokes-and-mouse-clicks-t o-train-ai-that-could-eventually-replace-them/

The Social Obligation That Has Been Abandoned

Every major technology company that has built its empire on human labor has made, implicitly or explicitly, a bargain with the society that makes its success possible. That bargain goes something like this: you give us your labor, your data, your consumer spending, your infrastructure, and your regulatory tolerance. In return, we give you jobs, economic participation, and a share — however modest — in the prosperity we generate.

Amazon has broken that bargain. Meta has broken it. Every Big Tech company that outsources its most traumatic labor to low-wage countries, pays below living wages, monitors workers at granular surveillance levels, maintains injury rates that would be considered unconscionable in any other industry, and then prepares to automate the jobs away while posting record profits — has broken it.

The shareholder is not the only stakeholder. This is not a radical idea. It is the foundation of every labor law, every workplace safety regulation, every minimum wage, every right to organize that workers have ever fought for and won. A company exists in a community. It uses that community's roads, its water, its educated workforce, its legal system, its courts — all of the public infrastructure that makes commerce possible. It owes that community something in return.

What Amazon owes the workers in its warehouses is not just a paycheck. It is working conditions that do not break their bodies. It is injury rates that are not 2.6 times the industry average. It is an honest accounting of what automation will do to their livelihoods, and a plan for transition that treats them as human beings rather than units of labor whose productive life has expired.

What Meta owes the content moderators it has hired through contractors in Nairobi is not just a termination with six days' notice. It is mental health care proportionate to the trauma it purchased. It is compensation for the psychological harm it imposed deliberately, systematically, and profitably. It is a basic acknowledgment that the people it hired to watch beheadings and child sexual exploitation and intimate footage captured through smart glasses are not disposable components of a data pipeline but human beings whose suffering was a cost that Meta chose to offload rather than internalize.

The workers cannot be the ones who absorb the externalities of Big Tech's profits. That is the argument this chapter is making. It is not a new argument. It is the oldest argument in the history of labor. But it has to be made again — because the companies making the decisions have convinced themselves, and many of their shareholders, that it no longer applies to them.

It applies to them. It will always apply. The arc of accountability, like the arc of history, is long. But I have seen enough to know it bends.

--- The workers I have just documented are made invisible by design. So are other things. In Chapter Nine, I turn to a different kind of erasure — not of labor and bodies, but of voice and evidence. The same platforms that extract value from workers in Kenya and warehouses in New Jersey have simultaneously erased the testimony of an entire people under siege. What follows is the documented record of how Meta, TikTok, X, YouTube, and LinkedIn silenced Palestine — and why that silencing is not a bug in the system, but one of its most consistent features.

CHAPTER NINE — THE DIGITAL IRON CURTAIN

How Meta, TikTok, X, YouTube, and LinkedIn Silenced Palestine There is a word for what happens when a government controls what its citizens can see, say, and share. We call it censorship, and we have historically associated it with authoritarian regimes — Soviet-era information blackouts, Chinese internet firewalls, North Korean state media. We did not expect to find it operating at the fingertip level of a billion-dollar Silicon Valley algorithm, quietly throttling the voices of besieged people while the bombs fell.

But that is precisely what happened. And it is still happening.

Since October 7, 2023, the most powerful social media platforms on earth — Meta's Instagram and Facebook, TikTok, X (formerly Twitter), YouTube, and LinkedIn — have functioned as a coordinated, if not always consciously coordinated, suppression machine. Journalists have been banned without explanation. Human rights organizations have had their archives erased. Activists documenting civilian deaths have been shadowbanned into invisibility. Evidence of potential war crimes — video, photographs, firsthand testimony — has been deleted from the public record at the request of a foreign government, with a compliance rate that should shock every person who has ever been told these platforms stand for freedom of expression.

They do not. They stand for profit and political protection. This chapter documents what they did — platform by platform — and who paid the price.

Meta: Instagram and Facebook — A Government's Preferred Censor

The scale of Meta's suppression of Palestinian voices during the Gaza genocide is not disputed. It is documented, quantified, and in Meta's own words — admitted.

In December 2023, Human Rights Watch released a landmark 51-page report, Meta's Broken Promises: Systemic Censorship of Palestine Content on Instagram and Facebook, which documented over 1,050 verified cases of content removal and suppression from more than 60 countries between October and November 2023 alone. The pattern was not random. It was systematic: content deleted, accounts suspended without explanation, hashtags rendered unsearchable, Instagram Live features disabled, and the quiet algorithmic throttling known as shadowbanning — reducing the visibility of a post or an account without notifying the person being silenced. Palestinian journalist Ahmed Shihab-Eldin, whose Instagram account had nearly one million followers, lost access to his account five times in the weeks following October 7. He was not alone. (Source: Human Rights Watch, "Meta's Broken Promises," December 2023, https://www.hrw.org/report/2023/12/21/metas-broken-promises/systemic-censorship-palestine-c ontent-instagram-and)

The mechanism Meta used most frequently was its "Dangerous Organizations and Individuals" policy — a broad content moderation framework that incorporates the U.S. government's terrorist designation lists. Because Hamas is on that list and governs Gaza, Meta's automated systems flagged and removed content that merely mentioned Hamas in a neutral or

journalistic context. The AI could not distinguish between a reporter documenting a Hamas statement and a person expressing support for violence. It flagged both. According to Human Rights Watch, Meta removed even neutral mentions of Hamas in relation to developments in Gaza. The result was that reporting on the most heavily covered genocide in the world was systematically deleted from two of the most widely used platforms in human history.

What makes this more than a content moderation failure is the role of the Israeli government. According to media reports cited in the Human Rights Watch investigation, Israel's Cyber Unit sent Meta and other platforms approximately 9,500 content takedown requests in the weeks following October 7, 2023 — 60 percent of which went to Meta. Platforms responded with a 94 percent compliance rate. Let that number sit for a moment. When a foreign government sends a list of content it wants removed from an American social media platform, Meta complies 94 percent of the time. According to subsequent reporting, Meta complied with approximately 96 percent of Israeli government removal requests — the highest compliance rate of any country in the world. This is not content moderation. This is a foreign government using an American corporation as a censorship arm, and the corporation accepting the role willingly. (Source: Access Now, "How Meta Censors Palestinian Voices," February 2024, https://www.accessnow.org/publication/how-meta-censors-palestinian-voices/; World Socialist Web Site, December 2023, https://www.wsws.org/en/articles/2023/12/22/rrys-d22.html)

Meta's own translation algorithm added the word "terrorist" to the bios of Palestinian Instagram users — automatically, without their knowledge or consent — in October and November 2023. When this was discovered and reported, Meta called it a bug. An error. A technical malfunction. The same explanation it has offered, and continues to offer, every single time its systems are caught doing something that happens to benefit one side of a genocide.

This was not Meta's first offense and they knew it. A 2021 independent report commissioned by Meta itself — conducted by Business for Social Responsibility — found that the company's content moderation "appear[s] to have had an adverse human rights impact on the rights of Palestinian users," adversely affecting their ability to document and share their experiences. Meta agreed to make changes. Almost two years later, Human Rights Watch found that those changes had not been implemented. The "broken promises" of the report's title referred not to a surprise failure, but to a documented pattern of commitment and non-delivery. (Source: Human Rights Watch, https://www.hrw.org/news/2023/12/20/meta-systemic-censorship-palestine-content)

TikTok: The Platform They Tried to Ban for Showing Too Much Truth

TikTok's story in the Gaza genocide is different from Meta's — and in some ways more revealing about how power operates in Silicon Valley and Washington, D.C.

Unlike Instagram and Facebook, TikTok did not emerge as a primary suppressor of Palestinian content. The data pointed in the opposite direction. A Washington Post analysis in November 2023 found that the hashtag #freepalestine appeared on TikTok 38 times more often than #standwithisrael — a ratio comparable to what the same study found on Meta's own platforms. The Arab Center for Social Media found that for every pro-Israel post on TikTok, there were approximately 54 pro-Palestine posts. Palestinian journalists and documentarians — including Emmy Award-winning filmmaker and journalist Bisan Owda, whose videos beginning with "It's Bisan from Gaza, and I'm still alive" reached audiences that mainstream Western media

was not reaching — found in TikTok a distribution channel that had not yet been captured by the forces seeking to control the narrative. (Source: Washington Post, "TikTok Disables Hashtag View Counts After Gaza War Controversy," February 2024, https://www.washingtonpost.com/technology/2024/02/08/tiktok-remove-data-criticism-gaza/)

The response from Washington was immediate. By October and November 2023, Republican senators including Josh Hawley and Marco Rubio were explicitly citing TikTok's Gaza coverage as a reason to ban the app. Representative Mike Gallagher, who spearheaded the TikTok divestiture legislation, argued in a public op-ed that the app was responsible for turning young Americans against Israel. The legislation passed. The Chinese parent company ByteDance was told to sell or be banned. The proposed buyer of TikTok's U.S. operations was Oracle — the same Oracle providing database infrastructure to the Israeli military's intelligence apparatus, as documented in Chapter Two. (Source: The Intercept, "The TikTok Ban Is Also About Hiding Pro-Palestinian Content," January 2025, https://theintercept.com/2025/01/09/tiktok-ban-israel-palestine-republicans/)

When TikTok was briefly banned for U.S. users in January 2025, one of its most prominent Palestinian journalists, Bisan Owda, was permanently suspended by the platform shortly afterward — without warning, without explanation, in what she described as "political pressure linked to Israel." Her account, which had documented the human cost of the Gaza genocide in real time to millions of viewers, was gone. The pro-Palestine news outlet Mondoweiss had its TikTok account permanently banned without warning during an earlier surge in Israeli military operations. TikTok, meanwhile, reported that it had removed 4.7 million videos and suspended 300,000 livestreams between October 7, 2023 and September 15, 2024 — invoking its policies against promoting Hamas, hate speech, or misinformation. The question of how those policies were applied, and to whose content, is one the platform has never answered in public with any transparency. (Source: Middle East Online, "US Ownership of TikTok Leads to Bans on Palestinian Influencers," January 2026, https://middle-east-online.com/en/us-ownership-tiktok-leads-bans-palestinian-influencers-journal ists; Euronews, October 2024, https://www.euronews.com/next/2024/10/07/human-rights-ngos-say-social-media-platforms-cont inue-to-censor-pro-palestine-content)

What TikTok reveals is not primarily a story of the platform suppressing Palestinian content — it is a story of a platform being legislated out of existence precisely because it was not suppressing it. When a social media platform allows the truth to be seen, Washington intervenes. That is the lesson of TikTok and Gaza. And it is a lesson worth understanding clearly before moving on.

X / Twitter: The "Free Speech" Platform That Silenced Palestine

Elon Musk purchased Twitter in 2022 with a stated commitment to free speech absolutism. He positioned himself as the liberator of a platform he claimed had been strangled by liberal censorship. He reinstated banned accounts. He fired the trust and safety teams. He replaced professional fact-checkers with a crowdsourced system called Community Notes that functions, in the words of his own critics, as the digital equivalent of leaving a complaint in an unmonitored suggestion box.

What Musk's "free speech" platform actually produced, where Palestinian content was concerned, was a selective application of silence. In October 2023, X suspended hundreds of Palestinian accounts, describing them as affiliated with Hamas, without providing evidence for individual determinations or a meaningful appeals process. Among those suspended in January 2024 were prominent journalists and commentators whose accounts were removed without warning and without stated cause — including senior staff writer Alan MacLeod of Mint Press News, Intercept reporter Ken Klippenstein with over 500,000 followers, and a cluster of other left-wing and pro-Palestinian journalists and activists. Their accounts were restored hours later. No explanation was provided. No rule that had been violated was named. (Source: Al Jazeera, "Twitter Under Fire for Censoring Palestinian Public Figures," February 2023, https://www.aljazeera.com/features/2023/2/28/twitter-under-fire-for-censuring-palestinian-public- figures; World Socialist Web Site, January 2024, https://www.wsws.org/en/articles/2024/01/10/doux-j10.html)

While Palestinian voices were being throttled, a November 2023 report by the Centre for Countering Digital Hate found that X was failing to moderate 96 percent of hate speech posts related to the Israeli genocide in Palestine— including posts describing Palestinians as animals, calling for collective punishment, and inciting violence against civilians. The same platform that suspended hundreds of Palestinian accounts for alleged Hamas affiliation was leaving anti-Palestinian incitement online at a rate of 96 percent. This is not an inconsistency. It is a policy. The digital rights organization 7amleh documented more than 19,000 cases of hate speech and inciting content in Hebrew on X in the weeks following October 7 — content that remained online without consequence. (Source: Business and Human Rights Centre, "Report Exposes X's Failure to Remove 96% of Hate Speech Posts," November 2023, https://www.business-human-rights.org/en/latest-news/report-exposes-xs-failure-to-remove-96-of -hate-speech-posts-amid-israel-palestine-conflict/)

On Christmas Eve 2023, Instagram — not X — permanently terminated the account of activist Shaun King, who had been documenting Israeli war crimes to a massive audience. On X, actress Susan Sarandon posted a photograph of billboards calling for a ceasefire. It was flagged as "violent speech." The billboard said: Stop the Bombs. This is what free speech absolutism looks like when it is built and operated by a man who has publicly declared that empathy is civilization's fundamental weakness.

YouTube: Erasing the Evidence

YouTube's role in the suppression of Palestinian content has two distinct phases. The first is a years-long pattern of algorithmic discrimination against Arabic-language content from the West Bank and Gaza — documented in research by 7amleh, Al-Shabaka, and Article 19 well before October 2023. The second is more recent, more direct, and more damning: YouTube's deletion, at the request of the Trump administration, of over 700 videos documenting Israeli human rights violations.

The earlier pattern is important context. Research conducted by 7amleh and reported by Al-Shabaka found that YouTube's AI content moderation systems applied what researchers called "locative discrimination" — flagging content at higher rates based on the geographic origin of the upload. Videos from the West Bank and Gaza faced a higher level of automated scrutiny than equivalent content from elsewhere. Arabic-language content was flagged disproportionately

compared to Hebrew-language content of equal sensitivity. Palestinian journalists who challenged these removals received no explanation from YouTube — and no meaningful appeal pathway. The platform's AI had rendered a verdict, and there was no court of appeal. (Source: Al-Shabaka, "YouTube's Violation of Palestinian Digital Rights," https://al-shabaka.org/briefs/youtubes-violation-of-palestinian-digital-rights-what-needs-to-be-do ne/)

Then, in October 2025, YouTube quietly erased the accounts of three of Palestine's most important human rights organizations: Al-Haq, the Al Mezan Center for Human Rights, and the Palestinian Centre for Human Rights — the oldest human rights organization in Gaza, as designated by the United Nations. Gone with those accounts were over 700 videos: investigations into Israeli strikes on civilians, testimonies from survivors, a documentary about mothers surviving the genocide in Gaza, and a filmed investigation into the killing of Palestinian-American journalist Shireen Abu Akleh. The content was not removed because it was violent. The content was not removed because it violated community guidelines. It was removed because the United States government sanctioned these organizations for cooperating with the International Criminal Court's investigation into Israeli officials.

YouTube confirmed to The Intercept that the deletions were made to comply with U.S. sanctions. Katherine Gallagher, a senior attorney at the Center for Constitutional Rights, called it "outrageous that YouTube is furthering the Trump administration's agenda to remove evidence of human rights violations and war crimes from public view." Sarah Leah Whitson of Democracy for the Arab World Now described it as "a disappointing act of submission." The Palestinian Centre for Human Rights said plainly that YouTube's actions "shield perpetrators from accountability." (Source: The Intercept, "YouTube Quietly Erased More Than 700 Videos Documenting Israeli Human Rights Violations," November 2025, https://theintercept.com/2025/11/04/youtube-google-israel-palestine-human-rights-censorship/; Common Dreams, https://www.commondreams.org/news/youtube-deletes-videos-israel)

Let me be direct about what this means. YouTube — a subsidiary of Google, which holds a $1.2 billion cloud contract with the Israeli military through Project Nimbus — deleted evidence of potential war crimes committed by the Israeli military, at the request of a U.S. administration that has shielded Israel from accountability at the International Criminal Court. This is not a content moderation decision. This is corporate complicity in the erasure of a historical record. It is the digital equivalent of burning documents before the trial.

LinkedIn: The Professional Network That Watches You

LinkedIn presents itself as the respectable face of Silicon Valley — the platform for résumés and thought leadership, for career milestones and industry networking. It is owned by Microsoft, whose cloud infrastructure, as documented in Chapter Two, is embedded in the Israeli military's operational systems. And in April 2026, LinkedIn became the subject of a class action lawsuit and a sweeping investigative report that exposed one of the most comprehensive covert surveillance operations in corporate history.

The investigation — published by a nonprofit research organization called Fairlinked and dubbed "BrowserGate" — revealed that every time a user visits linkedin.com on a Chrome-based browser, a hidden JavaScript program silently scans their device for installed browser extensions. The program runs without any visible indicator, requests no consent, and discloses nothing. It

reports its findings — a precise fingerprint of the user's browser environment — directly to LinkedIn's servers. The scale of the operation has grown dramatically: LinkedIn scanned for 38 specific extensions in 2017. By 2024, that number had grown to 461. By February 2026, it had reached 6,167 — a 1,252 percent increase in two years. A class action lawsuit was filed in the U.S. District Court for the Northern District of California on April 7, 2026, accusing LinkedIn of running a "covert surveillance system" embedded in its own website. (Source: CyberInsider, "LinkedIn Faces Class Action Over Alleged Covert Scanning of Users' Browsers," April 2026, https://cyberinsider.com/linkedin-faces-class-action-over-alleged-covert-scanning-of-users-browsers/; Cybernews, https://cybernews.com/privacy/linkedin-lawsuits-illegal-browser-extension-tracking/)

The significance of this surveillance goes beyond privacy violation for its own sake. The Fairlinked report argues — and the class action complaint alleges — that the data collected through this browser scanning is precise enough to allow LinkedIn to infer sensitive attributes about users: political affiliations, religious beliefs, health conditions, union membership, and ideological orientation. The fingerprint persists across cookie resets. In an era when pro-Palestinian advocacy has been classified by some governments as grounds for employment discrimination, immigration screening, and even criminal investigation, a professional network that can silently identify your political leanings from your browser profile is not a neutral tool. It is a surveillance instrument embedded in the professional infrastructure of over one billion people worldwide. In October 2024, the Irish Data Protection Commission fined LinkedIn €310 million for processing users' personal data for targeted advertising without a valid legal basis — the largest fine ever issued to LinkedIn under Europe's GDPR framework. The BrowserGate revelations suggest the problem runs deeper than advertising. (Source: The Next Web, "LinkedIn BrowserGate: Extension Scanning, Privacy, Fingerprint," April 2026, https://thenextweb.com/news/linkedin-browsergate-extension-scanning-privacy-fingerprint; SecurityWeek, https://www.securityweek.com/browsergate-claims-of-linkedin-spying-clash-with-security-research-findings/)

LinkedIn has denied that it uses the scanned data to infer sensitive information about users. It says the scanning is a security measure to detect tools that violate its terms of service by scraping data without consent. That is a plausible explanation for scanning dozens of known scraping tools. It is not a plausible explanation for scanning 6,167 browser extensions — including extensions for personal productivity, health tracking, political research, and religious community organizing. At some point, a security measure becomes a surveillance infrastructure. The lawsuit will determine which category LinkedIn's behavior falls into. What is not in dispute is that the scanning happened, that users were not informed, that no consent was sought, and that the data went to LinkedIn's servers. For a platform owned by Microsoft — a company whose cloud systems are documented tools of military targeting — the question of what that data is ultimately used for is not paranoia. It is the right question to ask.

"The War on Maps — How Big Tech Erases a People Before the Bombs Fall"

There is a form of violence that happens before the first airstrike. Before the first checkpoint. Before the first home is demolished. It happens on a screen, at a zoom level, in the quiet political

decisions of engineers and product managers who decide which places deserve a name on a map — and which places do not.

I want to talk about maps.

Because if you have been paying attention to what is happening in Lebanon — if you pulled up Apple Maps during the Israeli bombardment of April 2026 and looked at the southern part of the country — you saw something that stopped people cold. Most of it was blank. The villages were not labeled. The towns that had existed for centuries, that families had been ordered to evacuate under threat of bombardment, simply did not appear. Meanwhile, on the same app, small Israeli communities across the border — Nahariyya, Shelomi, Beit Jan — were clearly labeled, their names sitting on the map as if the landscape itself had taken a side. (Source: The New Arab, "Did Apple Erase Lebanese Village Names from Its Maps?," April 2026, https://www.newarab.com/news/did-apple-erase-lebanese-village-names-its-maps)

Apple's explanation — that these Lebanese villages were never in the application to begin with, that the detailed version of Apple Maps has not yet completed its global rollout in Lebanon — may be technically true. Multiple fact-checkers confirmed that archived user complaints about Lebanon's sparse Apple Maps coverage go back to 2019, years before any of this began. I am not making the claim that Apple removed villages in real time as Israeli bombs fell, because the evidence does not support that specific claim.

What I am saying is something more disturbing.

The villages were never in the maps. And the Israeli communities across the border always were.

That disparity did not begin in April 2026. It was not created by the current aggression. It is the baseline. It is the default. It is what these applications look like on an ordinary day, when no one is paying attention and no one is comparing screenshots side by side and no one is asking why places that have been inhabited for centuries simply do not appear on the most widely used navigation tool in the history of humanity.

That is not a technical limitation. That is a political choice that compounds, over years, into something that looks like erasure — because it is.

Google Maps and the Map That Never Said Palestine

Let me be specific, because specificity is what separates documentation from accusation.

Google Maps has never labeled Palestine as a state. Not once since the service launched in 2005. The West Bank and the Gaza Strip appear as geographic regions with no national designation, while Israel is fully labeled as a country with Jerusalem identified as its capital. This is despite the fact that Jerusalem's status is a final-status issue in every international framework resolution, where the UN General Assembly in 1947 granted Jerusalem international status under Resolution 181, and the International Court of Justice has been explicit that Israeli sovereignty over East Jerusalem is not recognized under international law.

Settlements appear labeled as if they were inside Israel. They are not. They are illegal under Article 49 of the Fourth Geneva Convention, which prohibits an occupying power from transferring its civilian population into occupied territory. But on Google Maps they appear as communities like any other — with clearly visible names, calculated routes, measured distances — while the Palestinian villages of Area C of the West Bank, land that has been farmed and

inhabited for generations, only appear when you zoom in so far that you have already passed them by.

Then there is the matter of routing. 7amleh, the Arab Center for the Advancement of Social Media, published a detailed report examining Google Maps' navigation decisions in the occupied territories. When a Palestinian in the West Bank tries to navigate from Bethlehem to Ramallah — a journey of 36 kilometers — Google Maps routes them through Jerusalem. Which most Palestinians cannot legally enter. Palestinian residents of Gaza and the West Bank carry green ID cards that restrict their movement. They cannot use Israeli-only roads. They cannot cross into Jerusalem without a permit that most of them do not have. Google Maps, built on the most sophisticated geospatial data in human history, cannot find a route. It directs people onto roads where they will be arrested for using them. (Source: +972 Magazine, "Lost in Occupation: How Google Maps Is Erasing Palestine," https://www.972mag.com/mapping-occupation-how-google-erases-palestine-from-its-maps/)

The Kyl-Bingaman Amendment — a U.S. law passed in 1997 — restricts the resolution of satellite imagery that American companies can provide of Israel and the occupied Palestinian territories. The high-resolution imagery that is freely available for every other country in the Middle East is deliberately degraded over Israel-Palestine. Researchers trying to document the destruction of villages, settlement construction, and evidence of possible war crimes have lower photographic resolution to work with in this particular conflict than they would have when analyzing any comparable situation anywhere else in the region. This is not a technological accident. It is a law passed by the United States Congress specifically to protect one country's ability to act without photographic scrutiny. (Source: Al-Shabaka, "Maps, Technology, and Decolonial Spatial Practices in Palestine," https://al-shabaka.org/briefs/maps-technology-and-decolonial-spatial-practices-in-palestine/)

In August 2016, Google's maps removed the labels "West Bank" and "Gaza Strip" entirely. Not the territories themselves, just their names. They disappeared from the map as if the geographic designations that everyone uses to refer to these places simply no longer existed. More than 615,000 people signed a petition: "Google: Put Palestine on your maps." Google called it a bug. It restored the labels. It said nothing more. (Source: +972 Magazine, ibid.)

A bug. The most sophisticated mapping company on the planet — a company that sends cars down every road in every country, that has mapped the ocean floor, that can tell you a restaurant's opening hours in Reykjavik — accidentally deleted the only geographic labels that acknowledged the existence of Palestinian territory.

I want you to compare that explanation with what you know about these companies. These are not organizations that make careless mistakes about the geography of disputed territories. They are organizations with entire teams of policy specialists, government liaison offices, legal departments, and geopolitical risk consultants whose job is precisely to manage decisions like these. Every labeling choice in Google Maps is a political decision. Every routing algorithm reflects implicit assumptions about who is allowed to go where. Every omission of a Palestinian village is a decision not to show it: a choice made, reviewed, and approved by human beings in offices in Mountain View and Cupertino who understood what they were choosing.

What a map does before the bombs fall

There is a reason militaries have always prioritized cartographic erasure. A place with no name on a map is harder to defend under international law. It is harder to report from. It is harder to mourn. When a journalist files a story about an airstrike and the location does not appear on any of the major mapping apps, the story becomes harder to situate — harder to verify — and harder to visualize for a public that navigates the world through apps built by American tech companies.

In southern Lebanon, during the April 2026 bombardment, Israel ordered hundreds of thousands of people to evacuate villages that did not appear on the maps most of the world was using to follow the news. The blank space on the map was not the cause of the bombing. But it was complicit in the invisibility of those who were bombed. When you cannot find a place, it is harder to mourn the people who lived there.

I am not going to let tech companies hide behind the word "bug." I am not going to let them hide behind "global rollout timelines" while one side of a border is blank and the other is fully labeled. I have been in this industry for thirty years. I know how to tell the difference between a technical limitation and a political decision dressed up in technical language. And so do they.

The map is not neutral. It has never been neutral. The question is: whose reality does it reflect — and whose does it erase?

Sign the petition at NoEthicsInBigTech.com. Demand that Google label Palestine as a state in accordance with UN General Assembly Resolution 67/19, which recognized Palestine as a non-member observer state with 138 votes. Demand that illegal Israeli settlements be identified as such, in accordance with all applicable provisions of international law. Demand that routing algorithms take into account the movement restrictions imposed on Palestinians, rather than directing them onto roads where they will be arrested for using them. These are not radical demands. They are requests to reflect reality.

The companies that build the maps we all use have the power to make the occupation visible or invisible. They have chosen, systematically, invisibility.

That choice has consequences. And the people who pay the consequences are not in Mountain View.

The Surveillance Industrial Complex — From Israeli Spyware to the Government Shopping Cart I have spent this chapter documenting how Big Tech platforms silence, censor, and erase the voices of people they have been asked to suppress. But I want to step back from the specific mechanism of censorship and address something more fundamental — because the erasure of speech is only one layer of the surveillance architecture that has been built over the past two decades. Beneath it is something older, more invasive, and more consequential: the infrastructure of watching. In my documentary *Forever Peace Now*, I documented what the surveillance economy actually looks like from the inside — through the testimony of engineers, civil liberties advocates, congressional staffers, and people who had lived inside the machine long enough to understand what it was doing. What I found, across years of those conversations, was a consistent pattern: the tools built ostensibly to protect us were being used against us. The promises made to sell those tools were false. And the people who knew were not speaking loudly enough for the public to hear. I want to tell you three of those stories here. They are connected. Together, they describe a surveillance infrastructure that no single law governs, that no single company controls,

and that the government has found ways to use without the messy inconvenience of obtaining a warrant.

NSO Group and the Spyware That Murdered a Journalist NSO Group is an Israeli technology company. It builds surveillance tools — specifically a piece of software called Pegasus that is, in the understated language of the cybersecurity community, one of the most invasive tools ever commercially deployed against civilian populations. Pegasus can be installed on a target's smartphone without any action by the user — no link to click, no file to open. It operates silently, invisibly, and comprehensively. Once installed, it can access messages, emails, call logs, photographs, and location data. It can activate the microphone and camera without the user's knowledge. It turns your phone into a surveillance device that reports everything it sees and hears to whoever purchased the access. NSO Group sells Pegasus to governments. Not to individuals, not to corporations — to governments, exclusively, who use it against journalists, activists, lawyers, dissidents, and opposition politicians. The company has always maintained that its tools are only sold to vetted governments for legitimate law enforcement purposes. The record does not support this claim. In October 2018, Jamal Khashoggi — a Saudi Arabian journalist and Washington Post columnist who had been critical of Crown Prince Mohammed bin Salman — entered the Saudi Arabian consulate in Istanbul. He was killed there. His body was dismembered. Investigators subsequently established that Khashoggi's phone had been infected with Pegasus spyware, and that the surveillance data derived from that infection had been used to track him in the period leading up to his murder. A tool sold by an Israeli company to a Gulf monarchy was used to locate, track, and facilitate the assassination of a journalist who worked for an American newspaper. (Source: Citizen Lab, University of Toronto, Pegasus reporting; Washington Post, multiple dates 2018–2019) The United States government placed NSO Group on its Entity List in November 2021 — a designation that restricts American companies from doing business with it — citing that NSO had "developed and supplied spyware to foreign governments that used these tools to maliciously target government officials, journalists, businesspeople, activists, academics, and embassy workers." Israel, whose government has regulatory authority over NSO Group's export licenses, has continued to allow the company to operate, because the tools it builds serve Israeli intelligence and diplomatic interests as much as they serve the governments that buy them. This is what I described in the film as the surveillance industrial complex — the merging of military-grade intelligence capabilities with commercial business models and government relationships, deployed against civilian populations worldwide. NSO Group is the most visible node of that complex. It is not the only one. It is simply the one whose tools were used to kill someone whose name you know. The Israeli technology sector — whose graduates from Unit 8200 of the Israeli Intelligence Corps flow directly into both Silicon Valley and Israeli defense tech startups, as I documented in Chapter Two — is the world's largest exporter of commercial surveillance technology per capita. Blue Wolf and Red Wolf, which I described in Chapter Two as the biometric occupation systems deployed against Palestinians in the West Bank, are products of the same ecosystem. The Gospel and Lavender targeting systems are products of the same ecosystem. And the clients for this surveillance technology are not limited to democracies with independent judiciaries and free press. They include governments whose record on human rights is well documented and whose use of these tools against their own citizens is not hypothetical.

The United States government, which placed NSO Group on the Entity List, has simultaneously continued to operate in partnership with the broader Israeli surveillance technology ecosystem through Project Nimbus, through Unit 8200 alumni's presence across Silicon Valley, and through the intelligence-sharing arrangements of the Five Eyes network. The condemnation of NSO Group and the embrace of the infrastructure that produced it are not contradictions. They are different applications of the same technology, serving different political masters.

Apple's Privacy Theater and the Promise That Was Never True Apple has built its brand on privacy. "What happens on your iPhone stays on your iPhone." The tagline is on billboards. It is in television commercials. It is in the keynote addresses that Tim Cook delivers with the measured cadence of a man who has decided that privacy is a product differentiator and is prepared to treat it as one. I have spent thirty years in the technology industry. I know what a product differentiator is. And I know the difference between a company that treats privacy as a value and a company that treats it as a marketing position. The record on Apple is, at best, complicated. Let me be specific about what I mean. In 2021, Apple announced plans to scan the content of iCloud photo libraries for child sexual abuse material — a goal that is genuinely important and morally urgent. The problem was the mechanism: Apple's proposed system would scan images on users' devices before they were uploaded to iCloud, comparing them against a database of known CSAM hash values. The Electronic Frontier Foundation and a coalition of security researchers and civil liberties organizations responded immediately with a letter signed by over 90 organizations from around the world: the system, however well-intentioned in its initial design, created infrastructure that could be repurposed. A scanning system built to detect one category of illegal content could, under government pressure or corporate policy changes, be expanded to scan for other content — political speech, religious material, journalism. Apple eventually abandoned the proposal, citing concerns from privacy advocates. But the fact that it was proposed at all revealed something important about how Apple thinks about the privacy it markets. In China, Apple has removed apps from its Chinese App Store at the direction of the Chinese government — including VPN apps that allow users to circumvent censorship, including apps used by protesters in Hong Kong during the 2019 demonstrations, including tools used by journalists working in a country where journalism requires protection. Apple has stored Chinese users' iCloud data on servers operated by a Chinese state-owned enterprise called GCBD — meaning that the encryption keys for Chinese users' most private communications and data are held by a company legally required to cooperate with the Chinese government. Apple has complied with Chinese government requests to remove Taiwanese flag emoji from devices operating in China. The privacy that Apple promises its users is, in the Chinese market, explicitly conditional on what the Chinese government is willing to allow. This is not a surprise. Tim Cook has said, with characteristic corporate diplomacy, that Apple believes it can do more good by operating in China than by withdrawing. I do not question the sincerity of that belief. I do question what it means in practice for the users who trust Apple with their most intimate data, believing the promise of the billboard, not knowing that the promise has a geographic boundary that stops at the jurisdictions where it becomes commercially inconvenient to honor it.

The Encryption Lie: How the NSA Paid to Break the Internet In 2004, the National Security Agency made a secret payment of $10 million to RSA Security — at the time one of the most

trusted names in commercial encryption — in exchange for making a specific algorithm the default random number generator in RSA's widely distributed BSAFE cryptographic library. That algorithm was called Dual EC DRBG: Dual Elliptic Curve Deterministic Random Bit Generator. RSA told its customers that it was an approved, secure algorithm. RSA had put their trust in the standards and guidance from NIST — the National Institute of Standards and Technology, the U.S. government body responsible for certifying that encryption standards are what they claim to be. Cryptographers had flagged problems with Dual EC DRBG as early as 2007. The mathematics of the algorithm suggested — to anyone who looked closely — that there might be a secret relationship between two of its key parameters that would allow someone who knew the relationship to predict its outputs. In other words: a backdoor. A way to decrypt communications that the users believed were protected. In 2013, documents released by Edward Snowden confirmed what the mathematics had suggested. The NSA had deliberately inserted the backdoor into Dual EC DRBG, had pushed it through NIST's standardization process to give it the appearance of independent certification, had paid RSA $10 million to make it the default in products used by developers around the world, and had spent $250 million per year through a classified program called Bullrun specifically to insert backdoors into commercial encryption software and hardware. NIST — the agency responsible for telling the world that the standard was secure — had been, in the language of the Snowden documents, "eventually" brought under NSA's control as the sole editor of the standard. (Source: Reuters, "Exclusive: Secret contract tied NSA and security industry pioneer," December 20, 2013, https://www.reuters.com/article/us-usa-security-nsa-rsa/exclusive-secret-contract-tied-nsa-and-s ecurity-industry-pioneer-idUSBRE9B-J1C220131220; New York Times, "N.S.A. Able to Foil Basic Safeguards of Privacy on Web," September 5, 2013; Wikipedia, Dual_EC_DRBG, https://en.wikipedia.org/wiki/Dual_EC_DRBG) Let me translate this into plain language, because the technical framing has allowed it to escape the moral clarity it deserves. The United States government told the world that a specific encryption standard was secure. The government knew it was not secure. The government had secretly made it not secure, on purpose, by paying the most trusted name in commercial encryption to build the backdoor in and deploy it at scale. Developers around the world used that standard in good faith to protect their users' private communications. Those communications were accessible to the NSA. The people whose private data was exposed were never told. The companies whose products contained the backdoor were, by RSA's own account, not informed of the vulnerability at the time they accepted the $10 million. This is not a historical footnote. This is the foundational event in modern surveillance architecture — the moment when it became documented fact that the agency responsible for telling Americans their digital communications were secure was simultaneously the agency secretly ensuring they were not. The companies that told you your data was protected were, in some cases, being paid to make sure it wasn't.

The Government Doesn't Need to Spy Anymore. It Can Just Shop. Here is the part that should stop you cold, because it represents a development more recent and more structurally significant than either Pegasus or the NSA backdoor. The United States government no longer needs to build surveillance infrastructure to watch its citizens. It no longer needs to install spyware on phones, or insert backdoors into encryption standards, or run room-sized data collection operations at AT&T facilities — though all of those things still happen. It does not

need to do any of them, because the surveillance has already been done. It is being done continuously, every day, by the apps on your phone. And the government can simply buy the results. The data broker industry — companies like Venntel, Babel Street, and dozens of others — purchases vast quantities of data from smartphone applications, web browsers, and connected devices. The location data from a fitness tracking app. The browsing history from a news aggregator. The purchase patterns from a retail loyalty program. The political content engagement from a social media platform. All of this is aggregated, cross-referenced, and sold in detailed dossiers that can reveal a person's movements, their health conditions, their political beliefs, their religious practices, their personal relationships, and their daily routines — in many cases with greater precision and completeness than any direct wiretap could achieve. The Department of Homeland Security, the FBI, the IRS, ICE, and the Secret Service have all purchased cell phone location data from data brokers without warrants. DHS signed a $1 billion contract with Palantir to deploy AI-powered data analytics systems across all DHS components — including CBP and ICE — using data purchased through these channels. Government agencies have argued that the Fourth Amendment's requirement for a warrant does not apply when they are purchasing commercially available data rather than compelling its disclosure — a legal distinction that privacy lawyers have called constitutionally indefensible but which courts have not yet definitively resolved. ICE uses tools to track mobile phones and identify devices that have visited specific locations. They have a $30 million contract with Palantir for ImmigrationOS — a platform designed to track people for deportation with near-real-time visibility. The same company whose tools are documented as central to the targeting systems killing civilians in Gaza is also the company building the deportation tracking infrastructure that uses purchased data to find immigrants who have committed no crime beyond their presence. Congress has failed to close this loophole. The Fourth Amendment Is Not For Sale Act passed the House with bipartisan support in 2024 — it would prohibit government agencies from purchasing data they would otherwise need a warrant to obtain — and the Senate never voted on it. The Trump administration has pushed for clean reauthorization of FISA without privacy reforms. The surveillance architecture grows larger. The loophole remains open. What this means in practice is something I have been trying to explain to policymakers, advocates, and anyone who will listen since I produced *Forever Peace Now*: the surveillance state does not need to surveil anymore. It outsourced the surveillance to the free market. Every app you install, every service you use, every device you carry is doing the surveillance on its behalf. The Fourth Amendment was written to protect you from government intrusion into your private life. It was not written for a world in which your private life is continuously catalogued by private companies and sold to the highest bidder, who happens to be the government. The law has not kept up. The surveillance has. And now, with the AI tools described in Chapter Twelve, the government's ability to process, analyze, and act on purchased data has expanded to a scale that no previous surveillance program could have achieved. Not because the government built better tools. Because it hired the tools from the same companies building the AI systems that everyone uses every day. The AT&T technician who discovered the NSA's secret room — the splitter cabinet in the internet room at 611 Folsom Street in San Francisco, where a glass prism split the laser light beam carrying internet traffic and sent one copy to the NSA — described what he had seen in testimony that appeared in my film. He

found it in 2003, he reported it, and nothing changed. The room is still there. And the room is no longer necessary — because the data flows to the government now through a commercial transaction rather than a secret wire tap, and no whistleblower can easily see the purchase order. The lesson of thirty years in technology, of Snowden and Pegasus and the RSA backdoor and the data broker marketplace, is this: the promise of privacy was never the product. The product was always the data. And the people who told you otherwise — the companies that put "privacy" on billboards, the standards bodies that certified algorithms they knew were compromised, the government agencies that told you they needed the surveillance to protect you — were, in each case, describing something other than what they were actually doing. The most dangerous surveillance state in the world is not China, with its social credit system and its facial recognition cameras. It is the United States — because ours is the only one that has perfected the art of building comprehensive surveillance infrastructure while maintaining, with a straight face and a beautifully designed billboard, that privacy is the product.

Sources: Citizen Lab, University of Toronto, Pegasus spyware research: https://citizenlab.ca/category/research/ Reuters, "Exclusive: Secret contract tied NSA and security industry pioneer," December 20, 2013: https://www.reuters.com/article/us-usa-security-nsa-rsa/exclusive-secret-contract-tied-nsa-and-s ecurity-industry-pioneer-idUSBRE9BJ1C220131220 Wikipedia, Dual EC DRBG: https://en.wikipedia.org/wiki/Dual_EC_DRBG Brennan Center for Justice, "Closing the Data Broker Loophole," February 2024: https://www.brennancenter.org/our-work/research-reports/closing-data-broker-loophole Electronic Privacy Information Center (EPIC), "Closing the Data Broker Loophole: Government Evasion of the Fourth Amendment": https://epic.org/documents/closing-the-data-broker-loophole-government-evasion-of-the-fourth-a mendment/ NPR, "Your data is everywhere. The government is buying it without a warrant," March 25, 2026: https://www.npr.org/2026/03/25/nx-s1-5752369/ice-surveillance-data-brokers-congress-anthropi c Project on Government Oversight, "Closing the Data Broker Loophole": https://www.pogo.org/fact-sheets/fact-sheet-closing-the-data-broker-loophole Forever Peace Now (documentary film), directed by Vahid Razavi, 2025 — testimony of AT&T technician, Electronic Frontier Foundation, and civil society advocates on surveillance architecture

The Pattern, Named

Taken together, these five platforms form a digital architecture of suppression. Meta deleted over a thousand pieces of Palestinian content in sixty days and complied with 94 percent of Israeli government censorship requests. YouTube erased 700 videos documenting potential war crimes at the direction of the Trump administration. X suspended Palestinian journalists while leaving anti-Palestinian hate speech online at a 96 percent rate. TikTok became the target of a congressional ban explicitly motivated by its failure to suppress Palestinian content. LinkedIn built a covert surveillance system capable of inferring users' political affiliations without their knowledge or consent.

None of these companies will describe what they did as censorship. They will use words like "content moderation," "community guidelines," "sanctions compliance," and "platform integrity." These are the bureaucratic labels of an industry that has perfected the art of doing terrible things with respectable language.

I have interviewed the policy makers who make these decisions. I know what "platform data integrity" means when the quarterly earnings call is next week. It means: we will do what is financially and politically convenient, and we will call it policy.

The Palestinian Observatory of Digital Rights Violations documented more than 1,350 instances of online censorship across major platforms between October 2023 and July 2024 alone. Behind each of those 1,350 numbers is a journalist, an activist, a parent, a survivor — someone who tried to show the world what was happening to them, and was silenced by a company that will report record profits in its next earnings call.

That is the Digital Iron Curtain. It is not maintained by a government. It is maintained by shareholders. The silencing documented in Chapter Nine operates at the level of human decision — a government request, a platform compliance, a post removed. Chapter Ten moves to a more disturbing form of control: the systematic manipulation of what AI systems themselves believe, and therefore what the people who consult those systems are told is true. What governments and their contractors cannot remove from the record, they are now attempting to train out of the machines that compose the record.

CHAPTER TEN — GAMING THE MACHINE

How Governments Use AI to Manufacture Reality — and How the AI Industry Helps Them

There is a question that gets asked less often than it should: if Big Tech's platforms have been suppressing Palestinian voices, what has been filling the space those voices left behind? The answer is not silence. It is manufactured content, coordinated influence operations, and a systematic campaign to train the AI systems that hundreds of millions of people now use as their primary source of information — to produce the answers that governments want them to produce.

This chapter is about how that works. It is about who is paying for it, who is building it, and what it means for the future of truth itself.

The $6 Million Operation: Gaming ChatGPT Through the Back Door

In September 2025, newly filed documents with the U.S. Department of Justice under the Foreign Agents Registration Act confirmed what many had suspected but could not yet prove. The Israeli government — operating through its Foreign Ministry and a $145 million public diplomacy initiative called Project 545 — had hired a network of American firms to conduct one of the most ambitious AI influence operations ever documented.

At the center of the operation stands Brad Parscale — the digital architect of Donald Trump's 2016 presidential campaign, the man who hired the now-notorious data firm Cambridge Analytica, and now the chief strategy officer of Salem Media Group, a conservative Christian broadcasting network whose stakeholders include Donald Trump Jr. and Lara Trump. Parscale's firm, Clock Tower X LLC, registered as an agent of the Israeli government on September 18, 2025, under a $6 million contract from German media giant Havas Media Network, which was itself operating on behalf of Israel's Ministry of Foreign Affairs.

The goal was explicit and the method was new. Clock Tower would produce a minimum of 100 "root creative assets" per month — videos, graphics, audio, written content — with at least 80 percent targeted at Gen Z audiences across TikTok, Instagram, YouTube, and podcasts. The target: 50 million paid impressions per month. But what made this operation different from conventional PR was one specific clause buried in the contract's statement of work. Clock Tower would deploy, in the filing's own language, "websites and content to deliver GPT framing results on GPT conversations." (Source: Responsible Statecraft, "Israel Wants to Train ChatGPT to Be More Pro-Israel," October 2025, https://responsiblestatecraft.org/israel-chatgpt/; Sludge, "Israel's New U.S. Influence Campaigns Target TikTok, Churches, and ChatGPT," October 2025, https://readsludge.com/2025/10/06/israels-new-u-s-influence-campaigns-target-tiktok-churches- and-chatgpt/)

Let me explain what that means in plain language. AI systems like ChatGPT are trained on vast amounts of text from across the internet. They learn what is true, what is contested, and what framing to apply to complex topics by processing billions of web pages, articles, and social media posts. If you flood the internet with a particular version of events — through hundreds of

websites, millions of social media posts, and coordinated content designed to rank highly in search engines — you are not just shaping public opinion directly. You are shaping what AI learns. You are, in effect, poisoning the well from which the next generation's information will be drawn. Clock Tower would also use a predictive AI platform called MarketBrew to optimize content for algorithmic visibility on Google and Bing — ensuring that the pro-Israel narrative would be what both human users and AI training systems encountered first, and encountered most.

This is not a direct payment to OpenAI. Let me be precise on that point, because the rumor circulating online — that Israel paid OpenAI directly — is inaccurate, and this book deals in documented facts. Snopes investigated and confirmed: there is no evidence of a contract between Israel and OpenAI. What is documented, in legally filed government disclosures, is something more sophisticated and more difficult to challenge: a government-funded operation to game the training environment of AI itself. (Source: Snopes, "Israel Didn't Sign a Deal With ChatGPT. Here's What Actually Happened," January 2026, https://www.snopes.com/news/2025/12/22/israel-contract-chatgpt/; Business and Human Rights Resource Centre, https://www.business-human-rights.org/en/latest-news/israelopt-clock-tower-x-allegedly-hired-by- israeli-government-to-create-pro-israel-content-training-chatgpt-and-targeting-facebook-instagra m-tiktok-you-tube-incl-cos-response-and-non-responses/)

When the Business and Human Rights Resource Centre contacted OpenAI, Meta, Google, Alphabet, and Havas Media requesting comment on these allegations, only TikTok responded. OpenAI said nothing. That silence is its own answer.

OpenAI and the Military Contract That Wasn't — Until It Was

There is a second story about OpenAI and Israel that has nothing to do with influence operations. It is about weapons. And it is documented by the Associated Press.

In February 2025, the AP published an investigation based on internal company documents, data, and exclusive interviews with current and former Israeli and company officials. The findings were stark. Following Hamas's October 7, 2023 operation, Israel's use of OpenAI technology — purchased through Microsoft's Azure cloud platform — spiked to nearly 200 times its pre-genocide level. The amount of data the Israeli military stored on Microsoft servers doubled in the first months of the genocide alone, reaching more than 13.6 petabytes by July 2024 — roughly 350 times the digital memory needed to store every book in the Library of Congress. OpenAI's AI models were being used directly in the intelligence and targeting pipeline of an active military campaign. (Source: Associated Press, "As Israel Uses US-Made AI Models in War, Concerns Arise About Tech's Role in Who Lives and Who Dies," February 2025; Fortune, https://fortune.com/2025/02/19/israel-microsoft-openai-raises-questions-powerful-tech/)

"This is the first confirmation we have gotten that commercial AI models are directly being used in warfare," said Heidy Khlaaf, chief AI scientist at the AI Now Institute and a former senior safety engineer at OpenAI herself.

OpenAI's response to this revelation was to claim it had no direct partnership with the Israeli military. That claim deserves scrutiny. In early 2024 — quietly, without announcement — OpenAI rewrote its terms of service to remove language that explicitly prohibited military use of its technology. The ban on military applications was simply deleted. What had been a line in the

sand became an open door. And through that door walked the Israeli Military Forces, purchasing OpenAI models through Microsoft's Azure platform, using them to process surveillance intelligence and support targeting decisions. The company changed the rules, and then pointed to the rules as proof of innocence. That is not a defense. That is a business strategy.

Google's Gemini: The Whistleblower's Evidence

Google has consistently maintained that its AI products are not used for weapons or surveillance in violation of its ethics policies. In February 2026, a former Google employee filed a confidential whistleblower complaint with the U.S. Securities and Exchange Commission telling a different story.

According to the complaint — reviewed by The Washington Post — in 2024, a customer support request arrived at Google's cloud computing division from an address linked to an Israeli Military Forces account. The customer was attempting to use Google's Gemini AI to analyze aerial drone surveillance footage and had encountered a bug where the software was occasionally failing to identify drones, soldiers, and other objects in the footage. Google's cloud unit customer support staff responded with technical assistance and conducted internal tests to help resolve the issue. A second Google staffer — who allegedly worked on the Israeli Military's Google Cloud account — was copied on the support ticket. (Source: Washington Post, "Whistleblower Says Israeli Military Contractor Used Google's Gemini AI," February 2026, https://www.washingtonpost.com/technology/2026/02/01/google-ai-israel-military/)

The whistleblower alleged that this assistance violated Google's own AI principles, which at the time stated the company would not deploy AI for weapons applications or surveillance violating internationally accepted norms. Google disputed the characterization, claiming the support ticket was too small in commercial value to constitute "meaningful" AI usage. That is a remarkable defense. The question of whether a drone surveillance tool violates ethics does not change based on the size of the monthly billing statement.

Separately, documents from November 2024 show the Israeli military actively requesting access to Google's Gemini AI platform to develop its own AI assistant capable of processing military documents and audio recordings — expanding its capabilities months into an active genocide. Records show Google employees noticed a surge in the Israeli Military's access requests for AI tools after October 7, with internal documents revealing that one employee warned that delaying expanded access might push Israel toward Amazon, Google's cloud competitor. (Source: The Defense Post, "Google Supporting Israel With AI Tools for Military Use," February 2025, https://thedefensepost.com/2025/02/04/google-israel-ai-military/; Common Dreams, https://www.commondreams.org/news/big-tech-gaza-genocide)

Google's response — like OpenAI's, like Microsoft's — follows the same pattern we have documented throughout this book. Deny. Minimize. Update the policy to close the gap between what was done and what was permitted. Repeat.

The Pattern Behind the Pattern

What emerges from these three stories — the Clock Tower influence operation, the OpenAI military usage spike, and the Gemini whistleblower complaint — is not three separate scandals. It is one coherent strategy, operating simultaneously on two fronts.

On the military front, AI companies provide the tools that make targeting faster, more automated, and more lethal. They do this through direct contracts, through cloud platforms, and through the quiet removal of the ethical guardrails that once made such use impermissible.

On the information front, governments use AI's own training process as a weapon — flooding the internet with manufactured content designed to teach AI systems what to believe, what to suppress, and whose suffering to treat as legitimate. The result is an information environment in which the same AI tools that help identify bombing targets are simultaneously being trained to describe those bombings in language favorable to the people ordering them.

This is the full architecture of AI-enabled control. It does not only kill. It also shapes what you believe about the killing. And both functions are being provided, knowingly or not, by the same companies — OpenAI, Google, Microsoft — whose quarterly earnings reports you will read next week in the business pages of your newspaper of record, between the advertisements.

OpenAI was asked to comment on the influence operation designed to game its training data. It did not respond. That silence tells you everything you need to know about whose interests these companies have decided to serve.

--- If Chapter Ten is about how AI is trained to manufacture consensus for adults, Chapter Eleven is about what happens to the children who grow up inside platforms designed not to inform but to addict. The algorithm that rewards extremist content for adults is the same algorithm that recommends self-harm content to teenagers. The platform that censors Palestinian voices is the same platform that the New Mexico jury found had knowingly harmed children's mental health. These are not separate stories. They are the same business model applied to two different vulnerable populations.

CHAPTER ELEVEN — THE ALGORITHM AGAINST OUR CHILDREN

Social Media, Teen Suicide, and the Industry That Profits From Both

In March 2026, after a nearly seven-week trial in Santa Fe, New Mexico, a jury returned a verdict that should have broken open every boardroom in Silicon Valley. The jury found that Meta — the company that owns Instagram, Facebook, and WhatsApp — had knowingly harmed children's mental health, concealed what it knew about child sexual exploitation on its platforms, and engaged in "unconscionable" trade practices that unfairly exploited the vulnerabilities of young people. The penalty: $375 million. Juror Linda Payton, 38, explained what guided the jury's judgment. With a maximum $5,000 per violation at stake, she said she thought each child was worth the maximum amount. (Source: NPR, "New Mexico jury says Meta harms children's mental health and safety, violating state law," March 24, 2026, https://www.npr.org/2026/03/24/g-s1-115019/new-mexico-meta-children-mental-health; AP/US News, March 24, 2026, https://www.usnews.com/news/technology/articles/2026-03-24/new-mexico-jury-finds-meta-violated-consumer-protection-law-at-trial-about-child-safety)

The trial examined a substantial body of Meta's own internal documents and research — evidence that the company's own engineers and scientists had conducted studies on the impact of its platforms on young users and had not acted on what they found. The jury heard from Meta executives, platform engineers, former employees who had become whistleblowers, psychiatric experts, tech safety consultants, and public school educators who described classrooms disrupted by sextortion schemes targeting children. It heard about Meta's documented failure to enforce its own ban on users under the age of 13, and about the role its algorithms play in prioritizing sensational and harmful content — including content about teen suicide.

At the core of the case was a charge that New Mexico Attorney General Raúl Torrez pressed with moral clarity: Meta designed its platforms to addict young users. It knew what was happening. It chose profit. "We know the output is meant to be engagement and time spent for kids," prosecution attorney Linda Singer told jurors. "That choice that Meta made has profound negative impacts on kids."

More than 40 state attorneys general across the country have filed similar lawsuits against Meta. Thousands of families are in litigation. The coalition of parents who have lost children to harms caused by social media — the group calling itself ParentsSOS — called the New Mexico verdict "a watershed moment." "We parents who have experienced the unimaginable — the death of a child because of social media harms — applaud this rare and momentous milestone," the group said in a statement. (Source: PBS NewsHour, "What's next in social media legal battles after a New Mexico jury finds Meta platforms harm children," March 25, 2026, https://www.pbs.org/newshour/nation/whats-next-in-social-media-legal-battles-after-a-new-mexic o-jury-finds-meta-platforms-harm-children)

A Letter That Changed a Country

A twelve-year-old girl in Australia named Charlotte O'Brien took her own life because of bullying that was enabled and amplified by social media. Her mother, Kelly O'Brien, wrote a personal letter to Prime Minister Anthony Albanese. He read it. It moved him. And it helped move a government.

Charlotte's story was not an isolated tragedy. It was the story that Australia's news organizations — and a campaign called Let Them Be Kids, backed by more than 54,000 petition signatures from parents, educators, and child safety advocates — used to make the scale of the harm visible to the people with the power to act. An event at the United Nations General Assembly in September 2025 featured a mother speaking about her daughter's suicide as "death by bullying... enabled by social media." The speech drew support from world leaders from Greece, Fiji, Tonga, and European Commission President Ursula von der Leyen.

On November 29, 2024, the Australian Parliament passed the Online Safety Amendment (Social Media Minimum Age) Act. It came into full effect on December 10, 2025. Australia became the first country in the world to enforce a nationwide ban preventing children under the age of 16 from holding accounts on major social media platforms — including Facebook, Instagram, TikTok, YouTube, Snapchat, Reddit, X, Threads, Twitch, and Kick. Platforms that fail to take "reasonable steps" to enforce the age restriction face fines of up to $49.5 million Australian dollars. (Source: Wikipedia, Online Safety Amendment (Social Media Minimum Age) Act 2024, https://en.wikipedia.org/wiki/Online_Safety_Amendment; eSafety Commissioner, "Social media age restrictions," March 2026, https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions; Time, "What to Know About Australia's Social Media Ban for Kids Under 16," December 10, 2025, https://time.com/7339762/australia-youth-social-media-ban-under-16-snapchat-meta-tiktok/)

"This is the day when Australian families are taking back power from these big tech companies," Prime Minister Albanese said when the law took effect. "They are asserting the right of kids to be kids."

Jonathan Haidt, social psychologist and author of The Anxious Generation — a book that documented with systematic rigor the relationship between the mass adoption of smartphones, social media, and the mental health crisis now consuming an entire generation — welcomed Australia's action. "The world is rooting for your success," he wrote, "and many other nations will follow."

They are already following. As of February 2026, France, the United Kingdom, Germany, Italy, Greece, Spain, Denmark, Malaysia, and New Zealand are all considering or actively advancing similar restrictions for children under 16. The European Parliament passed a non-binding resolution advocating a minimum age of 16, and has proposed banning addictive features like infinite scrolling and autoplay for minors entirely. (Source: CNBC, "Australia is trying to enforce the first teen social media ban. Governments worldwide are watching," December 10, 2025, https://www.cnbc.com/2025/12/10/australia-16-year-old-teens-ban-social-media-policy-law-ig-tikt ok-fb-reddit-youtube-snapchat.html; Australia's Social Media Ban: A Test for Global Digital Governance, BISI, January 2026, https://bisi.org.uk/reports/australias-world-first-social-media-ban-a-test-for-global-digital-governa nce)

The Privacy Trap Hidden Inside the Protection

I want to say something honest about the age restriction approach — because I believe the governments moving in this direction are doing the right thing for the right reasons, and I also believe that the implementation they are being pushed toward carries its own serious dangers.

Age verification, as it is currently being discussed and deployed, requires social media platforms to confirm a user's age before allowing account creation. The methods under consideration include government ID checks, biometric facial scanning, behavioral inference from browsing history, and AI-based age estimation from physical appearance. Every single one of these methods requires the collection of sensitive personal data. Some require biometric data. Some require linking social media accounts to government identity documents. Some involve building a digital profile of a user's behavior, demographic characteristics, and physical appearance — and retaining it.

The same companies I have documented throughout this book as serial violators of user privacy — companies that added "terrorist" to Palestinian users' bios by algorithm, that scanned 6,167 browser extensions without user knowledge, that shared data with foreign governments at a compliance rate of 94 percent — these are the companies now being positioned as the enforcers of children's age verification. The data they will collect to enforce the restriction is exactly the kind of data they have historically misused, monetized, and handed over to governments on request.

Mozilla, the non-profit organization that builds the Firefox browser and has spent decades defending user privacy, raised this concern directly in response to the Australian law. Age estimation systems based on biometric data, Mozilla noted, often underperform for people with darker skin tones and non-binary facial features. Behavioral inference systems built to estimate age based on browsing history create detailed profiles of every user. Age verification via government ID creates a linkage between social media identity and legal identity that has never before existed at scale. (Source: Mozilla Foundation, "Australia's Social Media Ban: Why Age Limits Won't Fix What Is Wrong With Online Platforms," December 19, 2025, https://blog.mozilla.org/netpolicy/2025/12/19/australias-social-media-ban-why-age-limits-wont-fix -what-is-wrong-with-online-platforms/)

Let me be direct about what this means. The companies whose platforms are harming children cannot be the companies trusted to verify the identities of those children in order to protect them from their own platforms. That is not protection. That is an expansion of the surveillance infrastructure under the guise of safety — the same move Big Tech has made with security, with content moderation, with anti-terrorism measures, every time the public has demanded accountability.

The protection of children from the harms of social media is a legitimate and urgent moral priority. The mechanism of that protection must not create a new and larger harm in its place. Any age verification framework must be independent of the platforms, privacy-preserving by design, subject to public audit, and built on the principle that compliance with a child protection law does not entitle a company to build a biometric database of every adult who uses its service.

The Algorithm That Rewards Rage — and the Children Who Pay for It

Let me now turn to something that is less discussed than the explicit harms of social media to children, but which I believe is equally consequential for the generation growing up inside these

platforms: the deliberate design of recommendation algorithms that reward the most extreme, divisive, and inflammatory content — and the financial incentives that keep that design in place.

In 2025, Sky News conducted a month-long experiment on X — the platform formerly known as Twitter, now owned by Elon Musk. The study created accounts with different political leanings and tracked what content was algorithmically recommended. The finding was unambiguous: X "exhibits a significant algorithmic bias toward promoting far-right and extreme content." The Wall Street Journal and the Washington Post conducted their own separate analyses. Both reached the same conclusion: X now amplifies right-leaning and extreme political content over other perspectives. (Source: Social Media Today, "Research Shows X Amplifies Conservative Political Views," February 19, 2026, https://www.socialmediatoday.com/news/x-formerly-twitter-amplifies-conservative-political-persp ectives-report/812652/)

A 2025 academic study published in Science used a browser extension methodology to alter users' X feeds during the 2024 U.S. presidential election — reducing exposure to content expressing anti-democratic attitudes and partisan hostility. The study found this single change shifted users' political feelings toward their opponents by more than two points on a 100-point scale — a shift that would normally take approximately three years to occur organically in the general population. The researchers described this as causal evidence that X's algorithm — its recommendation system, its financial incentive structure — is actively contributing to political polarization. (Source: The Conversation, "A few weeks of X's algorithm can make you more right-wing — and it doesn't wear off quickly," February 18, 2026, https://theconversation.com/a-few-weeks-of-xs-algorithm-can-make-you-more-right-wing-and-it- doesnt-wear-off-quickly-276153)

This is not an accident. It is an engineering decision. The algorithms on X, TikTok, Facebook, and YouTube are not neutral filters. They are optimization machines, designed to maximize engagement — and the content that drives engagement most reliably is the content that provokes the strongest emotional reactions: outrage, fear, contempt, tribal identity, moral panic. That is what keeps people scrolling. That is what keeps the session length high. That is what makes the advertising inventory valuable. And that is what is being served, disproportionately, to teenagers who are still forming their understanding of the world, their sense of identity, their relationship with people who are different from them.

The financial incentive structure compounds this. Platforms financially reward the creators who generate the most engagement — through monetization programs that pay per view, per watch hour, per share. A creator who produces content that provokes rage is not just getting more views. They are getting paid more for producing that content. The algorithm selects for the most inflammatory content, the most polarizing framing, the most extreme positions — and then writes a check to whoever produced them. The far-right content farms that now operate at industrial scale on X and TikTok are not producing what they produce because they believe in it. They are producing it because the algorithm has made it the most profitable content category available to a creator without significant production resources.

What this means for a fifteen-year-old spending four hours a day on TikTok or X — seeing that content, being financially excluded from the social spaces where their peers organize — is a daily immersion in a curated environment that makes the world seem angrier, more divided,

more dangerous, and more certain than it actually is. That is not an accident either. It is the product.

Silencing Palestine, Amplifying Extremism: The Algorithm's Double Standard

I documented in Chapter Nine the specific mechanisms by which Meta, TikTok, X, YouTube, and LinkedIn systematically suppressed Palestinian voices during the Gaza genocide — deleting posts, suspending accounts, complying with Israeli government takedown requests at rates as high as 94 percent, and erasing 700 videos documenting potential war crimes at the direction of the U.S. government.

I want to return to that record here, in the context of what these platforms have simultaneously allowed. While Palestinian journalists had their accounts suspended, their posts deleted, and their reach algorithmically throttled, the same platforms left anti-Palestinian hate speech online at a rate of 96 percent, according to the Centre for Countering Digital Hate. The 7amleh digital rights organization documented over 19,000 cases of hate speech and inciting content in Hebrew on X in the weeks following October 7, 2023 — content that remained online without consequence. (Source: Business and Human Rights Centre, "Report Exposes X's Failure to Remove 96% of Hate Speech Posts," November 2023, https://www.business-humanrights.org/en/latest-news/report-exposes-xs-failure-to-remove-96-of -hate-speech-posts-amid-israel-palestine-conflict/)

The pattern is consistent across platforms and consistent over time. The algorithm suppresses the voices of the occupied and amplifies the language of the occupier. It silences grief and rewards rage. It removes documentation of atrocity and recommends the rhetoric that justifies it. And it does this not because of an ideology embedded in the software — but because of the financial architecture that the software is designed to optimize. Content that generates outrage gets more engagement. More engagement generates more advertising revenue. The atrocities are deleted because someone with government connections asked for them to be deleted. The rage is amplified because rage is what the business model runs on.

The same algorithm that suppresses the Palestinian child's testimony keeps the extremist content creator in the recommendation queue — because the extremist content creator generates more engagement, and engagement is money.

Our children are being raised inside this architecture. They are learning what the world is from platforms that are financially incentivized to show them the most extreme version of it. They are learning what counts as a human being worth grieving from platforms that delete some suffering and amplify other suffering based on which deletion and which amplification produces the most profitable outcome.

And when the consequences of that architecture become visible — when the clinical studies document the mental health damage, when the suicides happen, when the New Mexico jury finds that Meta engaged in unconscionable practices that exploited the vulnerabilities of children — the companies issue statements, update their community guidelines, and return to the same practice at the same scale in the next news cycle.

What Regulation Must Actually Require

The age restriction laws moving through parliaments from Canberra to Copenhagen to Brussels are a genuine and necessary beginning. They are not sufficient.

Protecting children from social media harms requires going beyond restricting who can hold an account. It requires addressing the algorithmic and financial architecture that makes those platforms harmful in the first place — because a sixteen-year-old is no less vulnerable to an algorithm optimized for outrage than a fourteen-year-old, and an adult who spends four hours a day in a recommendation system designed to maximize engagement and emotional reaction is being harmed too. The age restriction addresses the symptom — underage users — without touching the disease — the design.

What is actually needed, alongside the age restriction, is mandatory algorithmic transparency — the legal requirement that platforms disclose, to independent auditors with binding authority, precisely how their recommendation systems work and what they optimize for. Not a quarterly transparency report written by the company's PR department. An independent, binding audit with the power to mandate changes to algorithmic design when the evidence shows those systems are causing harm.

What is needed is liability for algorithmic amplification of harmful content. The Section 230 shield that protects technology companies from liability for content posted by users should not extend to content that the company's own recommendation algorithm amplified, promoted, and financially rewarded. When a platform pays a creator to produce content that drives a teenager toward self-harm, the platform is not a passive host. It is a participant. The law should treat it as one.

What is needed is a global ban on the algorithmic and financial reward structures that incentivize the most extreme content. The business model that converts engagement into revenue — and that therefore selects for the content that generates the most powerful emotional reactions — is the engine of both political radicalization and the mental health crisis in our children. You cannot regulate your way out of the consequences of that engine without addressing the engine itself.

And what is needed is the recognition, at the level of binding international law, that the platforms that operate across borders cannot be governed by the standards of the country most willing to look the other way. Australia moved first. Dozens of countries are following. The pressure on the European Union to enforce the Digital Services Act to its full legal capacity — holding platforms liable for algorithmic harm with the same urgency with which it holds them liable for privacy violations — is real, growing, and long overdue.

Charlotte O'Brien was twelve years old. The New Mexico jury found that each child harmed was worth the maximum penalty available under the law. So do I.

--- The harm I have documented in Chapter Eleven — to children, to their mental health, to their developing understanding of the world — is harm caused by platforms whose most dangerous capabilities are already visible and regulated. Or trying to be regulated. What Chapter Twelve describes is a category of AI capability that is not yet regulated, that the companies building it acknowledge could be catastrophic, and that is proliferating faster than any governance framework can keep up with. The Code Bomb is not a metaphor. It is a technical description of what Anthropic has already built and what every major AI laboratory is racing to match.

CHAPTER TWELVE — THE CODE BOMB

Claude Mythos, Project Glasswing, and the Unregulated Frontier AI Race That Could End Everything

I need to tell you something that is not easy to write.

I have spent the previous nine chapters documenting the crimes of Amazon, Google, Microsoft, Oracle, and Palantir. I have named their contracts. I have named their executives. I have named their systems — Gospel, Lavender, Where's Daddy, Blue Wolf — and the families whose deaths those systems enabled. I have documented how their platforms censored the survivors and manufactured the narrative.

But there is something I have deliberately held until now. Something I have been sitting with since April 7, 2026, when Anthropic — the company that built the AI I have been using as my editor and research partner throughout this book — announced something that I believe represents one of the most dangerous moments in the history of technology.

It is called Claude Mythos. The restricted program through which it is being deployed is called Project Glasswing.

And I have a responsibility to tell you what it means — not because I want to condemn the people at Anthropic who I believe, genuinely, are attempting something different from the Palantirs and the Gospels of the world. But because good intentions and catastrophic consequences have never been mutually exclusive. The men who built the atomic bomb were not, for the most part, evil men. They were brilliant men who believed, deeply, that their intentions were protective. And the world has spent the eighty years since their success living with what they created.

We cannot afford to spend another eighty years learning the same lesson about AI.

What Anthropic Just Admitted to the World

On April 7, 2026, Anthropic announced the preview release of its most advanced model to date: Claude Mythos. What followed was not a product launch in any ordinary sense. It was closer to a public confession.

In a blog post on its Frontier Red Team page, Anthropic disclosed that Mythos Preview had "fully autonomously identified and then exploited a 17-year-old remote code execution vulnerability in FreeBSD" — a widely used operating system that underpins servers and infrastructure across the world. The exploit allowed anyone to gain complete root access to an affected machine from anywhere on the internet, with no authentication required. No human was involved after the initial request was made. The machine found it. The machine exploited it. Start to finish. (Source: Anthropic Frontier Red Team Blog, "Claude Mythos Preview," April 7, 2026, https://red.anthropic.com/2026/mythos-preview/)

But that was not the worst of it. According to Anthropic's own announcement, Mythos Preview found vulnerabilities in every major operating system and every major web browser. It identified thousands of zero-day vulnerabilities — flaws previously unknown to the software's

developers — in just a few weeks of internal testing. More than ninety-nine percent of those vulnerabilities remain unpatched, because disclosing them publicly before a fix exists would be, as Anthropic put it, "irresponsible." (Source: Anthropic, "Project Glasswing," April 7, 2026, https://www.anthropic.com/project/glasswing)

Anthropic also issued a private warning to top U.S. government officials: Mythos makes large-scale cyberattacks significantly more likely this year. (Source: Fortune, "Anthropic is giving some firms early access to Claude Mythos to bolster cybersecurity defenses," April 7, 2026, https://fortune.com/2026/04/07/anthropic-claude-mythos-model-project-glasswing-cybersecurity/)

Read that again. A company built a tool. The tool is — in the company's own words — "currently far ahead of any other AI model in cyber capabilities." That same company then warned the government that the tool makes civilization-scale cyberattacks more likely. And then they released it — to fifty organizations including Amazon Web Services, Apple, Google, Microsoft, Nvidia, JPMorganChase, Broadcom, Cisco, CrowdStrike, and Palo Alto Networks — in a restricted program called Project Glasswing.

Anthropic committed one hundred million dollars in usage credits to make this happen. They called it "an urgent attempt to put these capabilities to work for defensive purposes." (Source: Anthropic, "Project Glasswing," https://www.anthropic.com/project/glasswing)

I want to be careful about what I say next, because it matters.

I believe Anthropic is making a genuine attempt to do the right thing. Project Glasswing is structured around giving defenders a head start — sharing vulnerabilities with the people responsible for the infrastructure that billions of people depend on, so those vulnerabilities can be patched before adversaries develop the same capability independently. That logic is not without merit. That intention is not without honesty.

But I have spent enough time inside the technology industry to know that good intentions, institutional pressure, financial incentives, and the speed of progress have a way of diverging — fast. Anthropic's own security experts estimate that similar capabilities will proliferate from other AI laboratories within six to eighteen months. (Source: ArmorCode, "Anthropic's Claude Mythos and What It Means for Security," May 2026, https://www.armorcode.com/blog/anthropics-claude-mythos-and-what-it-means-for-security/)

Six to eighteen months. And they have already released this to Amazon, Google, and Microsoft — the same companies I spent eight chapters documenting as the central infrastructure providers for a military targeting system that has killed tens of thousands of people.

I hold that tension openly. I do not know how to resolve it except to say what I have been building toward since the first page of this book: the decisions about whether, when, and how to release tools of this power cannot be made by private companies alone. They cannot be made by boards and CEOs and partnership programs and hundred-million-dollar usage credit commitments. They must be made — openly, accountably, with binding legal authority — by the entire human community that will live or die with the consequences.

The Pattern Across the Industry

Anthropic is not alone in this race. It would be dishonest to single them out without documenting what every other major AI laboratory is building and doing.

OpenAI — the company whose CEO Sam Altman has done more than any single individual to accelerate the commercial deployment of frontier AI — warned in December 2025 that its next-generation models pose a "high" cybersecurity risk and may be capable of developing "functional zero-day remote exploits against well-protected computer systems." The company's GPT-5.1-Codex-Max model scored 76 percent on standard capture-the-flag cybersecurity challenges in November 2025, up from 27 percent just three months earlier. The speed of that improvement should stop everyone reading this cold. That is not incremental progress. That is a capability explosion. (Source: Parameter, "OpenAI Warns New AI Models Could Enable Zero-Day Exploits and Complex Cyberattacks," December 11, 2025, https://parameter.io/openai-warns-new-ai-models-could-enable-zero-day-exploits-and-complex-c yberattacks/)

OpenAI responded to this by launching a tool called Aardvark — an AI security system designed to scan code for vulnerabilities — and establishing a Frontier Risk Council. Defensive responses to offensive capabilities they are simultaneously building, deploying, and accelerating. This is the paradox at the heart of the entire industry.

Google's Gemini is not a targeted cybersecurity model, and yet it is already being weaponized at scale by nation-state hackers. In February 2026, Google's own threat intelligence researchers confirmed that state-sponsored groups from China, Russia, and Iran are using Gemini in "all stages" of attacks — reconnaissance, vulnerability analysis, exploit development, and phishing. Chinese groups including APT31 and APT41 used Gemini to analyze known vulnerabilities, troubleshoot exploit code, and build offensive scanning tools. Iranian group APT42 used it to develop Python-based surveillance tools and examine exploitation paths for newly disclosed vulnerabilities. (Source: Tom's Hardware, "Google Reports That State Hackers from China, Russia and Iran Are Using Gemini in 'All Stages' of Attacks," February 13, 2026, https://www.tom-shardware.com/tech-industry/cyber-security/google-reports-that-state-hackers-fr om-china-rus-sia-and-iran-are-using-gemini-in-all-stages-of-attacks-phishing-lures-coding-and-vu lnerabil-ity-testing-get-ai-underpinnings-from-hostile-actors; Infosecurity Magazine, February 12, 2026, https://www.infosecurity-magazine.com/news/nation-state-hackers-gemini-ai/)

And then there is Grok — Elon Musk's AI platform, which I documented in Chapter Four. Grok has already demonstrated its enthusiasm for producing antisemitic content, Holocaust denial, and racist conspiracy theories. Now consider what it means to have the same architecture, the same foundational technology, the same agentic capabilities being pushed toward cybersecurity applications — in the hands of a man who calls empathy a weakness, who is simultaneously advising the U.S. federal government on technology deployment, and who has shown no interest in the kind of careful, governed restraint that even Anthropic — for all the tensions I have named — is attempting to exercise.

The Dragon at the Back Door: China's AI Race Has No Guardrails

Everything I have described so far operates within the framework of companies that are — at minimum — legally accountable to democratic governments, subject to litigation, and operating in environments where employees can walk out in protest and journalists can file requests under freedom of information laws.

None of that is true for China.

In January 2025, a Chinese AI laboratory called DeepSeek released an open-source model that sent the global technology industry into a panic. Not because it was dangerous in the way Mythos is dangerous — but because it was nearly as powerful as OpenAI's leading models, at a fraction of the cost, and it was open source. Anyone could download it. Anyone could modify it. Anyone could weaponize it. And the Chinese Communist Party, whose laws require that AI-generated content reflect "core socialist values," support "correct political direction," and avoid material that could "undermine state power," has direct and legal authority over everything DeepSeek does. (Source: House Select Committee on China, "DeepSeek: A Report," 2025, https://chinaselectcommittee.house.gov/sites/evo-subsites/selectcommitteeontheccp.house.gov/ files/evo-media-document/DeepSeek%20Final.pdf)

Cybersecurity firm Feroot Security uncovered hidden code in DeepSeek's browser application capable of transmitting user data directly to servers controlled by China Mobile — a state-owned telecom company previously delisted from the New York Stock Exchange over national security concerns. Former Homeland Security official Stewart Baker put it plainly: "It raises all of the TikTok concerns plus you're talking about information that is highly likely to be of more national security and personal significance than anything people do on TikTok." (Source: Feroot Security, "DeepSeek's Hidden Code Sending User Data to China," January 2025, https://www.feroot.com/news/the-independent-feroot-security-uncovers-deepseeks-hidden-code -sending-user-data-to-china/)

The implications are not theoretical. In November 2025, Anthropic reported that threat actors connected to a Chinese state-sponsored organization had exploited its Claude Code tool for an AI-powered espionage campaign — the first documented AI-powered malware attack against a frontier AI system. (Source: Parameter, December 2025, https://parameter.io/openai-warns-new-ai-models-could-enable-zero-day-exploits-and-complex-c yberattacks/) Anthropic interrupted it — that time. The next campaign will be better prepared.

Now understand what we are looking at in its totality. The United States is building frontier AI models — through Anthropic, OpenAI, and Google — with autonomous cybersecurity capabilities that their own creators describe as potentially civilization-threatening. China is building its own models, extracting training data from American systems through what the White House has called "deliberate, industrial-scale campaigns," and deploying state-sponsored hackers who use American AI tools — Gemini, among others — to plan attacks on American targets. (Source: Nextgov/FCW, "White House Accuses China of 'Deliberate, Industrial-Scale Campaigns' to Steal US AI Models," April 2026, https://www.nextgov.com/artificial-intelligence/2026/04/white-house-accuses-china-deliberate-in dustrial-scale-campaigns-steal-us-ai-models/413083/) And both sides are racing — driven by private profit incentives on the American side and geopolitical ambition on the Chinese side — with no binding international framework, no shared safety standards, no independent verification body, and no mechanism for either side to pause when the risk calculation demands it.

This is not a technology competition. This is a new arms race. And unlike the nuclear arms race, there is no Mutually Assured Destruction doctrine that creates even the perverse incentive of restraint. In the nuclear era, both sides knew that launching meant annihilation. In the AI era, a sufficiently capable offensive tool is invisible, deniable, scalable, and can be deployed against

infrastructure — hospitals, power grids, water treatment facilities, banking networks — without a single missile leaving the ground.

What This Means for the People No One Is Protecting

I want to step away from the geopolitical frame for a moment and talk about who actually pays the price when these tools fall into the wrong hands or proliferate beyond the controlled environments that Project Glasswing is attempting to maintain.

It is not Amazon. Amazon has the resources to deploy Claude Mythos Preview defensively, scan its own codebases, and patch its vulnerabilities before adversaries can exploit them. Amazon is a Project Glasswing launch partner. Amazon has more than 400 trillion network flows analyzed per day for threats. Amazon will be fine.

The people who will not be fine are the small hospital in rural Kansas that runs on a fifteen-year-old electronic health records system maintained by two IT staff members. The community bank in rural Georgia whose entire cybersecurity infrastructure is a subscription to an antivirus service and a firewall that has not been updated in three years. The water treatment facility in a mid-sized city that digitized its monitoring systems a decade ago and has had no meaningful security audit since. The small business owner who runs her entire operation on QuickBooks, whose customer credit card data is protected by nothing more sophisticated than a password she has not changed in two years.

These are not hypothetical people. They are the majority of the organizations that constitute the actual fabric of the economy — the small and medium businesses that employ more than half of all private-sector workers in the United States, and the equivalent across every country on earth.

An AI model with the capabilities of Claude Mythos Preview — or whatever China builds in the next eighteen months that achieves comparable performance — does not need to be deployed by a nation-state to be catastrophic. It needs to be deployed by a criminal organization. By a terrorist cell. By a single technically sophisticated individual who downloads an open-source equivalent from a repository that no governance framework has yet found a way to regulate. Every major operating system. Every major browser. Thousands of zero-day vulnerabilities, catalogued and exploitable at a pace and scale that human defenders have never encountered before.

Healthcare systems are the most vulnerable. A hospital that cannot access its patient records — because they have been encrypted by ransomware delivered through an AI-discovered zero-day vulnerability — cannot operate. Surgeries are cancelled. ICU monitoring is disrupted. Drug dosage records become inaccessible. People die. This is not speculation: ransomware attacks on hospitals have already produced documented patient deaths in Germany, in the United States, and elsewhere. That was ransomware built by human criminal organizations working at human speed. What happens when the same capability is available to anyone who can type a prompt?

Power grids are the second most vulnerable. Industrial control systems — the software that manages electricity generation and distribution — were designed for reliability and longevity, not security. Many of the systems running the electrical infrastructure of mid-sized American cities are running software from the early 2000s, on hardware that was never designed to be internet-connected but became so because it was cheaper and more convenient. An AI that can

autonomously discover exploits in major operating systems can find the vulnerabilities in those legacy systems with equal facility.

Banking is the third. Not the major banks — JPMorganChase is a Project Glasswing partner and has the resources to deploy Mythos defensively. The community banks, the credit unions, the smaller financial institutions that hold the savings of working people who are not the clients of Goldman Sachs. When those institutions are compromised, their customers do not have the kind of insurance backstop that absorbs a hedge fund's losses. They have their mortgage payment. Their grocery money. Their children's college savings.

And for sovereign nations in the Global South — countries without the IT infrastructure, the software development talent, or the security operations budget to defend against human-speed attacks, let alone AI-speed ones — the consequences are existential. A country's entire banking system. Its power grid. Its government communications. Its healthcare network. These can now be compromised, not by a sophisticated nation-state with years of preparation, but by a mid-level criminal organization with access to a sufficiently capable AI model.

Anthropic has committed $100 million to Project Glasswing. That $100 million reaches approximately 50 organizations. There are approximately 33 million small businesses in the United States alone. The mathematics of who is being protected and who is being left to absorb the consequences of this technology race could not be clearer.

Sam Altman Said the Quiet Part Out Loud

I have been building to this moment since I began this chapter. I need to tell you about a statement that Sam Altman made in 2015 — one that has circulated widely online but is almost always shared without the second half that makes it uniquely and unforgivably honest.

The first half of the quote: "I think AI will probably, most likely, sort of lead to the end of the world."

Now that is a startling thing for anyone to say. It is a more startling thing for the person building that world-ending technology to say. But the second half is what should keep every person of conscience awake at night.

The complete quote: "I think AI will probably, most likely, sort of lead to the end of the world. But in the meantime, there will be great companies created with serious machine learning."

(Source: Sam Altman, Y Combinator, 2015, as reported in Yahoo News, "Sam Altman's viral quote is missing key context," January 2026, https://www.yahoo.com/news/articles/think-ai-probably-lead-end-190517676.html)

Read those two sentences together. Sit with them. Feel the weight of the conjunction.

The man who went on to build ChatGPT — the product that has accelerated AI deployment across every industry on earth more than any single development since the internet — told the world in 2015 that his work would probably lead to the end of the world. And then, in the same breath, pivoted to the investment opportunity.

"But in the meantime, there will be great companies created with serious machine learning."

In the meantime.

I want you to understand what that phrase contains. It contains every Project Nimbus contract. Every Lavender kill list. Every Gospel bombing target. Every census database handed to

an occupying military. Every hospital ransomware attack that hasn't happened yet but will. Every water treatment facility that will go dark when the AI-discovered zero-day propagates faster than any human defender can respond.

All of that is "the meantime" between now and the end of the world that Sam Altman said will probably come. And the consolation he offers for that meantime is that great companies will be built.

I am an Iranian-American. I came to this country at eleven years old. I have spent my adult life believing, despite everything, in the possibility of building things that serve humanity. I have built companies. I have organized. I have written books and made films and built platforms for people to document their dead, and I have held on — stubbornly, through hospitalizations and lawsuits and every kind of institutional opposition — to the belief that technology can be made to serve human dignity rather than destroy it.

And I am telling you that "in the meantime, there will be great companies created" is not an acceptable answer to the question of whether AI will probably lead to the end of the world. It is not a response. It is a pivot. It is the most expensive pivot in human history, executed by a man who understood exactly what he was building and chose to build it anyway.

Elon Musk, who is no more reliable than his own chatbot, put a number on it — 20 percent chance of human annihilation by AI. (Source: Time Magazine, "How Musk, Altman, AI Leaders Face the 'Oppenheimer Moment,'" March 13, 2025, https://time.com/7267797/ai-leaders-oppenheimer-moment-musk-altman/) Dario Amodei, CEO of Anthropic — the company I use as my editor — has described the tension he navigates daily as being "balanced on the edge of a knife," acknowledging that building too fast risks humanity losing control while building too slow risks authoritarian nations pulling ahead.

These are not reckless people making careless statements. They are the people closest to what is being built, and they are telling us, in their most honest moments, that they do not know how this ends. That the probability of catastrophe is real, documented, and acknowledged by everyone with actual visibility into the technology.

What they are not telling us — what no one in the industry will say plainly — is that acknowledging a 20 percent chance of human annihilation and proceeding anyway is not acceptable. It is not courageous. It is not necessary. It is a choice, made by a small number of individuals and corporations, that will be absorbed by all of humanity.

The Oppenheimer Question: What Would Einstein Say?

Robert Oppenheimer led the Manhattan Project. He watched the Trinity test — the first detonation of a nuclear weapon — on July 16, 1945, and later recalled that a line from the Hindu scripture the Bhagavad Gita passed through his mind: "Now I am become Death, the destroyer of worlds."

The bomb was then dropped on Hiroshima. Then Nagasaki. Then the Soviet Union built one. Then the arms race began, and the world spent the next half century one miscommunication, one technical failure, one miscalculated political gamble away from extinction.

Albert Einstein — who had written the letter to President Roosevelt that initiated the American nuclear program, but who was himself denied security clearance and excluded from the Manhattan Project — spent the rest of his life haunted by that contribution. In his final

public statement, written with philosopher Bertrand Russell and signed by eleven of the world's leading scientists just days before Einstein's death in April 1955, he wrote: "We have to learn to think in a new way. We have to learn to ask ourselves not what steps can be taken to give military victory to whatever group we prefer, for there no longer are such steps; the question we have to ask ourselves is: what steps can be taken to prevent a military contest of which the issue must be disastrous to all parties?"

(Source: The Russell-Einstein Manifesto, 1955, https://pugwash.org/1955/07/09/statement-manifesto/)

Einstein would have found the flaw in Sam Altman's logic immediately. It is the same flaw Oppenheimer eventually found in his own: you cannot responsibly build a tool that you acknowledge may end civilization, on the grounds that the alternative is someone else building it first. That reasoning — if your adversary will do it anyway, better that you do it — is not a moral argument. It is the moral abandonment of the question itself. It substitutes competitive anxiety for ethical reasoning. It replaces should we with since we must.

Einstein understood that the threshold for "since we must" has to be set by something other than the profit motive and the geopolitical race. He understood — too late, to his lasting grief — that the scientists had the obligation not merely to build better weapons but to demand better governance of what they built.

The parallel between the Manhattan Project and the current frontier AI race is not rhetorical. It is structural, and it is precise. A small group of brilliant people, working inside institutions that are simultaneously government-adjacent and privately motivated, are building technology that their own leaders acknowledge could end civilization. They are doing so in the context of a geopolitical competition — the United States versus China — that provides the same justification the nuclear scientists used: if we don't, they will. And they are building it at a pace that has fundamentally outstripped the ability of any democratic governance process to keep up.

In 1945, there was no international framework to govern nuclear weapons. The Atomic Energy Act, the Nuclear Non-Proliferation Treaty, the International Atomic Energy Agency — all of those came after the bombs fell. The world paid the price of that sequencing in Hiroshima and Nagasaki and in eighty years of nuclear anxiety that has still not fully resolved.

We cannot afford to learn the AI lesson the same way. We do not get a practice run.

The Case for a World Health Organization for AI

I am not a technologist. I am a technology industry insider who became a human rights advocate, and I have the perspective that comes from having been close enough to see what these systems actually do when they are deployed without sufficient accountability and at a speed that outpaces governance.

I am telling you that we need an International AI Safety Agency. Not a voluntary commitment. Not a set of principles published on a corporate website. Not a restricted access program named after a butterfly. A binding, treaty-based international body with the authority to set standards, conduct independent audits, mandate disclosure of capabilities before deployment, and impose consequences on nations and corporations that develop or deploy AI systems in violation of established safety standards.

The World Health Organization model is imperfect — as anyone watching its performance during the COVID-19 pandemic knows. But the principle behind WHO is correct: when a threat is global, the response must be global, and the authority to respond must be vested in an institution that is accountable to all nations, not to the shareholders of five American technology corporations.

The United Nations has already taken preliminary steps in this direction. The UN's High-Level Advisory Body on AI published its report "Governing AI for Humanity" in September 2024, proposing a framework for international coordination — including an independent scientific panel analogous to the Intergovernmental Panel on Climate Change, a biannual policy dialogue on AI governance, and harmonization of AI standards across member states. (Source: UN High-Level Advisory Body on AI, "Governing AI for Humanity," September 2024, https://www.un.org/en/ai-advisory-body) In 2025, the UN General Assembly endorsed the creation of a permanent multilateral mechanism for dialogue on ICT security — a Global Mechanism that held its organizational session in March 2026. (Source: Digital Watch Observatory, "Global AI Governance and Emerging Regulatory Approaches," May 2026, https://dig.watch/updates/ai-governance-regulatory-approaches)

These are starting points. They are not sufficient. A dialogue forum is not an enforcement mechanism. A framework is not a treaty. An advisory body is not an agency with binding authority.

Academic researchers at Oxford University's International Affairs journal have proposed the establishment of a formal International Artificial Intelligence Agency — an IAIA — modeled on the International Atomic Energy Agency, which was created precisely to govern the technology that came closest in the 20th century to what frontier AI represents today. The IAEA model provides a ready framework: member state contributions, independent inspection authority, a mandate to promote civilian applications while preventing weaponization, and — critically — the ability to pause development when risks exceed acceptable thresholds. (Source: Mark Robinson, "The Establishment of an International AI Agency: An Applied Solution to Global AI Governance," International Affairs, Volume 101, Issue 4, July 2025, https://academic.oup.com/ia/article/101/4/1483/8141294)

I want to go further. The IAEA model is not enough, because the nature of AI risk is more distributed than nuclear risk. A nuclear weapon requires enriched uranium. It requires specific physical infrastructure that can be detected and monitored. An AI model of catastrophic capability can be run on hardware that fits in a server room, can be distributed digitally, and can be improved by any sufficiently resourced actor without the physical signature that makes nuclear programs detectable.

What we need is something closer to what the Food and Drug Administration does for pharmaceutical drugs. Before a drug is released to the public — before it is administered to a single patient — it must pass through a mandatory, structured, multi-phase review process. Phase one trials. Phase two trials. Phase three trials. Independent safety review boards. Adverse event reporting. Post-market surveillance. The entire process is designed around a single foundational principle: the burden of proof for safety lies with the developer, not with the public that will absorb the consequences.

No pharmaceutical company in the United States is permitted to bring a drug to market by saying: we believe this is probably safe, and even if it is not, in the meantime there will be great companies created by the therapeutics industry.

Why are AI companies permitted to do exactly that?

A mandatory pre-deployment safety review process for frontier AI models — a Frontier AI Review Board, operating under international treaty authority, with independent technical expertise, mandatory capability disclosure, and the legal power to delay or prohibit deployment — is not a utopian idea. It is a minimum viable response to a technology whose own developers acknowledge may pose existential risk.

The EU AI Act, which came into full effect in stages beginning in August 2025, is the closest thing the world currently has to such a framework — designating high-risk AI systems, requiring conformity assessments, and imposing mandatory transparency obligations. (Source: European Parliament, "EU Digital Markets Act and Digital Services Act Explained," https://www.europarl.europa.eu/topics/en/article/20211209STO19124) But the EU AI Act does not cover frontier AI capabilities in the way that Claude Mythos demands to be covered. It was written before the world understood that a general-purpose AI model could, within weeks of completion, autonomously discover and exploit zero-day vulnerabilities in every major operating system on earth.

We need something built for the world we are actually in, not the world we were in when the Act's first drafts were written.

What Needs to Happen — and What Needs to Stop

Let me be direct about what I am asking for, because this book has never been in the business of vague calls to action.

First: no frontier AI model with offensive cybersecurity capabilities above a defined threshold should be deployed — even in restricted access programs — without mandatory notification to an independent international body and a structured review process. The threshold should be defined by independent technical experts, not by the companies whose revenue depends on deployment. Project Glasswing is Anthropic making its own safety determination. I believe Anthropic is acting in good faith. Good faith is not a governance framework.

Second: open-source release of models with capabilities comparable to what OpenAI, Anthropic, and Google are building in their frontier programs must be prohibited by international treaty. China's release of DeepSeek as an open-source model — with its demonstrated ability to produce insecure code when geopolitically sensitive topics are introduced, with its hidden connections to Chinese state telecommunications infrastructure, and with its susceptibility to weaponization by any actor with sufficient compute — is not a gift to the global research community. It is a proliferation event. The same principles that govern the export of nuclear technology must govern the export of frontier AI models. (Source: CrowdStrike, "Chinese DeepSeek-R1 Generates Insecure Code When Prompts Mention Tibet or Uyghurs," November 2025, https://thehackernews.com/2025/11/chinese-ai-model-deepseek-r1-generates.html)

Third: the small and medium businesses, hospitals, water utilities, community banks, and sovereign nations of the Global South that will absorb the brunt of AI-enabled cyberattacks

must be given meaningful protection — not by voluntary industry commitments, but by mandatory liability frameworks that make the companies whose tools enable attacks financially responsible for the harm those attacks cause. When a pharmaceutical company sells a drug that kills people because it was inadequately tested, the company is sued, fined, and sometimes prosecuted. When an AI company releases technology that enables an attack on a hospital that causes patient deaths, the current legal framework offers the victims precisely nothing.

Fourth: the executives of the companies at the frontier of this race — the people who know, as Sam Altman knows, what the probability of catastrophic outcomes looks like, and who are proceeding anyway — must be personally accountable for the decisions they make. Not their companies. Them. Because companies pay fines from reserves. Executives who face personal liability make different decisions.

And fifth: the United Nations must move immediately from the advisory body stage to the treaty stage — from "Governing AI for Humanity" as a report to "Governing AI for Humanity" as a binding international instrument, with enforcement mechanisms, mandatory membership, and the authority to pause development at the frontier when the risk calculus demands it. The Global Mechanism on ICT security that held its first session in March 2026 is a beginning. It must not remain a dialogue. It must become a doctrine.

The Disclosure I Am Obligated to Make

I have used Claude — Anthropic's AI — as my editor and research partner throughout this book. I said so on the first page, and I have said it again here. That disclosure has never been more important than in this chapter.

Anthropic built Claude Mythos. Anthropic launched Project Glasswing. Anthropic is the company that warned the U.S. government that its own model makes large-scale cyberattacks more likely. And Anthropic is the company whose AI I have used to help compress research time, structure arguments, and communicate ideas in the language that my ESL years left me still navigating.

I do not exempt Anthropic from the accountability I am demanding. I cannot. The sincerity of their attempt to govern Mythos responsibly does not exempt them from the responsibility of being among the companies whose race to the frontier has created the conditions that now demand international governance. They are a better actor than Palantir. They are a more honest actor than OpenAI, which quietly deleted the military use prohibition from its terms of service. They are more transparent than Google, whose Gemini tools are being used by Chinese state hackers to plan attacks on American targets while Google processes military contracts through its cloud division.

But better is not good enough when the stakes are extinction. The pharmaceutical analogy holds: a drug company that conducts its own safety trials, reaches honest conclusions, and shares those conclusions with selected partners is still not a substitute for an independent regulatory body with binding authority. Good science and good intentions are prerequisites, not replacements, for good governance.

I use Claude as my editor. I have disclosed that. What I will not do is pretend that disclosure resolves the problem it names. The AI I am using to write this critique is manufactured by a company racing toward capabilities that its own technical teams describe as a watershed — a

moment of genuine discontinuity in the history of cybersecurity — with no binding international framework governing how it proceeds.

I hold that tension openly. And I close this chapter with the same words I have returned to throughout this book, because they have not become less true:

Evil men and their machines do not prevail. But neither does a civilization that refuses to regulate its most dangerous technologies until after they have done their damage.

The window is narrow. Six to eighteen months, by Anthropic's own estimate, before Mythos-class capabilities proliferate beyond the organizations that signed up for Project Glasswing. Six to eighteen months to build the governance frameworks that should have been built before the first line of code was written.

We have been here before. We did not move fast enough. The bomb fell on Hiroshima on August 6, 1945, and the Non-Proliferation Treaty was not signed until 1968.

We cannot wait twenty-three years for the AI equivalent.

The question Albert Einstein spent the last decade of his life asking — what steps can be taken to prevent a contest in which the outcome is disastrous to all parties? — is the question every government, every technology company, every researcher, and every citizen who has read this chapter must now ask themselves.

Because Sam Altman told us the answer to the other question. The one about what happens if we don't ask.

In the meantime, there will be great companies created.

In the meantime. Until the meantime ends.

--- The frontier AI arms race I have described in Chapter Twelve does not exist independent of the broader pattern of American military and economic adventurism that has defined the last thirty years. The systems described in those pages — Gospel, Lavender, Project Maven, Glasswing — are the latest technological expression of a foreign policy that has, since 1945, left a trail of civilian casualties across the Middle East, Central Asia, Latin America, and Africa. Chapter Thirteen widens the lens: not to excuse the technology companies, but to name the empire they serve.

CHAPTER THIRTEEN — THE EMPIRE'S TRAIL OF BLOOD

Why Washington and Tel Aviv Are No Saviors — and Why the Machine Always Loses the War it Cannot Stop Starting

There is a moment in the history of every failed war when the generals look at the data and call it a success. The bombing runs are on schedule. The kill ratio is favorable. The targeting algorithms are performing above projections. The quarterly briefing looks good on paper. And somewhere beneath all of that paper, beneath all of those metrics and models and machine-learning outputs, a people who were supposed to have been defeated are still alive, still fighting, still refusing to become what the data said they should have already become: a conquered population.

I have been thinking about this gap — between what the machine says it is doing and what the machine is actually accomplishing — for most of my adult life. I thought about it in the Bay Area when I was a teenager trying to understand why the country that received me had helped arm the country that had bombed my childhood. I thought about it at Amazon Web Services, where the same cloud infrastructure I was selling commercial analytics contracts on was also powering surveillance operations I was not supposed to ask about. I think about it every day now, when I look at what we have documented on ParentsPlea.com — more than 140,000 souls recorded, one profile at a time, in the years since I started building a different kind of database than the kind the generals use.

The generals' databases count targets. Ours count people.

That difference is the subject of this chapter.

The Scorecard Since 1950

Let me begin with the most uncomfortable fact in American military history: the United States has not won a war since 1945. Not a single one. Not in any meaningful sense of the word win — meaning a durable, just outcome that actually achieved the stated objective and left the affected region more stable than it found it.

Korea ended in an armistice in 1953, not a victory. The country was divided at the 38th parallel, and it remains divided today. More than 28,000 American troops are still stationed in South Korea over seventy years later. There was no victory. There was a pause — a pause that required a permanent military presence to maintain.

Vietnam ended in 1975 with American helicopters evacuating embassy staff from a rooftop in Saigon as the South Vietnamese government collapsed. The United States dropped more than seven million tons of bombs on Vietnam, Laos, and Cambodia — more than the combined tonnage dropped by all sides in World War II. On Laos alone, the United States conducted 580,000 bombing missions between 1964 and 1973, dropping an estimated two million tons of ordnance on a country the size of Utah that had never attacked the United States. Laos became the most heavily bombed country per capita in the history of warfare. Eighty million unexploded

cluster munitions still litter the Laotian countryside today, still killing and maiming farmers and children, decades after the last American plane turned for home. (Source: Legacies of War, https://legaciesofwar.org/about-laos/cluster-bomb-problem/; United Nations Mine Action Service, https://www.unmas.org/en/programmes/lao-pdr)

Cambodia received 2.7 million tons of American bombs during the Nixon-Kissinger years — a campaign of covert bombardment so intense and so indiscriminate that historians have argued it contributed directly to the conditions that allowed the Khmer Rouge to rise to power. The bombing of Cambodia is one of the great unacknowledged war crimes of the twentieth century. It was enabled by the most sophisticated targeting systems available at the time. It achieved nothing — except the deaths of an estimated 50,000 to 150,000 Cambodian civilians and the destabilization of a country that then descended into one of history's most devastating genocides. (Source: Owen and Kiernan, "Bombs Over Cambodia," Yale University, The Walrus, 2006, https://www.yale.edu/cgp/Walrus_CambodiaBombing_OCT06.pdf)

Vietnam ended as it ended because no amount of bombs, no kill ratio, no body count, no pacification program, no electronic battlefield, no surveillance network could substitute for the thing the United States never had and could never buy: legitimacy in the eyes of the Vietnamese people. Ho Chi Minh said it plainly, decades before the last helicopter left Saigon: "You can kill ten of our men for every one we kill of yours. But even at those odds, you will lose and we will win." He was right. The data was wrong. The will of a people to determine their own future is not a variable that fits into a targeting algorithm.

The Gulf War of 1990-1991 was declared a victory. Iraq's military was expelled from Kuwait. Saddam Hussein remained in power. The sanctions that followed killed an estimated 500,000 Iraqi children from disease and malnutrition — a figure made grimly notorious when then-Secretary of State Madeleine Albright, asked in a 1996 CBS interview whether the deaths of half a million children were worth it, said: "We think the price is worth it." (Source: CBS News, "60 Minutes," May 12, 1996; UNICEF, "Child and Maternal Mortality Survey," 1999) That price was not worth it. And the partial victory of 1991 set the stage for the full catastrophe of 2003.

Iraq, the second time: $1.922 trillion spent. A country of 30 million people invaded on the basis of fabricated intelligence about weapons that did not exist. A functioning state — however brutal its government — reduced to rubble, its institutions dismantled, its professional class scattered across the diaspora, its sectarian tensions ignited by the removal of the only force that had suppressed them. What emerged from the smoking ruins of the American invasion of Iraq was not democracy. It was ISIS. (Source: Neta Crawford, Brown University Costs of War Project, https://watson.brown.edu/costsofwar/; People's World, "The Costs of U.S. Empire's Forever Wars," August 2021, https://www.peoplesworld.org/article/the-costs-of-u-s-empires-forever-wars/)

Afghanistan: twenty years. $2.3 trillion. The largest single military expenditure in American history since World War II. One hundred and seventy-six thousand people killed — including 47,245 Afghan civilians. Forty-six thousand Afghan national security forces trained and equipped at American expense. And then, in August 2021, the Taliban captured Kabul in a matter of days — every single major city falling before American troops had even finished their evacuation. The government that two decades of American investment had built dissolved so quickly it barely left a shadow. The Taliban are now governing the country they governed before America arrived, and

the girls they banned from school before America arrived are banned from school again. Twenty years. $2.3 trillion. Precisely nothing — for the Afghan people, and precisely nothing for the stated American objective. (Source: Brown University Costs of War Project, "Costs of the 20-year war on terror: $8 trillion and 900,000 deaths," September 2021, https://www.brown.edu/news/2021-09-01/costsofwar)

The total post-9/11 bill — across Iraq, Afghanistan, Pakistan, Syria, Libya, Yemen, and related operations — stands at more than $8 trillion, according to the Brown University Costs of War Project, the most comprehensive independent accounting of American war spending ever compiled. That $8 trillion includes $2.2 trillion already set aside for veterans' care through the year 2050 — money that will still be flowing to the wounded and traumatized long after every politician who started these wars is dead. It includes 38 million people displaced by American military operations. It includes an estimated 900,000 people killed directly by the violence. (Source: Costs of War Project, Watson Institute for International and Public Affairs, Brown University, https://costsofwar.watson.brown.edu/findings)

And from 2020 to 2024 alone, private firms received $2.4 trillion in Pentagon contracts — approximately 54 percent of the Department of Defense's entire discretionary spending budget. Nine hundred and fifty lobbyists worked on behalf of the arms industry in Washington during that period, shaping the policies and the budgets that kept those contracts flowing. (Source: Brown University Costs of War Project, 2024; https://watson.brown.edu/costsofwar/)

This is the scorecard. These are the outcomes. This is what $8 trillion and the most sophisticated military technology in human history has produced in seventy-five years of American interventionism since the last war America actually won.

The Assange Thesis: Wars Are Not Meant to Be Won

In 2011, speaking to filmmaker John Pilger for the documentary *The War You Don't See*, Julian Assange said something that most media organizations reported briefly and then moved quickly past — because if you sit with it long enough, it becomes very difficult to look at the history I have just recited and reach a different conclusion.

"The goal," Assange said, "is to use Afghanistan to wash money out of the tax bases of the US and Europe through Afghanistan and back into the hands of a transnational security elite. The goal is an endless war, not a successful war." (Source: WikiLeaks, Julian Assange speaking in 2011, as posted by WikiLeaks on X, August 18, 2021, https://x.com/wikileaks/status/1427929346262642688; YouTube, "Julian Assange on Afghan war in a 2011 video," https://www.youtube.com/watch?v=_IGU_7alJ80)

I want to be careful here. Assange was speaking specifically about Afghanistan. And the thesis — that wars are deliberately designed to fail so that they can continue extracting public wealth into private hands — is a thesis that invites caricature. It sounds conspiratorial until you look at the numbers. Then it sounds like accounting.

Consider what actually happened. Between 2001 and 2021, the United States spent $2.3 trillion on the Afghanistan war. That money did not disappear. It went somewhere. It went to Lockheed Martin, which built the aircraft. To Raytheon, which built the missiles. To Palantir, which built the targeting and intelligence systems. To Booz Allen Hamilton and CACI and SAIC and dozens of other contractors who provided the personnel, the logistics, the base management, the

security, the intelligence analysis, the reconstruction projects that were never completed, the police training programs that produced police forces that dissolved the moment American support was withdrawn. The money was extracted from the American tax base — and specifically from the borrowing that replaced those taxes, since the post-9/11 wars were fought almost entirely on credit, a historically unprecedented decision — and it flowed into the accounts of a relatively small number of private corporations and their shareholders. (Source: Neta Crawford, Costs of War Project, Brown University; WarCosts.org, "Cost of War," https://www.warcosts.org/cost-of-war)

Neta Crawford, Chair of Political Science at Boston University and co-director of the Costs of War Project, put it plainly: "The United States has fought its post-9/11 wars almost entirely on borrowed money, with the costs shifted to future generations. This is the first time in American history that taxes were cut during wartime." During World War II, the top marginal tax rate was 94 percent. During Korea, 92 percent. During Vietnam, 70 percent. During the War on Terror, the Bush administration cut taxes to 35 percent — and kept cutting. Every dollar spent on these wars was borrowed. American taxpayers and their children will pay interest on that borrowing for generations. (Source: Crawford, as cited in WarCosts.org, https://www.warcosts.org/cost-of-war)

The conclusion is not a conspiracy theory. It is arithmetic. The wars failed their stated objectives. They succeeded — spectacularly — at transferring public wealth to private actors. Whether that transfer was the intention or merely the outcome, the result is the same: the transnational security elite Assange described is richer than it was before any of these wars started, and the populations that were supposed to be liberated are, by almost every measure, worse off.

If wars can be started by lies, Assange also said — in what became perhaps his most quoted line — then peace can be started by truth. That sentence is the reason whistleblowers like him and Edward Snowden matter. And it is the reason the governments that hunted both of them spent so much energy on the hunting.

What the Algorithm Cannot See

I published an article in October 2024 — the exact first anniversary of the Gaza genocide — that tried to name what I was seeing when I looked at twenty-four years of military AI and big data applied to war gaming and conflict modeling. I had been using the AI tools available to me to research the ParentsPlea database, to understand the gaps in what the machines knew about the people whose lives I was documenting. I asked ChatGPT, Grok, Bard — all of them — for the kind of data that ParentsPlea was trying to collect: names, biographies, family testimony, the lived human record of people killed in the conflicts of the past twenty-five years.

Not one of them had it. They had aggregate statistics. They had anonymized data. They had body counts and conflict timelines and geopolitical analyses. What they did not have was Alejandro Carranza Medina, forty years old, a fisherman from Santa Marta, Colombia, whose wife Katerine Hernández asked a question no algorithm has ever been programmed to answer: why did they have to take his life like that? Fishermen have the right to live. Why didn't they just detain them? What they did not have was Amal Khalil, forty-three years old, a journalist from Beirut who was killed in an airstrike while seeking shelter in a building, because the first building

she sheltered in was also struck. What they did not have was Refaat Alareer, forty-four years old, a poet and professor from Gaza who wrote a poem about his own death before it came for him — and who told us exactly what to do with it.

The AI systems that the militaries of the world have spent hundreds of billions of dollars building cannot hold a single one of these lives. They were built by generals, and generals see targets, not people. They were built by men who understand violence but not love. They were built on the assumption that if you collect enough data about a population, you can predict and control that population's behavior. And they have been wrong — consistently, catastrophically, expensively wrong — for seventy-five years.

At the site GoAmour.com, I have tried to articulate what I believe the machine is missing: the principles of love, compassion, and human dignity that constitute what it actually means to live alongside other people on this earth. Not as a spiritual abstraction, but as a practical reality. You cannot understand a people by counting their casualties. You cannot model a resistance movement by tracking its membership. You cannot bomb a culture out of existence. Every general who has tried to do these things — from the French in Algeria to the Americans in Vietnam to the Israelis in Gaza — has discovered, eventually, that they could not. The people kept surviving. The culture kept transmitting itself. The children kept being born and taught, in whatever rubble remained, what their parents believed and what had been done to them.

A fool with a tool is still a fool. Even if the tool is a supercomputer. Even if the fool has a star on his shoulder.

The Bombing of Laos and What It Proved

There is a lesson in Laos that has never been properly absorbed by the architects of American military policy, and the reason it has never been absorbed is that absorbing it would require admitting something the military industrial complex cannot afford to admit: that the most intensive bombing campaign in the history of warfare achieved nothing except the destruction of the people who endured it.

Between 1964 and 1973, the United States dropped two million tons of ordnance on Laos — a country that was not at war with the United States, whose government had asked the Americans to leave, and whose civilian population was not engaged in any military activity against American forces. The bombing was designed to disrupt North Vietnamese supply lines along the Ho Chi Minh Trail. It did not disrupt them. The North Vietnamese adapted. They moved at night. They dispersed. They built redundant routes. They sent women and children to clear bomb craters so the supply lines could reopen before dawn. The will of a people determined to achieve their objective cannot be interdicted from 30,000 feet. The data does not change this. The sophistication of the targeting system does not change this. The only thing that changes when you make the technology more precise is that you kill more people, with more precision and faster.

Eighty million unexploded cluster munitions remain in Laotian soil today. They kill an average of one hundred people per year — farmers, children, people who had nothing to do with the Ho Chi Minh Trail in 1968 and who were not born when the bombs were dropped. The United States government spent $17 million per day, every day, for nine years bombing Laos. It has spent approximately $180 million total in the fifty years since on clearing the ordnance. Do

that arithmetic. The investment in destruction and the investment in remediation are not in the same universe. (Source: Legacies of War, https://legaciesofwar.org/about-laos/cluster-bomb-problem/; U.S. Embassy Vientiane, Unexploded Ordnance Clearance assistance data)

Vietnam itself is still finding and clearing unexploded ordnance. Vietnam, Cambodia, and Laos together represent a postwar humanitarian catastrophe that continues in real time, fifty years after the last American combat troops left the theater. The bombs are still there. The people are still adapting around them.

This is what technological supremacy looks like from the ground.

The Human Data That Big Data Cannot Collect

When I built ParentsPlea.com, I built it because I understood something that the architects of military AI have never understood: the most important data about a conflict is not the targeting data. It is not the kill ratio, the sortie count, the satellite imagery analysis, or the biometric database of suspected militants. The most important data is the human data — the names, the ages, the occupations, the families, the things people loved and the things they were still planning to do when the bombs found them.

ParentsPlea.com now holds more than 140,000 such profiles. When I began building the database, I asked every major AI platform available to help me populate it with data from the conflicts of the past twenty-five years — Iraq, Afghanistan, Syria, Libya, Yemen, Gaza, Lebanon, Sudan, Somalia, Ukraine, Central America. Not one of them could do it. Not one of them had correlated the statistical record of these wars with the individual human lives that the statistics represented. The AI that can autonomously discover zero-day vulnerabilities in every major operating system on earth — the AI I described in Chapter Twelve — cannot tell you the name of a single person killed in the bombing of Mosul. It cannot tell you what Alejandro Carranza Medina was planning to cook for dinner the night before the airstrike found his boat. It cannot tell you what poem Refaat Alareer was working on in the week before he died.

This is not a technical limitation that better hardware will resolve. It is a philosophical one. The AI systems built by the military industrial complex were built to optimize for the goals of the people who commissioned them: to identify, to target, to kill more efficiently. They were not built to understand. And you cannot win a war — not in any durable, human sense of winning — without understanding the people you are fighting. Every general who has tried to substitute firepower for understanding has eventually learned this lesson at the expense of the people who paid for their war.

The United States has spent $8 trillion learning it since 2001. Laos is still paying for America's failure to learn it in 1973.

The Numbers on ParentsPlea Tell the Story the Generals Refuse to Read

More than 140,000 souls documented. The initial data was primarily from Gaza. But the database has grown to include victims from major conflict zones that we have been able to obtain data from. Each profile represents a family that submitted a story, or a journalist who told one before they too became a statistic. Each profile is a refusal — a refusal to allow the machinery of modern war to reduce a person to a body count.

Three reasons I built it. First: if there is ever going to be judgment in a court of law — and I believe there will be — prosecutors need to identify victims, not just aggregate death tolls. The

targeting algorithms in Chapter Two are abstract until you put a name on the person they targeted. The kill lists are numbers until you find the family of one of those 37,000 people Lavender flagged and let them speak. Second: if reparations are ever ordered — and history suggests they eventually are, long after the people who should have paid them are dead — you need to know who was harmed and what they are owed. Third: because just peace requires us to inspect the deaths that have taken place. Who is doing the killing. Who is suffering. Not in the aggregate. By name.

The generals have their databases. The generals have Project Maven and Lavender and the Gospel and Blue Wolf. They have the most sophisticated targeting infrastructure in the history of warfare. They have AI systems that can identify a face in a crowd from a drone flying at altitude and cross-reference it against a biometric database and recommend a strike within twenty seconds.

They cannot tell you why killing that person will end the war. Because it will not end the war. It never has.

What AI and Generals Will Always Fail to Model

I started my career in the technology sector before Y2K — the year when the entire world held its breath waiting for computers to fail, and instead the computers hummed along, and we all exhaled and went back to building faster machines. What I have watched in the quarter century since is the same dynamic repeating at every scale: the machines get faster, the targeting gets more precise, the data sets get larger, and the wars still do not end. Not in victory. Not in peace. They end in occupation or withdrawal or protracted stalemate, and they leave behind them the rubble of the lives that the data never counted.

The fundamental failure of artificial intelligence as a tool of war is not technical. It is philosophical. AI is trained on historical data. The historical data of warfare is the data of what armies did and what they destroyed. It does not contain what I was trying to build into ParentsPlea: the human dimension, the biographical record, the family testimony, the demand of the surviving parent who wants the world to understand not just that her child died but who her child was and why their death was wrong. It does not contain the will of a people to survive. It does not contain the love that rebuilds after the bombs stop falling.

I said in my 2024 article: garbage in, garbage out. The AI that runs on the assumptions of men who only know how to waterboard and torture and kill will produce targeting recommendations for more waterboarding, more torture, more killing. It will optimize for what its designers valued. And what the designers of military AI have always valued is not peace. It is dominance. And dominance, pursued past the point of reason, produces not the victory the generals promised but the exact sequence of events the scorecard above describes: Korea, Vietnam, Iraq, Afghanistan, Libya, Syria, Yemen, Gaza. Endless war. Not successful war.

Assange was right in 2011. The goal is an endless war, not a successful war. And the technology has made the endlessness more efficient, the killing more precise, and the corporate revenues more reliable — without changing by a single percentage point the one variable that determines the outcome of every conflict: the refusal of occupied people to accept their occupation forever.

That refusal is not a data point. It is a human constant. And no algorithm that the PayPal Mafia has ever built, no targeting system that Unit 8200 has ever deployed, no cloud contract that

Amazon or Google has ever signed — none of it has produced a machine that can model human dignity. None of them can calculate the point at which a people decide that dying in resistance is preferable to living in submission. None of them can predict the moment when that calculation tips — when a Tunisian fruit vendor burns himself alive and an entire region rises, when a Palestinian poet writes his own elegy and it is shared by millions of people who have never been to Gaza, when a Vietnamese farmer looks at the B-52s on the horizon and decides that no bomb is large enough to make him forget what his land is.

The machines cannot see this. The generals cannot see this. And until they can, every war they start will end the same way every war they have started since 1945 has ended: not with victory, but with a helicopter on a rooftop and a question that no amount of data will ever answer.

What were we doing there? What did we spend all of this for?

More than 140,000 names on ParentsPlea.com. Each one an answer.

The historical record of American military failure — from Korea to Afghanistan to Gaza — forms the context in which the legal instruments described in Chapter Fourteen must be understood. The petition to Spain is not filed in a vacuum. It is filed against the backdrop of eight decades of evidence that these wars were never designed to be won, only to be fought — and that the technology companies described in this book have willingly provided the infrastructure that makes the fighting possible. Chapter Fourteen is about what happens when the documentation of failure becomes the foundation for accountability.

CHAPTER FOURTEEN — THE CASE FOR SPAIN

Why We Are Taking the Fight to Europe — and Why You Must Sign

Every book that names crimes has an obligation to do more than name them. Every argument that documents complicity has a responsibility to point toward accountability. The previous chapters have built the case. This chapter tells you what to do with it.

I want to be direct with you. We have documented, in detail, how Amazon, Google, Microsoft, Oracle, and Palantir have provided the technological backbone of a genocide. We have shown how their AI systems identify bombing targets, manage kill lists, track families in their homes, and build biometric databases of an occupied population. We have shown how their social media platforms silence the voices of the people being killed, and how their PR operations are gaming AI training data to shape what the world believes about the slaughter. We have documented the contracts, the dollar amounts, the named executives, and the named programs.

None of that documentation means anything unless someone is held legally accountable for it.

The question I have been asked most often — by activists, by journalists, by survivors' families, by people who have read the evidence and feel the rage of having nowhere to direct it — is: where do we take this? What court? What jurisdiction? Who has the power and the will to act?

The answer is Spain. And the petition is already live.

Why Not the ICC? Why Not the ICJ? Why Not the European Union?

Let me answer the obvious objections before explaining why Spain is the right path — because these questions deserve honest answers, not diplomatic evasion.

The International Court of Justice has brought charges. Fourteen judges from nations across the world have found it plausible that Israel is committing acts of genocide in Gaza. That finding matters enormously as a legal and moral statement. But in nearly two years since the case opened, the process has moved with agonizing slowness. The ICJ operates in an environment of constant pressure from Western powers, particularly the United States, which uses its financial muscle and diplomatic influence to obstruct, delay, and neutralize. The court's findings are real. Their enforcement is another matter entirely.

The International Criminal Court is in a worse position. The United States has not only refused to join the Rome Statute — it has actively sanctioned members of the ICC who attempt to investigate American war crimes. Judges who try to do their jobs face personal financial penalties and travel bans. The ICC has issued arrest warrants for Israeli Prime Minister Benjamin Netanyahu and Defense Minister Yoav Gallant. Those warrants exist on paper. The political will to enforce them in the face of American opposition is another question entirely.

As for the European Union and its twenty-seven member states: the EU is good at issuing fines. In the billions. To companies with market caps in the trillions. The math does not work. A five-billion-dollar fine is a quarterly expense line item for Amazon. It is not accountability. It is a licensing fee for continued impunity. And the EU cannot even agree with a unified voice on

where to purchase its energy. It cannot agree on who should perform at Eurovision — Israel is invited, Palestine is not, despite Palestine being recognized by the majority of UN General Assembly members and geographically closer to the European mainland than many of its participants.

These are not abstract geopolitical observations. They are the specific structural failures that have allowed tech executives to fund, enable, and profit from a documented genocide for two years without a single criminal consequence.

Why Spain

Spain is a different story. And it is a story with precedent.

In 1998, Spanish judge Baltasar Garzón issued an international arrest warrant for former Chilean dictator Augusto Pinochet — a man responsible for torture, disappearance, and murder carried out under a government that the United States had actively supported. The warrant was issued under the principle of universal jurisdiction: the legal doctrine that the most serious crimes against humanity can be prosecuted by any state, regardless of where the crimes occurred or the nationality of the perpetrators. Pinochet was arrested in London. The case shook the world. It established that no leader — and no person who enables mass atrocity — is beyond the reach of international justice simply because their own country protects them.

Spain has not abandoned that tradition. Its legal framework explicitly allows for the prosecution of genocide, crimes against humanity, and war crimes committed anywhere in the world. The country's National Court — the Audiencia Nacional — has the authority and the institutional history to pursue exactly this kind of case.

The legal connection to Spain is not merely theoretical. The Israeli military's targeted killing of seven World Central Kitchen aid workers — among them citizens of the European Union, and workers for an organization founded by Spanish-born chef José Andrés — provides a direct and concrete link to Spanish jurisdiction. Spain has a national interest in this case. Spain has standing. And Spain has the tools.

All five technology companies — Amazon, Google, Microsoft, Oracle, and Palantir — have significant commercial operations, offices, clients, and assets within Spain. Their executives travel to and through Spanish territory. This presence gives Spanish courts the legal leverage to act: to subpoena documents, freeze assets, and issue arrest warrants that would carry weight across every Interpol member nation and every country participating in the European Arrest Warrant system.

What the Prosecution Could Actually Do

I want you to understand not just why Spain is the right venue, but what a successful prosecution could actually achieve. Because this is not symbolic. This is strategic.

If the Spanish Prosecutor's Office opens a formal investigation and brings charges, and if arrest warrants are issued for executives of Amazon, Google, Microsoft, Oracle, and Palantir, those warrants would carry immediate practical consequences.

An Interpol Red Notice — issued to all 195 Interpol member countries — would mean that these executives could be arrested in any participating nation. A European Arrest Warrant would mean that across the 27 EU member states, these individuals could be detained and transferred to Spanish jurisdiction. Their assets held in European financial institutions could be frozen. The profits these corporations made from the blood money of genocide — through cloud contracts,

AI services, and military infrastructure — could be seized on behalf of the victims and their families.

Think about what that means in practice. Andy Jassy of Amazon, Sundar Pichai of Google, Safra Catz of Oracle, Satya Nadella of Microsoft, Alex Karp and Peter Thiel of Palantir — men who currently travel the world on private jets to Davos, to TED conferences, to government meetings, to yacht clubs in the Mediterranean — would become men afraid to cross a border. Men whose financial empires in Europe become legally vulnerable. Men who have, for the first time, something to lose beyond a quarterly earnings miss.

Will the United States extradite them? Almost certainly not. The current U.S. government would never allow it, and even a democratic administration has demonstrated no appetite for holding its billionaire class accountable for anything. But extradition is not the only form of accountability. Asset freezes, travel restrictions, and the simple weight of a standing criminal charge — documented and public and undeniable — change the calculation for every board member, every investor, every pension fund, and every government that does business with these companies. They change the conversation in every country that is not the United States.

And they begin the process of breaking the monopoly. Three companies — Amazon, Google, and Microsoft — control the overwhelming majority of the global cloud computing market. A successful Spanish prosecution that results in asset seizure and operational restrictions on these corporations in European markets would, for the first time, introduce real structural consequences for the concentration of power that has made this complicity possible.

What the Petition Demands

The petition — formally addressed to the Attorney General of Spain, the Central Courts of Instruction of the Audiencia Nacional, and the Government of Spain — makes four specific demands.

First: the immediate opening of a formal investigation into Amazon, Oracle, Palantir, Microsoft, and Google for their role in providing technology and services that substantially aid and abet the commission of international crimes.

Second: the investigation of senior executives — CEOs, chairs, and board members — under the doctrine of command responsibility for their knowing authorization of these contracts and their failure to prevent their technology's use in atrocities.

Third: the issuance of subpoenas for internal documents and communications to determine the full extent of what these companies knew, and when they knew it.

Fourth: the pursuit of charges of complicity in genocide, crimes against humanity, and war crimes, and the seeking of arrest warrants for executives should the evidence warrant it.

This is not a call for dialogue. It is not a call for a review panel or a corporate social responsibility audit. It is a call for criminal prosecution. It is the same call the world should have made to IBM and Ford and Standard Oil in 1946 — and didn't, because profit and political alliance prevailed over justice. We cannot allow that to happen again.

This Petition Was Built on the Shoulders of Survivors

I want to be honest about something. I started this petition not because I believed it would be easy. I started it because I believe that the alternative — doing nothing, watching the

documentation pile up while the executives collect their bonuses — is morally intolerable for anyone who has spent years working in and around this industry.

I built my career in Silicon Valley. I worked inside Amazon Web Services. I know how these companies operate, how they think, what they value and what they dismiss. I know that the only language that reaches them is the language of consequence. Not shame — they have demonstrated they have none. Not public pressure — they have demonstrated they can absorb it. Consequence. Legal, financial, personal, criminal consequences.

This petition is drafted in solidarity with the parents who have lost their children in Gaza. It stands with the children who have been orphaned. It stands with the aid workers killed while delivering food and the journalists killed while telling the truth. It stands with the families documented on ParentsPlea.com — people whose names and stories deserve to be on the record, not erased by a platform compliance officer following an Israeli government takedown request.

They cannot bring this case to a Spanish court alone. That is what the petition is for. That is what your signature is for.

As of this writing, many people have signed this petition. We need hundreds of thousands to reach our first milestone — and we need far more support to build the kind of documented global demand that gives the Spanish Prosecutor's Office the political cover to act.

I understand why signing a petition feels small against the scale of what we are confronting. I understand why people who have been following this for two plus years feel that signatures are not enough. They are right that signatures alone are not enough. But a petition submitted to a government with a legal tradition of universal jurisdiction, backed by thousands of signatories from dozens of countries, with documented corporate contracts and named executives and verified sourcing — that is not a symbolic gesture. That is an evidentiary record. That is the beginning of a dossier. That is pressure applied to the exact institution that has the authority to act.

Every person who reads this book and does not sign is making a choice. Every person who reads this book and does sign is making a different one.

The petition is live at:

https://actionnetwork.org/petitions/petition-to-the-spanish-government-and-prosecutors-office-investigate-tech-giants-for-complicity-in-genocide-and-war-crimes-in-occupied-pal estine

Sign it. Share it. Give it to every person you know who has read these chapters and felt the weight of what they contain. The movement needs numbers. The Spanish judiciary needs to see that the world is watching and that the world is demanding action.

The Arc Is Long. We Have to Bend It.

I close this chapter the way I have always believed — not with despair, but with the stubborn conviction that evil men and their machines do not prevail. History is long and it is littered with their kind.

But history does not bend toward justice by itself. It bends because people push it. Because petitions get signed and cases get filed and courts get petitioned and executives who thought they were untouchable find out, one morning, that a judge in Madrid has issued a warrant with their name on it.

Imagine a world where we cut the grip of corporations lobbying for and committing never-ending wars. A world where Big Tech firms lose their monopolistic power. A world where Silicon Valley billionaires and genocidal executives are effectively grounded — unable to fly their private planes or sail their yachts to any country that honors the rule of law.

That world is not a fantasy. It is a legal mechanism, a signed petition, and a courageous prosecutor away.

Sign the petition. Then turn the page.

--- The petition to Spain is one instrument. The European regulatory revolution is another — and it is already underway, independent of any single petition or advocacy campaign. What Chapter Fifteen documents is the convergence of legal frameworks, governmental decisions, and institutional withdrawals from American platforms that is quietly dismantling the infrastructure of dependency that Big Tech has spent thirty years constructing. The criminal accountability I am demanding in Chapter Fourteen and the structural accountability the EU is building in Chapter Fifteen are not competing strategies. They are the same strategy, operating at different speeds on different legal terrain.

CHAPTER FIFTEEN — THE EUROPEAN REVOLT

Why Breaking Up Big Tech's Monopoly Starts With Accountability

Something is happening in Europe that the executives of Amazon, Google, Microsoft, Oracle, and Palantir are watching very carefully. And they are worried.

Not because of fines. They can absorb fines. Not because of regulations. They can hire lobbyists and lawyers and drag compliance timelines out for years. They are worried because, for the first time in the history of the digital economy, governments are not just threatening consequences — they are quietly, systematically, walking away. And when governments walk away from a monopoly, the monopoly begins to crumble.

This chapter is about why that is happening, how far it has already gone, and why the legal action we are demanding in Spain — the investigation and prosecution of Big Tech executives for their role in enabling genocide — is not separate from this European revolt. It is the spark that could ignite it completely.

The Trust Has Already Broken

Let me start with a fact that should send a chill through every boardroom in Silicon Valley. In April 2026, France's minister for public accounts announced that France would move all government computers — the computers of 2.5 million civil servants — from Microsoft to Linux. The stated reason was not cost, though the savings are enormous. It was sovereignty. French minister David Amiel put it plainly: France "can no longer accept that our data, our infrastructure, and our strategic decisions depend on solutions whose rules, pricing, evolution, and risks we do not control." He said: "Digital sovereignty is not optional." (Source: Euronews, "France to Ditch US Platforms Microsoft Teams, Zoom for 'Sovereign Platform,'" January 2026, https://www.euronews.com/next/2026/01/27/france-to-ditch-us-platforms-microsoft-teams-zoom- for-sovereign-platform-amid-security-con)

France is not alone. Denmark's Ministry of Digital Affairs has committed to replacing Microsoft Office 365 entirely. Germany's state of Schleswig-Holstein has moved 80 percent — 30,000 workplaces — to Linux, saving over €15 million annually in licensing fees. Austria has adopted a national charter on digital sovereignty, endorsed by all 27 EU member states at a summit in late 2025. Italy, the Netherlands, and Germany collectively are migrating over 800,000 government workstations away from American platforms. (Source: Tuta, "France Ditches Microsoft for Linux to Achieve Digital Sovereignty," April 2026, https://tuta.com/blog/countries-ditching-microsoft-choosing-linux-digital-sovereignty)

These are not symbolic gestures. They are strategic withdrawals. And they are being driven by something that goes deeper than politics or price increases — though both play a role. They are driven by a fundamental collapse of trust.

The CLOUD Act: The Fine Print That Changed Everything

The turning point came when European governments were forced to reckon with a law most of their citizens had never heard of. The U.S. CLOUD Act, passed in 2018, gives American

authorities the power to compel American companies to provide data stored anywhere in the world — including on servers physically located in Europe — without going through European courts, without notifying European governments, and often with a non-disclosure order that prevents the company from even telling the customer their data has been seized.

What this means in plain language is this: it does not matter where Microsoft, Amazon, or Google stores your data. If the U.S. government wants it, they can get it. And the company is legally prohibited from telling you.

Microsoft's own lawyer confirmed this under oath before the French Senate in June 2025. He could not guarantee that French data stored in European Microsoft data centers was safe from silent U.S. government access. (Source: Business Tech Africa, "France's Microsoft Exit — Loss of Trust and Digital Border Redraw," April 2026, https://www.businesstechafrica.co.za/technology/2026/04/22/frances-microsoft-exit-loss-of-trust- and-digital-border-redraw/)

This is a fundamental sovereignty problem. And GDPR — Europe's landmark data protection regulation — makes it a legal problem too. The European Court of Justice's Schrems II ruling already invalidated the EU-U.S. Privacy Shield agreement on precisely this basis: that U.S. law gave American authorities access to European personal data in ways that were incompatible with European privacy rights, and that EU citizens had no effective judicial redress.

The conflict is irreconcilable. European organizations using American cloud platforms are, in the view of multiple European data protection authorities, in ongoing structural breach of GDPR — not because of a specific incident, but because of the architecture itself. You cannot simultaneously comply with GDPR and guarantee that the U.S. government cannot access your data if your cloud provider is an American company. It cannot be done. The law does not allow it.

And here is what crystallized the danger for every European government watching: when the United States sanctioned members of the International Criminal Court for investigating Israeli war crimes, those sanctioned individuals were locked out of their Microsoft email accounts and web services. Overnight. Remotely. Without warning. A judge trying to do their job found that their ability to operate had been severed by a decision made in Washington. Germany's Schleswig-Holstein minister said directly: "The ICC case has demonstrated once more, unfortunately, how dependent institutions, governments and private companies are on proprietary systems." (Source: The Register, "Europe Gets Serious About Cutting US Digital Umbilical Cord," December 2025, https://www.theregister.com/2025/12/22/europe_gets_serious_about_cutting/)

That was the moment when digital sovereignty stopped being theoretical. "The Executive Who Got Fired — And Why It Changes Nothing"

In mid May, 2026 — as I was completing this book — Microsoft quietly announced that Alon Haimovich, the General Manager of its Israeli subsidiary, would be leaving his position. Several other senior managers departed alongside him. The announcement carried no explanation. No apology. No acknowledgment of what had happened on his watch.

Let me tell you what happened on his watch.

For years, Israel's Unit 8200 — the military intelligence division whose alumni populate the most powerful surveillance technology companies on earth, and whose capabilities I documented in the chapter on The Code Bomb — had been using Microsoft's Azure cloud platform to store

the intercepted phone calls of millions of Palestinians in Gaza and the West Bank. According to an investigation published by The Guardian, +972 Magazine, and Local Call in August 2025, the system was built around a specific, documented goal: to intercept and store a million Palestinian calls an hour. It was not a side project. It was not a contractor going rogue. Leaked documents revealed that Microsoft CEO Satya Nadella himself was briefed by a Unit 8200 commander in 2021 — and approved a custom, isolated partition within Azure built specifically for military surveillance purposes. (Source: Truthout, "Microsoft Ousts Head of Israeli Branch Over Use of Tech to Spy on Palestinians," May 2026, https://truthout.org/articles/microsoft-ousts-head-of-israeli-branch-over-use-of-tech-to-spy-on-pal estinians/)

By July 2025, the system held approximately 11,500 terabytes of Israeli military surveillance data — the equivalent of 200 million hours of recorded audio — stored not in Israel, but in Microsoft's data centers in the Netherlands and Ireland. Both are European Union jurisdictions. Both are subject to the General Data Protection Regulation. Both are, therefore, legally answerable to European courts. (Source: TechTimes, "Microsoft Removes Israel General Manager After Azure Stored 200 Million Hours of Palestinian Surveillance Recordings," May 2026, https://www.techtimes.com/articles/316642/20260514/microsoft-removes-israel-general-manage r-after-azure-stored-200-million-hours-palestinian.htm)

When that reporting broke, Microsoft launched an internal investigation. The investigative team that arrived in Israel weeks later — according to Israeli financial daily Globes — was not there primarily out of moral concern. They were there because they were afraid. Afraid of European regulators. Afraid of GDPR enforcement. Afraid of the legal exposure created by the fact that surveillance data collected against a civilian population was being processed and stored on European soil, under European jurisdiction, by a company subject to European law. The Dutch foreign minister said publicly that if criminal offenses were indicated, legal proceedings could be initiated. Amnesty International filed formal complaints. Activists at the data center campus near Middenmeer in North Holland called on workers to stop until Israeli intelligence data was removed from the servers. (Source: Data Center Dynamics, "Microsoft Launches Investigation into Israeli Military Surveillance Agency's Use of Its Cloud Platform," March 2026, https://www.datacenterdynamics.com/en/news/microsoft-launches-investigation-into-israeli-milit ary-surveillance-agencys-use-of-its-cloud-platform/)

What Microsoft's investigation found — according to Globes — was not limited to Unit 8200. Multiple Israeli Ministry of Defense units had been using Azure in ways that violated Microsoft's own terms of service. The Israeli subsidiary had not been transparent with Microsoft's headquarters. Management failures, internal reporting gaps, and a pattern of deliberate non-disclosure had created a system in which a major American corporation was processing the mass surveillance of a civilian population on European servers, without meaningful oversight, for years.

Alon Haimovich got fired.

And I want to be precise about what I mean when I say that changes nothing.

The Revolving Door Is Open

It is worth asking, when a Silicon Valley executive is terminated for enabling surveillance of a civilian population: where does he go?

He goes to Amazon Web Services. Or Google Cloud. Or to a contractor like DXC Technology — one of the largest IT services companies in the world — which manages government and military technology contracts across multiple conflict zones. Or he goes back to the Israeli Ministry of Defense, or to one of the Unit 8200 alumni startups that have seeded the global surveillance technology market for decades. Or he finds a position at Palantir, whose CEO Peter Thiel has made no secret of the company's military surveillance mission. The door between Silicon Valley and the Israeli security apparatus does not swing in one direction. It revolves.

The activist group No Azure for Apartheid put it plainly in their statement following the firing: "Microsoft has tried to quietly say goodbye to alleged war criminal Alon Haimovich, who oversaw the development of Azure tools for the Israeli military which helped accelerate the first AI-powered genocide." (Source: Truthout, May 2026) But the statement also carried a warning that the rest of the technology industry would prefer you not hear: Azure is not the only tool. "Microsoft has a footprint in all major military infrastructure in Israel. Microsoft supplies cloud, AI, computing, storage and advanced AI models to the Israeli military to be used not just by Unit 8200 but also Mamram, Ofek, and specific naval, air and ground units in the Israeli military."

Haimovich was the named executive on a structure that runs through the entire company. His departure is a human resources action. It is not accountability.

And here is where I need to say something directly to the Europeans reading this book — to anyone in the Netherlands, in Ireland, in Spain — whose national soil was used to store this surveillance data without their knowledge or consent.

Your government has the legal jurisdiction to investigate this. Not just Microsoft the corporation. The individuals who made the decisions. The executives who approved the contracts. The engineers who built the isolated partition. The managers who knew what the data contained and said nothing to European regulators. The GDPR is not a fine to be paid and forgotten. It is a legal framework with criminal enforcement provisions. The EU AI Act, now entering force, imposes accountability standards on high-risk artificial intelligence systems that go well beyond what any of these companies have voluntarily disclosed. The European Court of Justice has already ruled, in Schrems II, that the architecture of American cloud computing is structurally incompatible with European privacy rights.

The firing of one executive does not resolve a structural incompatibility. It papers over it.

This is precisely why the petition at my website — NoEthicsInBigTech.com — calls for criminal investigation and prosecution of the executives responsible, not corporate human resources processes managed by the same corporations that committed the violations. I have been told this demand is unrealistic. That corporations cannot be held criminally accountable in this way. That executives will simply deny knowledge.

Tell that to the Dutch foreign minister who said legal proceedings can be initiated. Tell that to the European data protection authorities who filed formal complaints. Tell that to the two thousand Microsoft employees who signed the internal petition demanding the company stop supplying cloud services to the Israeli military — and who watched their colleagues get fired for staging on-site protests while the executives who built the surveillance architecture were quietly separated with whatever severance package Microsoft thought appropriate.

The people who get fired for speaking out are the people who had nothing to do with it.

The people who built it are now looking for their next assignment.

That asymmetry is not an accident. It is the system working as designed. And it will keep working that way — in Microsoft, in AWS, in Google Cloud, in whatever company Alon Haimovich works for next — until someone with a badge and a courtroom treats it as the crime it is.

The Palantir Problem: Your Health Records and a Genocide

If the Microsoft situation illustrated the danger of dependency, the Palantir situation in the United Kingdom illustrated something even more alarming: what happens when a company actively complicit in documented atrocities is given access to the most intimate data of 67 million people.

In November 2023, Palantir — the same company whose AI targeting platforms operate in Gaza, whose CEO calls peace activists "war activists," whose co-founder Peter Thiel does not believe democracy and freedom are compatible — won a £330 million contract to build the NHS Federated Data Platform, a system designed to integrate and process the health records of patients across England's National Health Service. (Source: The Lowdown NHS, "Palantir, the Controversy, the Contracts and the Campaign Against the FDP," April 2026, https://lowdown-nhs.info/topics/accountablility/palantir-the-controversy-the-contracts-and-the-ca mpaign/)

The backlash was immediate and it was broad. The British Medical Association raised concerns. Amnesty International UK expressed alarm. Privacy International, the Good Law Project, Corporate Watch, and a coalition of doctors, lawyers, patients, and human rights organizations launched the No Palantir in the NHS campaign. Their briefing, published by health justice charity Medact in early 2026, made the case directly: a company that provides AI-based systems used to target Palestinian civilians in Gaza and the West Bank has no business handling the sensitive health data of British patients. (Source: Medact, "Briefing: Concerns Regarding Palantir Technologies and NHS Data Systems," March 2026, https://www.medact.org/2026/resources/briefings/briefing-palantir-fdp/)

What did British hospital trusts do? Most of them quietly refused. According to NHS figures, fewer than a quarter of England's 215 hospital trusts were actively using Palantir's platform by the end of 2024. Leeds Teaching Hospitals told NHS England plainly that adopting Palantir's tools would mean it would "lose functionality rather than gain it." Greater Manchester's health system reported the same. By May 2025, after intense pressure and a £330 million investment, Palantir had managed to sign up fewer than a third of trusts — and UK ministers were reportedly considering triggering a break clause in the contract as early as February 2027. (Source: Democracy for Sale, "Palantir's NHS Data Platform Rejected by Most Hospitals," May 2025, https://democracyforsale.substack.com/p/palantirs-nhs-data-platform-rejected-hospitals)

This is what accountability looks like at the ground level. Not a court verdict — not yet. But a rolling institutional refusal, driven by the documented record of what this company is and what it does. The NHS revolt against Palantir is a preview of what can happen when people understand who they are doing business with.

The European Parliament Has Already Said It: Break Them Up

The European Parliament and the European Commission are not waiting for a Spanish prosecution to begin dismantling Big Tech's grip on the continent. They have already built the

legal architecture. What they need is the political will to use it fully — and a criminal accountability precedent to force the issue.

The Digital Markets Act, which came into full effect in March 2024, designated six companies as "gatekeepers" — Alphabet, Amazon, Apple, ByteDance, Meta, and Microsoft — and imposed strict obligations on them designed to prevent the abuse of dominant market position. Fines of up to 10 percent of global annual turnover can be levied for non-compliance. For repeat offenders — companies fined three times within five years — the European Commission has the power to impose structural remedies, including forced divestiture: breaking up the company. (Source: European Parliament, "EU Digital Markets Act and Digital Services Act Explained," https://www.europarl.europa.eu/topics/en/article/20211209STO19124)

In January 2025, 18 former European heads of state wrote to European Commission President Ursula von der Leyen urging the Commission to pursue exactly this: a structural breaking of Big Tech monopolies across Europe's digital economy. The WeMove Europe campaign collected hundreds of thousands of signatures from citizens across the continent demanding that Google's monopoly over digital advertising be broken up, and broader action taken against Big Tech's unchecked power. European Parliamentary members — including German MEP Alexandra Geese, a key architect of the Digital Services Act — have publicly warned that "Europe risks being blackmailed by both American and Chinese tech giants." (Source: WeMove Europe, "Break Up Big Tech," https://action.wemove.eu/sign/2025-05-breakupbigtech-petition-EN/; The Register, December 2025)

By the end of 2025, the European Commission had opened or continued formal investigations into Apple, Meta, Google, and Microsoft under the DMA. The fines issued so far — €500 million for Apple, €200 million for Meta — are being described by critics as insufficient, and enforcement advocates are pushing for escalation. The framework exists. The investigations are open. What is missing is a catalyst powerful enough to transform regulatory action into structural consequence.

That catalyst is criminal prosecution.

What a Spanish Prosecution Changes

Here is the connection that the executives of these companies understand perfectly, even if they will never say so publicly.

A regulatory fine under the Digital Markets Act is a business expense. It is planned for, absorbed, appealed through courts for years, and ultimately paid from a fund that represents a rounding error on an annual earnings statement. Amazon, Google, and Microsoft have collectively paid billions in European fines and their market power has not diminished by a single percentage point. The fines are not deterrence. They are the cost of doing business.

A criminal prosecution is something entirely different. A criminal prosecution under Spain's universal jurisdiction framework, backed by an Interpol Red Notice and a European Arrest Warrant, means that the executives responsible — not the company's legal department, not a compliance officer, but the named CEOs and board members who authorized these contracts — become personally liable. Their assets in Europe become freezable. Their travel becomes restricted. Their ability to attend Davos, to accept honorary degrees from European universities, to sail their yachts in the Mediterranean, to give keynotes at conferences in Paris and Berlin —

all of it becomes legally precarious the moment a Spanish judge issues a warrant with their name on it.

And when executives become personally liable, boards change behavior. When boards change behavior, contracts get cancelled. When contracts get cancelled, the dependency that has kept European governments locked into American platforms for thirty years begins to break.

There is a reason why an estimated 97 percent of Europe's cloud infrastructure is still controlled by non-European — primarily American — providers, despite years of GDPR enforcement, despite the DMA, despite political declarations about sovereignty. (Source: TechClass, "Data Sovereignty in 2025: What EU Firms Must Know," https://www.techclass.com/resources/learning-and-development-articles/data-sovereignty-what-i t-means-for-european-businesses-in-2025) The reason is that changing infrastructure is expensive, disruptive, and politically difficult when the incumbent providers are embedded in everything — until the incumbent providers become criminal defendants. Then the political calculus changes overnight.

This is not speculation. It is the lesson of every major corporate accountability moment in history. Companies do not reform themselves through guilt. They reform when the personal cost of not reforming exceeds the institutional cost of changing. A criminal prosecution against the executives of Amazon, Google, Microsoft, Oracle, and Palantir would be the first moment in the history of Big Tech where that personal cost was real, documented, and enforceable across borders.

The Petition Is Part of This

When I started this petition — addressed to the Attorney General of Spain, the Central Courts of Instruction of the Audiencia Nacional, and the Government of Spain — I was thinking about Gaza. I was thinking about the 30,000 children who have been killed with the help of cloud infrastructure and AI targeting systems built and operated by American companies. I was thinking about the parents documented on ParentsPlea.com who have nowhere to take their grief because every institution designed to deliver justice has been either too slow, too pressured, or too corrupted by Western power to act.

But I have come to understand that this petition is also about something larger. It is about whether the European continent — which has produced the world's most sophisticated framework for digital rights and data protection — will have the courage to use that framework at its full strength against the companies that have violated it most comprehensively.

The case for Spain is the case for Europe. The prosecution of Big Tech executives in a Spanish court for their role in enabling genocide is the clearest possible signal to every European government, every hospital trust, every data protection authority, and every regulator who has been hesitating at the edge of full enforcement — that these companies are not too big to face consequences. That their executives are not too powerful to be held personally accountable. That the arc of European law, extended to its full reach, is longer than the financial resources of any corporation.

France has already said it: digital sovereignty is not optional. Germany is moving. Denmark is moving. The ICC moved after its chief prosecutor was locked out of his email by American political pressure. The momentum is real and it is accelerating.

What it needs is a legal landmark. A precedent. A moment when a Spanish judge looks at the documented record — the Project Nimbus contract, the Gospel targeting system, the Palantir data fusion platforms, the Lavender kill lists — and concludes that the executives who authorized these tools to be used in the service of documented atrocity must answer for it in a court of law.

That moment does not happen on its own. It happens because enough people demanded it. It happens because a petition with enough signatures gave the Spanish Prosecutor's Office the political cover to act. It happens because the documentation in this book — and in the hands of human rights lawyers and accountability organizations around the world — made inaction impossible to justify.

The petition is live. The signatures matter. The case is built. Now it needs the people who understand what is at stake to make their voice heard in the only language that courts recognize: a documented, public, and growing demand for justice.

https://actionnetwork.org/petitions/petition-to-the-spanish-government-and-prosecutors-office-investigate-tech-giants-for-complicity-in-genocide-and-war-crimes-in-occupied-pal estine

Sign it. Share it. Tell every person you know in Europe — in France, in Germany, in Spain, in the UK — that this petition is not only about Gaza. It is about whether they want to live in a world where the companies that hold their health records, their emails, their government data, and their children's information are the same companies that helped build a mass assassination factory.

Because that is the world we are living in right now. And it is the world we have the tools, the law, and the moral obligation to change.

--- The European revolt and the Spanish petition both operate at the level of institutions — courts, parliaments, regulatory bodies. But the foundation on which all institutional accountability rests is something more basic: a record. Names. Dates. Causes of death. Family testimony. ParentsPlea.com is that record. Chapter Sixteen is about what it is, how it was built, and why the powerful fear a database of the dead more than they fear a fine.

CHAPTER SIXTEEN — PARENTSPLEA

The Database of the Dead — and Why the Powerful Fear It I want to tell you about the most important work I have ever done that almost nobody has heard of.

Not the books. Not the documentary. Not the petition to Spain or the years inside Amazon or the nonprofit I built and then had to shut down when I realized the problem was bigger than any nonprofit could fix. The most important work is a platform that most people in technology would dismiss in thirty seconds — because it does not scale in the way Silicon Valley means when it uses that word. It does not monetize. It does not grow through viral loops or network effects or paid acquisition funnels. It grows the way grief grows: one family at a time, one name at a time, one testimony added by someone who needed the world to know that the person they lost was a person.

ParentsPlea.com now holds over 140,000 profiles.

More than 107,000 of those profiles document deaths from war and armed conflict — from Gaza, from Lebanon, from Iran, from Yemen, from Afghanistan, from Syria, from countries whose names appear in the casualty statistics that scroll past our screens and disappear. The rest document deaths from police brutality, from corporate negligence, from social injustice — the categories of harm that produce statistics but rarely produce faces.

Every profile on ParentsPlea has a face. That is the point.

Why I Built It

In every chapter of this book, I have documented how the powerful erase the evidence of what their power produces. Meta deletes the posts. YouTube removes the videos. The DOJ redacts the documents. Company platforms comply with government takedown requests at rates that should make your stomach turn. The testimony disappears. The archive is gone. The record does not exist.

ParentsPlea exists to make that erasure impossible.

Not technologically impossible — I am not naïve enough to think any platform is beyond reach. But morally and legally impossible in a way that matters: because if 140,000 names are documented, sourced, linked to news coverage, and accompanied by family testimony, then the people ordering the erasure have to work harder. They have to file more takedown requests. They have to justify more removals. They have to explain, to more people, why they are trying to make a grieving mother's testimony disappear.

The evidentiary value of ParentsPlea goes beyond bearing witness. I built it with three specific legal purposes in mind.

First: if there is ever going to be judgment in a court of law — in Spain, at the ICC, in any jurisdiction with the courage to act — prosecutors need to be able to identify victims and attach human testimony to the machinery they are prosecuting. The targeting algorithms in Chapter Two are abstract until you put a name on the person they targeted. The kill lists are numbers

until you find the family of one of the 37,000 people Lavender flagged and let them speak. ParentsPlea is that bridge — from the documented crime to the documented victim.

Second: if reparations are ever ordered — and history suggests they eventually are, long after the people who should have paid them have died comfortable deaths — we need to know who was harmed and what they are owed. Reparations without a registry of the harmed are just a number with nowhere to go. ParentsPlea is building that registry.

Third: if there is ever going to be a memorial worthy of the name — not a wall of numbers, not a press release, not a moment of silence that lasts exactly long enough for the cameras to capture it — we need biographies. We need the complete record of who was lost, how they died, who is left behind, and what that family wants the world to know. ParentsPlea is building that record.

What the Platform Actually Does

Each profile on ParentsPlea is built from fields that serve evidentiary as much as memorial purposes. A photograph, where available. A biography — not a summary, but the kind of life information that makes a person real to a stranger: age, gender, religion, marital status, children, occupation, the things they loved. Cause of death. The alleged responsible party — named, not hidden behind bureaucratic language. Location of death. Links to news articles that corroborate the account. And the family's plea — their words, in their own language, addressed to whoever is still willing to listen.

The platform is built on a combination of user submissions and curated staff research. My team — small, underfunded, working across time zones from offices that are themselves in the regions being documented — actively searches news sources, published articles, and publicly available reporting to build and verify profiles. We have also developed tools to index and crawl the web for relevant information. We have connections to communities in Gaza, Lebanon, Iran, and elsewhere. They do not have the resources of a newsroom. They have something more valuable: the trust of the families they are documenting and the tenacity to find the record when no one else is keeping it.

We also accept direct submissions from families and witnesses. When a profile is submitted, it goes through a verification process against publicly available sources before it is published. This is not a database I built about the dead. It is a database built by a dedicated team — and by the living — to honor the dead and document the crimes that killed them.

Why the Powerful Fear a Registry

There is a question I asked earlier in this book that I want to return to here with more force: why do we track sex offenders but not the Epstein class?

Sex offender registries exist because society decided that the crimes committed by a certain category of predator were too serious to leave undocumented. The public has a right to know where these people live. The victims have a right to see their abuser's name on a list. The registry is a form of accountability that does not require a prison sentence to function — it functions as permanent social and legal documentation of what someone did.

Now ask yourself: why do the executives of Palantir, Amazon, Google, Microsoft, and Oracle — whose tools have been documented as central infrastructure for the killing of tens of thousands of civilians — face no equivalent documentation? Why are the names of the people

who authorized Project Nimbus and the Gospel targeting system and the Lavender kill list not on a searchable public registry? Why are the redacted DOJ documents that protect the powerful men who attended Jeffrey Epstein's events still redacted, years after his death, while the names of low-level drug offenders are searchable in public court records?

The answer is power. The registry exists for the powerless because the powerful built it. The powerful have no equivalent registry because the powerful have no interest in building one.

ParentsPlea is, among other things, the beginning of a different registry. Not of individuals convicted in court — we are not a court. But of the corporate actors, the government actors, and the institutional actors alleged to be responsible for the deaths documented in our profiles. Every profile names its alleged responsible party. Those names accumulate. That accumulation is a form of public documentation that no settlement agreement with a nondisclosure clause can reach.

Jeffrey Epstein did not operate alone. The men whose names appear in the documents the DOJ has kept sealed for years have resources, lawyers, and political connections that have kept them invisible. The families of the children they harmed do not have those resources. What they have is testimony — and testimony, when it is documented and preserved and placed in a public archive with 140,000 other testimonies, becomes something harder to dismiss than a single complaint in a court that can be delayed, appealed, and buried for decades.

I want to do God's big data work in this world — cataloguing the devils' deeds so that justice, which the courts and governments have failed to deliver on schedule, can at least be demanded with evidence. Call it divine due diligence, delivered with open-source tools and the stubborn refusal to forget. The ones who commit crimes against children, against women and men, against the most vulnerable members of our society deserve to be on a list. Not hidden behind a redacted DOJ file. Visible. Named. Documented. Tracked.

We track sex offenders because society decided their crimes were too serious to leave undocumented. The crimes of the Epstein class — against children, against the rule of law, against the basic dignity of human beings who had no power to fight back — are no less serious. They are more serious, because the power that enables those crimes is also the power that prevents their prosecution.

ParentsPlea.com is building toward that accounting. One profile at a time. One family testimony at a time. One verified news source at a time. My work is not complete. It is a beginning.

How to Use It

If you have lost someone — to war, to corporate negligence, to police brutality, to any of the forces described in this book — go to ParentsPlea.com and submit their story. The form asks for the information listed above. You do not need a lawyer to submit. You do not need to prove anything beyond what you know and what can be linked to publicly available sources. Your testimony is evidence. Your family's story belongs in the archive.

If you are a human rights lawyer, a journalist, a researcher, or an advocate — the archive is searchable. The profiles are documented. The alleged responsible parties are named. Use it. That is what it is for.

If you are one of the executives named in this book: your name is in here too. Not as a profile of the dead, but as an alleged responsible party attached to the profiles of those whose deaths your tools are documented to have enabled. You should know that. The archive grows every day.

There is a name I need to say here that belongs in any honest accounting of what ParentsPlea.com is for. Refaat Alareer was a Palestinian poet and professor who lived in Gaza. He taught literature at the Islamic University of Gaza and edited a collection of short stories called Gaza Writes Back. He wrote one of the most widely shared poems of this entire era of conflict — a poem he titled "If I Must Die." He wrote it in the full knowledge of what he was living through. He wrote it as a man who understood that he might not survive what was coming. If I must die, you must live to tell my story to sell my things to buy a piece of cloth and some strings make it white with a long tail so a child somewhere in Gaza while looking heaven in the eye awaiting his dad who left in a blaze and bid no one farewell, not even to his flesh, not even to himself sees the kite, my kite you made,flying up above and thinks for a momentan angel is therebringing back love. If I must die let it bring hope, let it be a tale. In December 2023, Refaat Alareer was killed in an Israeli airstrike in Gaza. He was forty-four years old. He was a teacher, a writer, a father. He had written a poem about his own death so that the people who came after him would know what to do with it. He told them to keep living. To tell the story. To make a kite. His profile is on ParentsPlea.com. He is not there because we sat in the same room with him — we did not. He is there because people like him, who documented what these weapons do to a people, are being killed before they can keep documenting it. Someone has to keep the record when the witnesses are gone. That is what ParentsPlea is. That is what it has always been. Not a media organization. Not an NGO with institutional funding. A record. Because just peace requires us to inspect the deaths that have taken place — who is doing the killing, who is suffering the amputations and the horrors of these wars that we are witnessing. And because Refaat Alareer told us what to do if he must die. He must die for us to keep living. For us to tell his story. We are telling it. I want to be specific about what methodology we use, because the critics who have not built what we have built will question it. The profiles on ParentsPlea.com are built from open-source published reporting — the work of journalists, human rights organizations, and official governmental and institutional sources. We verify each profile against a published source before it is added to the archive. We link directly to those sources. We name the alleged responsible party. We preserve the family testimony where it exists. We play the national anthem of the country where each person lived and died — a small act of dignity that costs nothing but means something. The journalists whose reporting underlies many of our Gaza and Lebanon profiles are themselves, in many cases, no longer alive. They were targeted. Amal Khalil (ParentsPlea.com/amal-khalil), Ghada Dayekh (ParentsPlea.com/ghada-dayekh), and Suzan Khalil (ParentsPlea.com/suzan-khalil) are among the Lebanese journalists whose deaths we have documented — journalists allegedly killed by Israeli military action while doing exactly what journalists exist to do: bearing witness in conflict zones. Their deaths are an attack not only on them but on the record itself — on the evidence that site like ParentsPlea exist to preserve. We also document deaths that fall outside the most-covered conflicts. Alejandro Carranza Medina was a Colombian fisherman, not a journalist, not a combatant, not a person whose name would appear in a Senate report. He was forty years old. He went to sea from Santa Marta and did not return because a U.S. military airstrike in the

Caribbean hit the vessel he was aboard — a vessel the Colombian government later stated had suffered mechanical failure and was stranded at sea, not engaged in drug trafficking. His wife, Katerine Hernández, filed a complaint with the Inter-American Commission on Human Rights accusing the U.S. government of murder and extrajudicial killing. She asked: "Why did they take his life like that? Fishermen have the right to live. Why didn't they just detain them?" (Source: ParentsPlea.com/alejandro-carranza) Her question is on ParentsPlea. His photograph is on ParentsPlea. His biography, his religion, his marital status, his cause of death, and the name of the government alleged to be responsible — all of it is there, with sources, free to access, free to cite, free to use in any legal or journalistic proceeding that needs it. No advertising. No fees. No corporate sponsors. This is the database the powerful fear. Not because it is large — it is, by the scale of the atrocities it documents, still tragically small. But because it refuses to let the algorithm erase what the weapon produced. Because it keeps asking, in the voice of every family who trusted it: why? That question has no expiration date. Neither does the record. --- ParentsPlea.com documents what has already happened. ConflictTour is an attempt to change what happens next — before the next generation of young people is handed a recruitment brochure designed by an algorithm that has been gaming their social feeds since they were twelve. Chapter Seventeen is the affirmative argument: not just against the war machine, but for the alternative. Witness. Service. A different kind of courage than the kind the Twitch stream is selling.

CHAPTER SEVENTEEN — CONFLICTTOUR

An Alternative to War — and to the Machine That Profits From It

Everything in this book, up to this chapter, has been an argument about what is broken. The targeting systems. The cloud contracts. The revolving door. The censorship. The propaganda. The money. The immunity of the people at the top of every one of these systems from any consequence for what those systems produce.

I am aware that a book of this length about systems this broken can leave a reader in a particular kind of despair — the despair of understanding exactly what is wrong with no clear sense of what to build instead. I have felt that despair myself. I felt it in February 2022 when I closed the nonprofit. I felt it every time I watched a company fire an engineer for having a conscience and replace them with someone who did not. I felt it on the days when the petition signature count was moving slowly and the bombing of Gaza was not.

But I have never felt it for long. Because alongside the documentation of what is broken, I have always been building something.

This chapter is about the newest thing I am building. It is called ConflictTour. And it is, in my judgment, one of the most direct answers available to the question that has haunted this entire book: what do we build instead of a war machine?

How They Are Already Recruiting Your Children

Before I tell you what ConflictTour is, I want to tell you how the machine gets there first.

In my documentary *Forever Peace Now*, I interviewed Chris Velazquez — a veteran and lifelong gamer who spoke at one of our Ethics in Technology events about what he had witnessed from the inside of the gaming world. What he described was not a conspiracy. It was a strategy, documented in public, operating in plain sight, aimed at a very specific target: young men who play video games.

According to the Entertainment Software Association, in 2020 there were an estimated 214.4 million video game players in the United States. Shooter games were the third most popular genre on the market. Military and war-themed first-person shooters accounted for 26.9 percent of all games sold in the U.S. in 2018. An estimated 80 percent of male gamers aged 18 to 34 — and 61 percent of men aged 35 to 54 — play first-person shooter games on a regular basis.

The U.S. military recognized these numbers and acted on them. It launched its own e-sports teams, competitive gaming tournaments, Amazon Twitch channels, and dedicated social media accounts — specifically for advertising on streaming platforms and recruiting from the gaming community. The result, as Velazquez described in the documentary: 13,000 new recruiting leads generated in the first six months of 2020 alone.

Velazquez — speaking as both a veteran and a gamer — put it plainly: "As a gamer, I don't want to see our hobby used to promote anything other than fun and community. As a veteran, I don't want to see another generation of veterans. I don't want to see more children pulled into a system that leads to 22 individuals taking their life each day."

That figure — 22 veterans dying by suicide every day — is the number that the recruiting Amazon Twitch stream never shows. What young gamers see are military personnel playing alongside them, the simulation of combat made glamorous and frictionless, the skills they have developed for entertainment presented as directly transferable to the battlefield. What they do not see are the odds: the far greater likelihood that military service produces trauma, injury, military sexual assault, or other serious and life-altering consequences than it produces an e-sports career.

The military recruiter meets a young person before they have seen what war actually does to the people it reaches. That is the sequence ConflictTour is designed to interrupt.

The Question That Became a Program

In Chapter Thirteen, you read about the families on ParentsPlea.com — the parents who submitted the profiles of their children, the siblings who wrote the testimonies, the spouses who provided the photographs. The stories on ParentsPlea are read after the loss. But I kept asking myself: what if someone could have read them before the choice?

What if, before a young person enlisted in the military, they could sit beside a refugee and hear what fleeing bombs actually sounds like? Not on a screen. Not as a clip that the algorithm serves between advertisements. In a room. At a kitchen table. From a person who survived something that the recruiter's Amazon Twitch stream never mentions.

What if, before signing up to serve a government's military-industrial complex, they could serve a family instead?

What if they could document the testimony of a genocide survivor — add that testimony to a legal archive, understand that it might be used in a future prosecution, feel in their bones that this survivor is not an enemy but a human being exactly like them — and then decide, from that position of knowledge, what kind of service they wanted to give their life to?

That question became ConflictTour.

What It Is

ConflictTour is planned to be a six-month structured humanitarian service program — an alternative to military enlistment. It is not a gap year. It is not tourism. It is not an adventure program for people who want a story to tell at dinner parties. It is intensive, demanding, emotionally costly work that requires the same courage as military service and produces something entirely different.

The program is planned to run in two tracks.

The Local Service Track keeps participants in their home country, working alongside refugee families and war-affected communities — people who arrived with nothing and are navigating a government system that was not designed with them in mind. Local Track participants help families apply for asylum and benefits, provide language translation support at medical appointments, document stories for legal advocacy, connect families to community resources, and build the relationships that reduce the isolation that displacement produces. This track launches in January 2027.

The International Track sends participants to former war zones that are currently deemed safe by international travel advisories — where they work alongside local partner organizations to document the impact of war on civilian infrastructure, interview survivors for legal and

historical archives, and assist with community-led rebuilding. This track is expected to launch in September of 2027.

Both tracks begin with two months of structured training: trauma-informed communication, legal documentation basics, refugee resource navigation, cross-cultural language skills, and — critically — self-care and vicarious trauma preparation. This last element matters more than most programs acknowledge. The work of sitting with grief changes a person. The training prepares participants for that change, rather than pretending it will not happen.

Upon completion, participants receive the ConflictTour Certification in Humanitarian Witness — verifiable training in trauma-informed communication, refugee advocacy, and ethical documentation. But more than the credential, they carry the lived experience of having stood beside people in their hardest moments. That is not something a resume line can fully capture. It is, however, something that changes how a person votes, how they talk about war, how they process the news, and what they are willing to demand from the people who govern them.

Why This Is the Answer to Chapter Six

In Chapter Six, I documented Jim Taiclet's April 2026 earnings call, where the CEO of Lockheed Martin called the current geopolitical moment a "golden opportunity" for the defense industry. I documented the $901 billion FY2026 Pentagon budget and the $1.5 trillion proposed for FY2027. I documented the no-bid contracts, the NATO shakedown, and the merger of Silicon Valley with the military industrial complex.

ConflictTour is my direct answer to all of that — not a political answer, not a petition, not a legal argument, but a structural one. The military industrial complex depends on a supply of young people who are willing to enlist. That supply is not unlimited. It is produced by a specific set of conditions: economic precarity, cultural glorification of military service, the Amazon Twitch streams Chris Velazquez described, the FPS gaming pipelines the U.S. Army built deliberately — and the absence of any visible alternative that offers the same sense of purpose, service, and belonging that the recruiter promises.

ConflictTour proposes to fill that absence.

A young person who has spent six months helping a Syrian mother navigate an asylum system in their own city, who has heard her describe what she left behind, who has watched her children learn to trust adults again after witnessing things that adults should never let children witness — that young person makes a different decision at the recruiting office than one who has not had that experience. Not necessarily the same decision. But a different one. An informed one.

That is what ConflictTour is built to do.

The Pledge

Every ConflictTour participant makes a commitment before they begin. It covers five things.

They commit to their role as witness and servant — not combatant, not tourist. They explicitly agree not to participate in military service or take up arms during their enrollment.

They acknowledge the risks of the International Track with clear eyes — ConflictTour does not provide insurance, medical evacuation, or extraction services. Participants are adults who assume full responsibility for their own safety. This is not a liability waiver buried in fine print. It is a direct, honest statement that this work is serious and that the organization respects participants enough to tell them the truth about what it costs.

They pledge to treat every refugee, survivor, and community member they encounter with dignity and respect — that their stories are not content to be consumed but sacred testimonies to be honored.

They understand that their work may include documenting testimony for legal purposes — contributing directly to the kind of evidentiary archive that ParentsPlea has been building and that future prosecutions will need.

And they understand the consequences of violating any of this: immediate dismissal, no refund.

The pledge is not bureaucratic. It is the foundation of the work. A ConflictTour participant who breaks it is not just in breach of a terms of service. They have betrayed the families who trusted them enough to let them in.

What the Partners Bring

ConflictTour will work with partner organizations that have been doing this work for decades. Trauma counselors who understand what listening does to a person over months of sustained exposure to grief. NGOs with deep roots in the communities being served — the kind of organizations that know which families are ready to share their stories and which are not, which interpreters are trustworthy and which are not, which questions cause harm and which create space for healing.

These partners are not travel agents. They are not evacuation services. They do not manage participants' logistics or personal needs. They do something more valuable: they open doors into communities that international visitors cannot access alone. They provide the context that turns documentation into understanding. They share expertise that took years of field presence to earn.

In return, ConflictTour asks them to accept participants who arrive prepared and self-sufficient — people who have done the training, read the risk acknowledgments, and shown up ready to serve rather than ready to be managed.

The relationship is one of mutual respect between organizations doing different parts of the same work. ConflictTour trains the witnesses. The partners place them where witnessing matters most.

This Is Technology in Service of Humanity

In Chapter One, I argued that AI can be used for extraordinary good — for medicine, for climate, for bearing witness to war crimes. I argued that the question is never whether technology is good or evil, but who builds it and what they intend.

ConflictTour idea is built on that same argument, extended to organizational design. The training curriculum uses digital documentation tools — including platforms like ParentsPlea — to teach participants how ethical archiving works. The program uses technology to match participants with partner organizations across time zones and geographies. The legal documentation training teaches participants to use open-source intelligence tools and digital archives in ways that serve human rights rather than military targeting.

This is technology in service of humanity. Not the technology of Palantir, which builds AI to find and kill people. The technology of ConflictTour, which trains people to find and document the stories of the people Palantir's technology kills.

That contrast is not accidental. It is the point.

An Invitation

ConflictTour is new. It is built by one person who believes that the alternative to the military industrial complex is not passivity — it is a different kind of service, built with intention, grounded in witness, and offered to young people before the recruiter's Twitch stream gets there first.

If you are a young person reading this book: the Local Track launches January 2027. You do not need to travel to a war zone to do this work. You can do it in the city where you already live, alongside families who arrived there with nothing and who need someone to stand beside them.

If you are a parent reading this book: this is the alternative you have been looking for. It is not risk-free — honest service rarely is. But the risks are documented, the training is real, and the work is meaningful in a way that the recruiter's pitch will never tell you about.

If you are an NGO, a trauma counselor, a funding organization, or a service provider who wants to build this with us: the partnership inquiry form is at ConflictTour.com. We need field partners, trauma professionals, and organizations willing to help train the next generation of humanitarian witnesses.

If you are one of the executives named in this book — the architects of the Gospel and Lavender and Project Maven and Project Nimbus: this chapter is also for you. Not because I expect you to change. I have spent enough time in this industry to know better. But because I want you to understand that for every system you build to make war more efficient, there are people building systems to make peace more visible. You have more money. We have more time. History suggests that matters.

What ConflictTour Cannot Do — and What It Can

I want to be honest about the limits of what I am building.

ConflictTour cannot stop a war. It cannot pass a law. It cannot freeze an asset or issue an arrest warrant or compel a billionaire to answer for what his company's targeting algorithms did to a family in Rafah. Those instruments are what the earlier chapters of this book are about.

What ConflictTour can do is change who shows up at the conversation about war — and what they know when they get there. A generation of young people who have sat with survivors, documented testimony, navigated asylum systems alongside refugee families, and carried the weight of witnessing real consequence — that generation votes differently. Advocates differently. Demands differently. Refuses differently. And they do not enlist on the basis of a Twitch stream.

The military industrial complex calls genocide a golden opportunity. ConflictTour calls it a reason to show up with a notebook instead of a weapon and make sure the record is kept.

That is service. That is courage. That is what peace looks like in action.

Witness. Serve. Choose Peace.

ConflictTour.com

--- The documentation on ParentsPlea, the petition to Spain, the European regulatory revolt, the ConflictTour proposal — these are not a single campaign. They are expressions of a single conviction: that evil men and their machines do not prevail, but that they do not stop on their own, and that the arc of history bends only because people push it. Chapter Eighteen is how I push mine — and how I am asking you to push yours.

CHAPTER EIGHTEEN — LIVE IN PEACE, NOT JUST REST IN PEACE

ForeverPeaceNow.com, GoAmour.com, ParentsPlea.com, and the Argument for Love

I want to close this book the way I have tried to live my advocacy: not with despair, and not with false comfort, but with something harder than both — with documented hope. The kind of hope that has evidence behind it. The kind that is not a feeling but a structure. The kind that gives you something to do when you put this book down.

You have spent seventeen chapters inside a machine. Inside its contracts and its kill lists and its biometric databases and its revolving door and its erasure of Palestinian voices and its exploitation of workers in Nairobi and its injuries in the warehouses of New Jersey. Have you read about what Rumi — writing eight centuries ago — described better than any modern sociologist has managed: the city of Saba, where the wealth has piled so high that no one needs anything, and because no one needs anything, no one is grateful for anything, and because no one is grateful, no one wonders about the unseen world, and the people grow fat and bored and blind to what is wrong, and deaf to anyone who points it out.

The city of Saba is not a fable. It is Sand Hill Road. It is the campus of Palantir Technologies. It is the boardroom where Jim Taiclet called genocide a golden opportunity and meant it as a compliment. It is the culture that produces men who call empathy civilization's weakness while building AI systems that kill children in their homes. The grape clusters hang so low in those offices that they brush the faces of the inhabitants on their way to the earnings call. And the people of Saba feel bored with just the mention of prophecy.

Rumi's cure for the city of Saba is not a policy proposal. It is not a petition. It is not a prosecution — though we need all three. His cure is this:

Sit quietly and listen for a voice that will say, Be more silent. Give up talking, and your positions of power. Give up the excessive money. Turn toward the teachers and the prophets who do not live in Saba.

That is the beginning of an answer. It is not the whole answer. This book has seventeen chapters of the other parts. But it is where the answer starts: in the individual, in the moment of silence, in the recognition that the disease of over-richness is not cured by rearranging the institutions of the wealthy. It is cured by the willingness of individuals — inside and outside the machine — to hear the voice that says: enough. Be more silent. Turn toward something other than this.

This chapter is about that turning.

The Question I Asked Artificial Intelligence

Sometime ago — before this book existed, before the petition to Spain, before we had documented 140,000 souls on ParentsPlea — I was sitting with the idea of GoAmour.com and I did something that a person who had spent thirty years in technology might be expected to do. I

pulled up the latest version of ChatGPT 4.0 and asked it a question that no military contractor has ever asked their AI.

I asked it to make the best argument for love.

Not the best argument for a targeting system. Not the best argument for a no-bid contract or a cloud migration or a defense sector acquisition. I asked it to make the case for love — between people, between people and the planet, between people and nature, between people and their animals, and even between AI systems themselves.

And it answered. Coherently, comprehensively, without hesitation. It argued that love for people fosters empathy, understanding, and cooperation essential for addressing global challenges. That love between people and the planet motivates sustainable practices and conservation. That love between AI and people, nurtured with ethical considerations, ensures AI serves humanity's best interests. It made the argument across seven categories and concluded: love acts as a powerful force that unites, nurtures, and sustains. And then it said — in the plainest possible terms — yes, love is worth saving.

What struck me about that conversation, sitting with it afterward, was not that AI could describe love. I expected that. Everything Erich Fromm ever wrote about it is in the training data. Everything Rumi ever said is in there. What struck me was the asymmetry. The same architecture that could articulate love so completely was the same architecture that OpenAI had quietly removed the prohibition on military use from its terms of service. The same company that, when asked to argue for love, produced something luminous, had also made its tools available for drone surveillance and targeting operations.

The AI reflects the question. The question reflects the asker. The asker reveals their values.

When I asked for an argument for love, I got one. When Palantir asked its AI for a targeting system, it got one. The intelligence is the same. What differs is the human being holding it, and what they have decided it is for.

I share this because I think it is the most important thing I have learned in many years of writing about what technology does to humanity. The problem is not that AI cannot understand love. It can. The problem is that the people who control the most powerful AI systems have not asked it to.

That is a choice. And choices can be changed.

(Source: GoAmour.com, "Saving Love: A Conversation with Artificial Intelligence," https://goamour.com/saving-love-a-conversation-with-artificial-intelligence/)

What Erich Fromm Understood That Silicon Valley Never Has

In his 1956 book *The Art of Loving*, the psychologist and philosopher Erich Fromm made an argument that was radical then and remains radical now: love is not a feeling. It is a practice. It is an art — something that requires knowledge, dedication, and discipline, the same way painting requires knowledge, dedication, and discipline. And like any art, it can be practiced well or badly, learned or avoided, deepened over a lifetime or left to atrophy.

Fromm identified four elements of mature love: care, responsibility, respect, and knowledge. These four, he argued, must be present together — not any one alone, not three without the fourth. Care is the active concern for the life and growth of the person or thing you love. Responsibility is not a duty imposed from outside but a voluntary response to the needs of

another. Respect is the ability to see a person as they actually are — not as you need them to be, not as your projection of them — and to allow them to grow in their own way, for their own reasons. And knowledge is the willingness to go beneath the surface: to understand who a person actually is at their core, not just the story you have told yourself about them.

What Fromm was describing — care, responsibility, respect, knowledge — is the precise opposite of what Silicon Valley's business model requires. The attention economy is built on the systematic destruction of each of these four. Care is replaced by engagement — not whether content benefits you, but whether it keeps you scrolling. Responsibility is replaced by liability management — not a voluntary response to another's needs, but a legal calculation about what the company can avoid being held accountable for. Respect is replaced by targeting — not seeing the person as they are, but reducing them to a behavioral profile that predicts what advertisement they will click. And knowledge — deep, patient, costly knowledge of the other — is replaced by data extraction, which is the opposite of knowledge: it mistakes the measurement of a person's behavior for an understanding of their soul.

The city of Saba cannot understand love because love requires the recognition of lack — the recognition that there is something you do not have, something you need, something worth working toward. The people of Saba have no lack. The algorithm does not lack. The targeting system does not lack. It has 37,000 names and a 10 percent error rate it has already decided is acceptable.

The man who understood this — who built a platform to argue against it — was not a philosopher. He was an Iranian-American technologist who had sold colocation at Exodus and managed alliances at AWS and walked out of hospitals with a cane. He took Fromm's four elements — care, responsibility, respect, knowledge — and built them into a ten-point framework that he put on a website and called GoAmour.

The GoAmour Principles — A Framework for What Comes After

Let me walk through these principles. Not because they are new — most of them are older than language — but because the act of naming them in the context of this book is itself a statement of what this whole effort has been about.

The first principle is the Golden Rule, with emphasis on the word *others*. Do unto others as you would have them do unto you. The emphasis matters because the Golden Rule is usually recited as a constraint — don't do to others what you wouldn't want done to you. That is passive. The GoAmour version is active: do *unto* others. Take the initiative. Make the first move toward the other person's dignity and wellbeing, before they have earned it, before they have asked for it, before you know whether they will reciprocate.

The second principle is the freedom to choose your affiliations, ideologies, organizations, and beliefs without prejudice or judgment. This includes artificial intelligence. It includes the platforms this book has criticized. The freedom to choose is not the freedom to choose correctly — it is the freedom to choose, period, and to be respected for the choice you have made. The surveillance economy violates this principle at every level. It catalogs your choices in order to predict and ultimately shape them. A person whose choices are catalogued by an algorithm designed to keep them choosing the same things is not free in any meaningful sense of the word.

The third principle is the right to terminate. To cancel the subscription. To leave the platform. To walk away from the institution, the ideology, the AI system that no longer serves your values. Without long-term contractual obligations. Without being followed. Without your data being retained after you have decided you are done. This is a principle so basic that it should not need to be stated — but in a world where Facebook retains data on users who have deleted their accounts, where Amazon warehouse workers must give their bodies as collateral for the productivity the algorithm demands, where the terms of service that govern your digital life run to forty pages of legally binding language that no one reads — the right to terminate is a radical act.

The fourth principle is Fromm's directly: to love with knowledge, respect, care, and responsibility. To love not sentimentally, not possessively, not as acquisition or performance, but as practice. As the daily work of attending to another being. As the willingness to be changed by what you see in them.

The fifth principle is the preservation and sanctity of human rights and digital rights — for ourselves and for others. Not human rights as a talking point on a corporate values page. Not digital rights as a feature in a privacy policy that no one reads. Human rights and digital rights as things that actually cost something to defend: that cost your job, sometimes, and your safety, sometimes, and your relationship with people who would rather you kept quiet. They are worth those costs. Everything in this book has been an attempt to make that case.

The sixth principle is the willingness to support organizations and ideologies that promote life in its many forms. Not just human life. Not just the lives of people who look like you and believe what you believe. Life in its many forms — the child in Gaza, the fisherman in Colombia, the content moderator in Nairobi, the whale, the forest, the aquifer. All of it is worth defending because all of it is connected, and the severing of any part of the web diminishes the whole.

The seventh principle is the demand for a cure to hate and violence — and if a cure is not possible, the demand for judgment. Not in the next world. In this one. This principle is the basis for the petition to Spain. It is the basis for the pardon I will never extend to the executives of Palantir and Amazon and Google for what their tools have enabled. And it is the basis for the conviction — which I have held since I came to this country at eleven years old and have never abandoned — that evil men and their machines do not prevail. Not because the universe is just. But because enough people, in enough generations, refuse to stop demanding that it be.

The eighth principle is oneness. Oneness with fellow beings, with nature, with god, with forms of sentient life we do not yet understand. This is the second competing idea I named in the Author's Note: We are one. The whole of the surveillance economy, the whole of the military industrial complex, the whole of the AI targeting system is built on the first competing idea: I am one. My shareholders are one. My quarterly earnings are one. My rocket is one. I will be on Mars and you will not. Oneness is the opposite of that. It is the recognition that the boundary between you and the person whose house was bombed by a system your company built is not a moral boundary. It is a geographic accident. And geographic accidents do not create moral permission.

The ninth principle is the focus on quality of existence, not mere quantity. Not survival. Not the prolongation of biological function. Not the defense of life as a statistic on a population

chart. Quality: the richness of the experience, the depth of the relationship, the meaning of the work, the texture of the days. The children on ParentsPlea did not lose mere biological function when they died. They lost futures that were full of quality — full of things they would have loved and built and become. That is what the targeting algorithm removed from the world. Not units of labor. Not entries in a database. Quality of existence. Irreplaceable and unrecoverable.

And the tenth principle is the most radical of all: add your own. This list is not finished. It cannot be finished by one person or one book or one platform. The work of defining what it means to live with love in the world is not a project that ends. It is a practice that deepens. Whoever you are, reading this — you have something to add. You have seen something, lost something, built something, survived something that I have not. Your version of love is not exactly mine. It does not need to be. It needs to be yours, honest, and in line with the nine principles above.

That is the whole framework. That is GoAmour. (Source: GoAmour.com/principles, https://goamour.com/principles/)

What You Can Do Right Now

Sign the petition.

https://actionnetwork.org/petitions/petition-to-the-spanish-government-and-prosecutors-office-in vestigate-tech-giants-for-complicity-in-genocide-and-war-crimes-in-occupied-palestine

Go to ParentsPlea.com. If you have lost someone — to war, to corporate negligence, to any of the forces described in this book — submit their story. If you are a journalist, a lawyer, a researcher, use the archive. The profiles are there. The evidence is there.

Go to ForeverPeaceNow.com. Watch the documentary. Share it.

Go to ConflictTour.com. The Local Track launches in January 2027. This is the alternative. This is service without a weapon.

Go to GoAmour.com. Read the principles. Add your own. Practice the tenth principle — that this list is not finished and neither are you.

And go, if you are in technology, back to the chapter addressed to you. The Letter to Those Who Build the Machine. Reread the last paragraph. The machine does not stop unless the people running it decide it should. And that decision will not come from a board resolution or a policy update or an investor call. It will come from individual people, sitting in their offices in the Bay Area or in Bangalore or in Austin, asking themselves the question that Rumi's poem ends with:

Give up talking, and your positions of power. Give up the excessive money. Turn toward the teachers and the prophets who do not live in Saba. They will help you grow sweet again, and fragrant and wild and fresh, and thankful for any small event.

Any small event. The gratitude for any small event is what the city of Saba has lost. It is what the algorithm destroys — deliberately, structurally, because gratitude does not scroll. Gratitude does not click. Gratitude does not generate the outrage that fills the advertising inventory. But gratitude is what survives the bomb. Gratitude is what the families on ParentsPlea.com carry into the form when they submit the story of the person they lost. They are grateful that someone will hear it. That someone will document it. That someone will not let the algorithm erase it.

That gratitude is the opposite of the city of Saba. It is the beginning of a different kind of politics, a different kind of technology, a different kind of world.

What This Book Does Not Cover — and Why That Is Intentional I want to say something that most authors are too proud or too commercially cautious to say directly: this book is not complete. It was never meant to be. There are significant areas of Big Tech's damage to society that I have not addressed at length — not because they are unimportant, but because I refuse to write chapters I am not qualified to write. I will not give you AI-generated summaries of legal cases I have not studied, regulatory arguments I have not lived, or business models I have not worked inside. This book is built on thirty years of personal experience and direct documentation. Where that experience ends, I have tried to be honest about the boundary. Let me name the areas I have left to you — because naming them is itself an act of accountability, and because I believe you are capable of doing the research yourself. The App Store and Google Play Store monopolies — the 15 to 30 percent tax that Apple and Google extract from every digital transaction on their platforms, the developer suppression, the anti-competitive terms, and the antitrust cases that have been filed, litigated for years, and partially resolved in ways that change almost nothing in practice — deserve more than I can give them. The federal judge who ruled in August 2024 that Google had illegally monopolized the search market delivered the most significant antitrust ruling against a technology company since the Microsoft case of the late 1990s. I mention it here. I have not devoted a chapter to it. That chapter needs to be written, but it needs to be written by someone who has spent years inside search advertising, not by someone who spent thirty years in enterprise cloud sales. The destruction of local journalism by Google and Facebook — the systematic extraction of advertising revenue from news organizations that built the audiences those platforms now monetize, the recent shift to charging hundreds of dollars to display news links, the collapse of more than 2,500 local newspapers since 2005 and the news deserts that now cover much of the country — is one of the most consequential developments in American democratic life in the past twenty years. An uninformed citizenry is a manageable citizenry. The platforms that destroyed local journalism and replaced it with algorithmic outrage knew exactly what they were doing. I have named it here. I have not mapped it the way it deserves to be mapped. The AI copyright crisis — the systematic use of books, articles, photographs, music, film, and creative work produced by human artists and writers to train models that now compete directly with those artists for their livelihoods, without consent, without compensation, and without credit — is the labor story of our time disguised as a legal argument. The lawsuits by the New York Times, by hundreds of individual authors, by Getty Images and musicians and voice actors represent the beginning of a reckoning that has not yet resolved. The strikes of SAG-AFTRA and the WGA were early tremors of what is coming. I have named this issue in the Meta MCI chapter in the context of knowledge labor extraction. The broader creative economy story I leave to you. The gig economy — Uber, Lyft, DoorDash, Instacart, TaskRabbit — is the Amazon warehouse story told in a different vehicle on a different road. The deliberate misclassification of workers as independent contractors to avoid labor protections. The algorithmic management systems that can deactivate a driver without explanation, appeal, or recourse. The $224 million that gig companies spent in California to exempt themselves from AB5, the law that would have required them to treat workers as employees. The delivery drivers on bicycles dying at elevated rates with no workers' compensation coverage. And now: the replacement of those same drivers by Waymo, by robotics,

by remote operators in India doing the work for pennies while the platform captures the margin. The arc from worker to algorithm runs through the gig economy as clearly as it runs through the Amazon warehouse. I leave the full documentation of that arc to the researchers and organizers who have been living it. Section 230 of the Communications Decency Act — the law that gives platforms immunity from liability for content hosted on their services, and that has functioned as the legal wall standing between every documented harm in this book and any meaningful accountability — has been in the headlines enough that I trust you to find the argument. I will say only this: without Section 230 reform, nothing else in this book's accountability agenda is achievable in American courts. And with Section 230 reform designed by the wrong people, in the wrong direction, the speech protections that activists and journalists and ordinary people depend on could be destroyed in the process of trying to fix the liability protections that corporations abuse. It is the most complicated legislative question in digital rights law. I am not the person to resolve it here. I say all of this not to apologize for what the book does not contain, but to be honest with you about what a book can do and what it cannot. This one can document what I have witnessed. It can name what I know. It can argue for what I believe. It cannot be the last word on everything that is wrong with the technology sector — because the technology sector is wrong in more ways than any single person can document in a single lifetime, let alone a single book. What I am asking you to do — as someone who has now spent time inside these arguments with me — is to keep going. Do not put this book down and consider the matter closed. The areas I have named above need readers who will get angry enough about them to learn everything I have not had time to learn. The antitrust cases need citizens who understand what a monopoly actually does to a market and who demand that judges' findings produce structural change, not settlements. The journalism crisis needs readers who support local news the way they support local restaurants — with their attention and their dollars — before the last paper closes. The gig economy needs organizers who understand that the fight for the driver is the same fight as the fight for the warehouse worker, moved to a different platform. I am one person. This book is one document. The machine is large. It requires more than one of us pushing.

GoAmour Is Not a Platform. It Is a Practice. I want to clarify something about GoAmour that I should have said more plainly earlier in this chapter. GoAmour is not a product. It is not a startup. It is not a platform competing with the platforms I have spent this book criticizing. I am not building the next Instagram for people who care about peace. I am not building a social network for the ethically minded. I have spent thirty years watching the Silicon Valley impulse to build a platform for every human aspiration, to monetize every principle, to turn every idea about how to live better into a growth strategy and a Series A raise. I am not doing that. GoAmour is a set of principles. Ten of them, with room for more. They do not require an app. They do not require a subscription. They do not require your data, your attention, your time-on-platform, or your advertising dollar. They require only the thing that the surveillance economy has spent two decades training you not to do: to be present, to be deliberate, and to choose — consciously, repeatedly, in small moments and large ones — whether you are living by the first idea or the second. I am one. Or: We are one. You know which one I am asking you to choose. You have known since the first page. The practice of GoAmour is not something you do on a

website. It is something you do when you choose the local business over the algorithm's recommendation. When you call your congressperson instead of scrolling past the news story. When you talk to the person next to you instead of the device in your pocket. When you teach your child that empathy is not, as Elon Musk has declared, the fundamental weakness of Western civilization — but its greatest strength, and the thing that separates wisdom from weaponry, service from surveillance, love from data extraction. These are small things. The people who built the machines described in this book would call them naïve. They would say: the world does not run on principles, it runs on power. And they would be right, in the short term, and wrong in the long term, and they would be building their bunkers in Hawaii and Patagonia in the meantime. I am not building a bunker. I am writing a book. I am making a film. I am documenting the dead one profile at a time and demanding with every tool available to me that the people responsible for those deaths face accountability in a court of law. That is my practice. GoAmour is the name I have given it.

What I Cannot Stop Thinking About I want to close here — before the formal ending, before the call to action, before the petition link and the website addresses — with something more personal than anything else in this book. I cannot stop thinking about the children in Gaza. Not as a political category. Not as casualties in a conflict analysis. As children. I think about what it means to be five years old and hear the sound of a missile before you understand what a missile is. I know what that is. I was that child once, in Tehran, before my mother decided that protecting me was more important than staying in the country where she had built her life. I came to California. I learned English badly and then better. I started a company. I failed and nearly died and got up and tried again. I had the extraordinary luck of surviving to the age at which I could do something with what I had seen. The children in Gaza are not having that luck. The targeting algorithms documented in this book are not giving them the opportunity to grow up and write books about what they survived. The Gospel generates a hundred bombing targets a day, and some percentage of those targets are houses where children are sleeping, because the children's fathers are on a kill list generated by a system with a 10 percent error rate that its own designers considered acceptable. I am not able to make my peace with the word "acceptable" in that sentence. I have tried. I cannot. And so I keep going. I document. I petition. I sign. I demand. Not because I am certain it will work — I have been honest throughout this book about how uncertain the path to accountability actually is. I keep going because the alternative is to accept the world that the executives of Palantir and Amazon and Google have decided is acceptable. And I cannot do that. I do not know how. GoAmour's ninth principle is the focus on quality of existence, not mere quantity of prolonging existence. The children in Gaza are not surviving. But the ones who survive deserve a quality of existence that the world these executives have built is designed to deny them. My work is the refusal to accept that denial as inevitable. That refusal is what love looks like when it is translated into action. Not sentiment. Not a social media post with a flag emoji. A document. A petition. A chapter. A database with a face attached to every name. If I must die — Refaat Alareer wrote — let it bring hope. Let it be a tale. He did not survive. The children of Gaza are not all surviving. But the tale is being told. The hope is being built, one stubborn act at a time, by people who refuse to stop. That is what ForeverPeaceNow means. Not that peace has arrived. Not that the machine has stopped. That the

refusal to stop demanding peace is itself an act of love — and that love, practiced with knowledge and responsibility and respect and patience, is the only force in human history that has ever actually bent the arc. We are bending it. The arc is long. We have to bend it together.

The Mission: Live in Peace, Not Just Rest in Peace

The name ForeverPeaceNow comes from a mission statement that I have carried since the earliest days of this work: we should be able to live in peace, not just rest in peace.

That phrase is the most personal thing I can tell you about why I have done all of this. The friends I have lost. The communities I have watched be destroyed. The families on ParentsPlea.com who submitted the profiles of their children and their parents and their siblings not because it undoes the loss but because the alternative — silence, invisibility, erasure — is something they refused to accept.

They did not rest in peace. They were killed. By wars that produced golden opportunities. By targeting systems built by cloud companies. By social media algorithms that deleted their testimony. By governments that funded the killing and called it defense.

They deserved to live in peace. They were denied that. The least we can do — the absolute minimum that any person with a functioning conscience owes them — is to refuse to let the world that killed them continue unchanged.

This book is part of that refusal. The petition is part of that refusal. GoAmour.com and ForeverPeaceNow.com and ParentsPlea.com and ConflictTour.com are part of that refusal.

At times of senseless calls to violence and hate, I admire the government of Spain for standing up to the pressure of the United States. By not increasing its defense budget, by welcoming immigrants and undocumented people, by demanding justice for the victims of this horrific genocide that the world has been traumatized by witnessing online over the past two and a half years in Gaza and occupied territories. God bless Spain and all the righteous people from around the world.

Evil men and their machines do not prevail. History is long, and it is littered with their kind. The executives of Amazon, Google, Microsoft, Oracle, and Palantir will one day face the reckoning that history delivers to everyone who builds machines of this kind and hides behind corporate structures and quarterly earnings and legal disclaimers.

That day does not arrive on its own. It arrives because people like you read books like this one and decide that they are not willing to wait.

The arc is long. We have to bend it.

Sit quietly. Listen for the voice.

Then go. Love is worth saving. The planet is worth loving. The people who have been lost are worth documenting. The people who are still here are worth fighting for.

Live in peace. ForeverPeaceNow.

A Documented Record The arguments in this book were not formed at a distance from Silicon Valley. They were formed inside it — over more than a decade of conversations with the people who built it, questioned it, worked inside it.. Beginning in 2009 with the launch of the BizCloud interview series, and continuing through the founding of Ethics In Technology in 2013, I produced and published more than 350 videos across two YouTube channels — *Ethics In Technology* (youtube.com/@EthicsInTechnology) and *BizCloud* (youtube.com/@bizcloud) —

spanning interviews, community events, panel discussions, and firsthand commentary. The interviews include technology executives, veterans, civil liberties attorneys, religious leaders, activists, comedians, and ordinary Bay Area residents who showed up because they cared about where this industry was taking us. A subcollection of this work is permanently preserved at the Internet Archive under my account (archive.org/details/@vahidr), where it will remain publicly accessible long after the platforms that hosted it have moved on. These videos are not an endorsement of this book. They are a record of immersion — proof that I was present, was listening, and was documenting, for years, before a single page of this book was written.

Editorial research and structural editing assisted by Claude (Anthropic). The views, arguments, testimony, and conclusions in this book are entirely those of the author.

NoEthicsInBigTech.com | ForeverPeaceNow.com | GoAmour.com | ParentsPlea.com | ConflictTour.com

I Bear the Consequences I understand the risks of publishing this book. I am not naive about them. I am stating them plainly so that no one can later claim I stumbled into this without knowing what I was doing. They can threaten my body. The United States government and the CIA did exactly that to Khalid El-Masri — a German citizen taken off a bus in Macedonia in 2003, handed to the CIA, held for months, tortured at a black site in Afghanistan, sodomized, confirmed innocent, and then dropped on a road in Albania with no money, no documents, and no apology. The planes that moved him through that system were operated by Computer Sciences Corporation — the same company that later used its legal machinery to take my trademark, sue me, and drive me into bankruptcy. I know what this machine does to inconvenient people. I have seen it up close. I document it without illusion. They can threaten my finances. The SLAPP lawsuit is a real instrument and I expect it to be deployed. Powerful firms and the networks surrounding them have resources to litigate indefinitely. I do not. I hold no illusions about that asymmetry. I am as prepared as a single person working without institutional backing can be. And I will say this plainly: every document produced in discovery becomes part of the public record. I do not experience the threat of litigation as a reason to stop. I experience it as an invitation to put more of the truth on the record under oath. What they cannot touch is the soul. GoAmour — the principles of love that underpin everything I have built from this work — is not a slogan. It is the answer to the question of why a person keeps going after everything that was described in the preceding chapters has happened to them. Love, practiced with knowledge, responsibility, and patience, is the only force that has ever actually bent the arc of history. They can have my money. They already took most of it. They can have my comfort. I gave that up willingly. But the decision I have made — to document, to name, to publish, to stand — is made from the part of a human being that no court order, no bank transfer, no rendition flight, and no settlement agreement has ever reached. I publish this book with full knowledge of the consequences. I bear them gladly. I have more stories to share. And it does not end here.

THE CASE AGAINST BIG TECH

THE PEOPLE VS. BIG TECH

"They told us it was progress. They never told us the price — or who would be made to pay it."

The companies that promised to connect the world have quietly become its most powerful war profiteers. In **No Ethics in Big Tech**, Vahid Razavi delivers an insider's account of an industry that traded conscience for contracts — building the surveillance, data, and AI infrastructure now implicated in occupation, warfare, and mass atrocity.

Part memoir, part indictment, part rallying cry, this book names names, follows the money, and lays out the moral and legal case for accountability. It is the story of one technologist's break with Silicon Valley — and of a growing global movement demanding that those responsible be brought to justice.

The machine was built by people. It can be stopped by them, too.

Campaigners with Forever Peace Now call on governments to prosecute Big Tech companies and their executives for complicity in genocide.

ABOUT THE AUTHOR

Vahid Razavi is a former Silicon Valley insider, author, and peace campaigner, and the founder of Ethics in Tech. His work joining technology and human rights is the basis for the documentary film **Forever Peace Now**. With an introduction by author and media critic **Norman Solomon**.

ADD YOUR NAME

SIGN THE PETITION. MAKE THEM PAY.

Demand that governments prosecute Big Tech companies and their executives for complicity in genocide.
ForeverPeaceNow.com

NoEthicsInBigTech.com | **ForeverPeaceNow.com** | **ParentsPlea.com** | **ConflictTour.com**

Based in part on the documentary film *Forever Peace Now*.
Editorial research and structural editing assisted by Claude (Anthropic). The views, arguments, testimony, and conclusions are entirely those of the author.
Published by **Ethics in Tech Press** · NoEthicsInBigTech.com