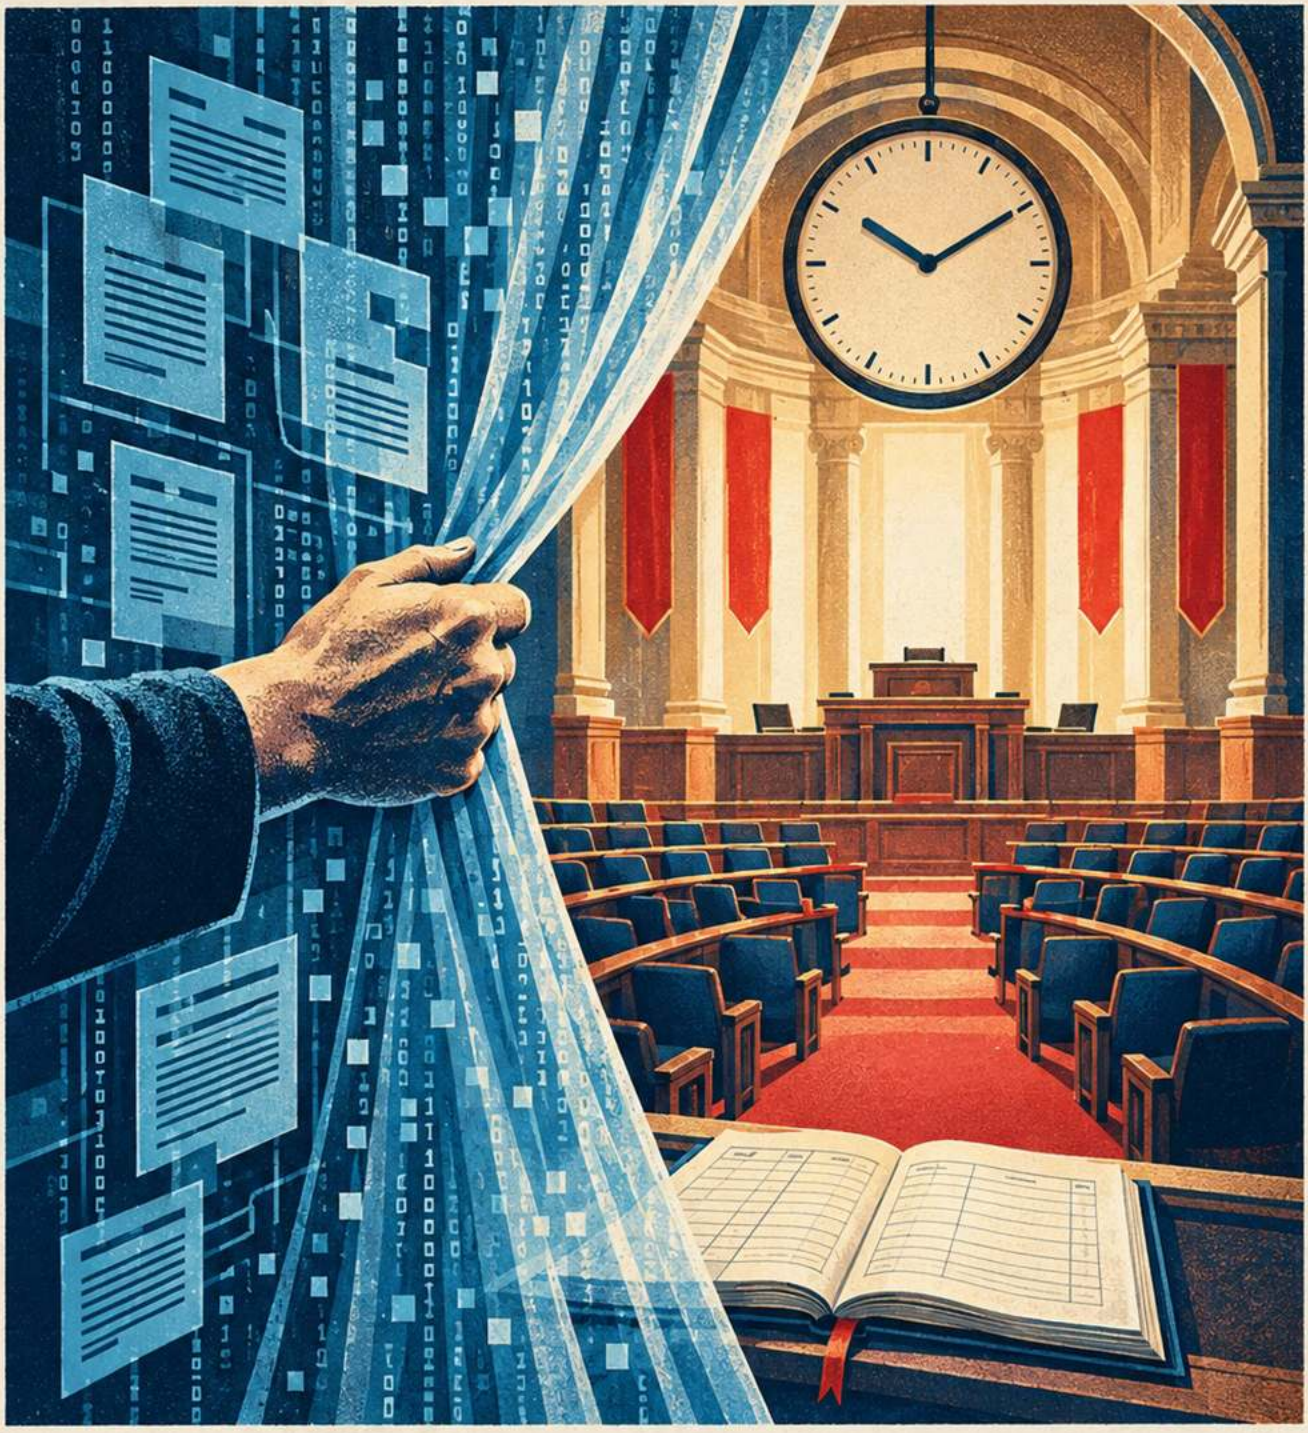
A CITIZEN'S MANIFESTO FOR THE AGE OF AI

PART TWO — WILL



POLITICS AND DEMOCRATIC PRACTICE

Governing at machine speed

A CITIZEN'S MANIFESTO FOR THE AGE OF AI

A Civic Blueprint for Human Agency, Democratic Self-Government, and a Workable Earth

PART TWO — WILL

POLITICS AND DEMOCRATIC PRACTICE

Governing at Machine Speed

DAVID M. DOYLE

The Series Charter

Book Zero—The Unoccupied Center—sets out the charter for this project: technology must enlarge human agency, not replace it. The six books that follow examine the conditions required to keep that commitment real.

Part Two concerns Will: the practical power to consent, refuse, appeal, organize, deliberate, and require an answer from those who exercise consequential power.

The movement of the six-part manifesto is deliberate:

See → Act → Protect → Share → Draw the Line → Build Together.

The Three Tests

Witness Test. Who sees this decision from outside the authority making it, and can their testimony survive?

Burden Test. Who bears the cost, and do they have meaningful voice, remedy, or refusal?

Return Test. How does this authority end, reverse, repair harm, or return value to those it affected?

How to Read Part Two

Perception makes public truth possible. Will asks whether a person can do anything with that truth. A citizen may discover why a benefit was denied and still find no human being authorized to reconsider it.

Will is the capacity to act, consent, refuse, appeal, organize, deliberate, vote, speak, and require an answer from those who exercise power.

The standard is simple: no one should lose liberty, livelihood, housing, health, education, benefits, or standing because "the system determined" and no person can identify the rule, disclose the evidence, hear the appeal, or accept responsibility for being wrong.

PART TWO — WILL

The capacity to act, decide, and defend

Its corruption: force that justifies itself, and emergency that never ends.

Contents

Chapter Four — Politics and Democratic Practice: Governing at Machine Speed

Chapter Five — The Emergency: Urgency as a Technique of Government

Chapter Six — Force and Autonomous Harm: Policing, Borders, Weapons, and the Accountability Gap

Part Opening

Two things happened on the fourteenth of September, twenty-five years apart.

In 1976, Congress approved the National Emergencies Act.

Its purpose was to end a condition that had embarrassed the republic: a country operating under emergency declarations so old that no one could produce a complete list of them, conferring powers no one could fully enumerate, on a basis nobody had voted on in living memory.

The Act was a housekeeping statute with a constitutional ambition.

It said, in effect: emergencies may be declared, but they must be declared by someone, in writing, and they must end.163 In 2001, the President declared a national emergency by reason of the terrorist attacks of September 11 and the continuing and immediate threat of further attacks.164 That emergency has not ended.

It was continued for another year by a notice signed in late August 2025, published on the third of September, stating that the emergency must continue in effect beyond the fourteenth of September, 2025.165 So this month the statute written to stop permanent emergencies turns fifty, and the emergency it did not stop turns twenty-five.

I do not think that is hypocrisy.

I think it is a design problem, and design problems can be fixed by people who are willing to look at the drawing.

Let me be clear about what I am not arguing, because this is the part of the book most easily mistaken for anti-government sentiment, and it is not that.

Will is not the enemy.

Someone must be able to decide.

A levee is failing, a pathogen is doubling every four days, a substation is on fire — in each case the deliberative machinery I spent Part One defending is too slow, and a person has to act on incomplete information and be answerable for it later.

I have made decisions on a Friday afternoon that a committee would have improved and a committee would not have made in time.

A state that cannot act is not a free state; it is a failed one, and its citizens are not sovereign, merely unprotected.

Part One asked whether you can see what is being done to you.

Part Two asks whether anyone can act on what you see — and whether the people who act can be reached.

The corruption is not decisiveness.

The corruption is a specific slide, and it happens in a single sentence: from a person who must decide quickly to a person who therefore has the right to decide forever.

The word "therefore" is where a republic is lost.

Everything in this part concerns that word.

The mechanism is worth stating plainly, because it is almost boring, and its boringness is the whole trick.

The statute's default is termination.

Under section 202(d), a declared national emergency terminates on its own anniversary unless, within the ninety days before that anniversary, the President publishes a notice in the Federal Register and transmits it to Congress saying the emergency will continue.166 Read that as a citizen rather than as a lawyer and it is genuinely reassuring: emergencies die of natural causes, annually, unless someone signs a piece of paper.

Then look at how many pieces of paper get signed.

In the twelve months ending the day I write this, the Federal Register published forty-three separate notices continuing a national emergency in effect for one more year.167 Not one required a debate, a hearing, a vote, or a finding reviewed by anyone outside the executive branch.

Each is a page or two of prose, filed at two in the afternoon, and each renews a grant of statutory authority for a year.

Congress's side of the ledger is thinner.

The Act directs each House to meet, within six months of a declaration and every six months thereafter while it continues, to consider a vote on whether to terminate it.168 The only instance I can document of Congress actually terminating a national emergency by joint resolution is Public Law 118–3, approved the tenth of April, 2023, which terminated the COVID-19 emergency declared on the thirteenth of March, 2020 — three years and a month after the declaration, and well after the emergency it addressed had ceased to be one.169 The most serious earlier attempt, over the 2019 southern-border declaration, passed the House 245–182 and the Senate 59–41, was vetoed, and failed on override, 248–181.170 That declaration was eventually ended the ordinary way: a later President signed a proclamation revoking it.171 I cannot prove from a primary source that Congress has never once held the six-month meeting the statute requires, and I will not assert it.

What the record does support is narrower and sufficient: the executive renewal is routine, automatic in practice, and unreviewed; and the legislative termination is so rare that I can document exactly one, arriving three years late.

There is a second ending, more common than termination and much harder to see, and France is the clearest recent case.

The French state of emergency declared after the attacks of the thirteenth of November, 2015 was renewed several times and finally expired on the first of November, 2017.

It is often described as having been lifted.

What actually happened is that two days before it expired, a statute was promulgated — Law No. 2017-1510 of the thirtieth of October, 2017 — which made a set of its central measures part of ordinary domestic security law: prefect-declared protection perimeters with searches at the boundary, administrative closure of places of worship, individual monitoring measures in three-month increments with reporting and geographic restrictions, and judicially authorized administrative searches of premises a terrorism suspect is believed to frequent.172 That is the pattern I would ask you to learn to recognize.

An emergency power is not usually repealed and not usually renewed forever.

It is naturalized.

The exceptional measure is exercised, found administratively convenient, and then written into the permanent law at the moment the exception lapses — at which point the sunset provision that everyone pointed to during the debate has done its job perfectly and protected nothing at all.

Why does any of this belong in a book about artificial intelligence?

Two reasons, and the second is the one that made me write this part.

The first is that speed has become the argument.

The case for suspending process is no longer only that the threat is grave; it is that the adversary is fast, the technology is moving, the window is closing, and consultation is a luxury of a slower age.

That claim is not always false.

It is, however, unfalsifiable as usually stated, and it is made by the party that benefits from being unexamined.

The second is that the emergency instrument is now being pointed at the venue itself.

On the twentieth of January, 2025, an executive order declared a national emergency concerning insufficient domestic energy production, transportation, refining, and generation; it was continued for a further year on the twelfth of January, 2026.

The continuation notice attributes the emergency in part to "many harmful State and local policies, especially in our Nation's Northeast and West Coast."173 Read that alongside the preemption campaign in the previous chapter and the shape becomes visible.

The declared threat is not only foreign, and not only physical.

It includes the decisions of state and local governments — which is to say, the bodies where a citizen still has standing, a hearing date, and a commissioner who has to run for re-election.

When the exercise of local authority is itself named as a component of the emergency, the emergency power is no longer a tool for acting fast in a crisis.

It is a tool for acting past the only forums most people can actually reach.

So the chapters that follow ask four questions of every fast decision, and I would ask you to carry them into any room where urgency is being offered as a reason: Who decided.

At what speed.

Under what claim of necessity.

And where do I go to argue with them.

And here is the standard I would hold an emergency power to, stated as a single test: an emergency authority is legitimate to the extent that it expires by default rather than by permission, is exercised by a named person who can be reached, and is reviewable while it operates rather than only after it stops.

Each of the three is falsifiable, which is the point of stating it that way.

An authority that renews itself on a signature, is exercised by a system rather than an officer, and can be challenged only years later in a damages action fails all three — and that description now fits a great deal of what is done to people in the name of moving quickly.

The emergency is the hinge of this book.

Not because the threats are imaginary — some are entirely real, and the chapters that follow take them seriously — but because urgency is the one argument that can suspend every protection in the other seventeen chapters, and the only one that gets stronger the less time you are given to evaluate it. — D.M.D.

Chapter Four: Politics and Democratic Practice

Governing at Machine Speed

Governing at Machine Speed

Who decided this, and where do I go to argue with them?

1.

At eight minutes past four in the morning on the first of July, 2025, the United States Senate voted ninety-nine to one to strike from a budget reconciliation bill a provision that would have barred the states from enforcing their own artificial intelligence laws. The amendment's stated purpose, in the official record, was "to strike the section relating to support for artificial intelligence." The lone Nay was Senator Thom Tillis of North Carolina.174 Ninety-nine to one. In a chamber that cannot agree on the naming of a post office. I remember reading the tally and feeling something I had not felt about this subject in a while, which was encouragement. Five months later, on the eleventh of December, 2025, the President signed Executive Order 14365, published at 90 Federal Register 58499. It directs the Attorney General to establish an AI Litigation Task Force within thirty days to challenge state AI laws. It directs the Secretary of Commerce, within ninety days, to evaluate which state laws are "onerous." It directs that states with such laws be made ineligible for the non-deployment portion of federal broadband funds, and authorizes the conditioning of other discretionary grants. It directs the Federal Communications Commission to open a proceeding toward a federal reporting and disclosure standard, and the Federal Trade Commission to issue a policy statement. It carves out child safety, data-center infrastructure, and — this detail is worth holding onto — state procurement and state government use of AI.175 I am not going to tell you those two events are the same thing. They are not. One is a statute that did not pass; the other is an executive instrument whose legal durability is genuinely contested and will be litigated for years. But set them beside each other and you can see the shape of the problem. A near-unanimous legislative judgment was rendered in July. In December an executive order pursued a substantially similar end by other means. And the citizen who wants to know where to go and argue about it discovers that the venue has moved, that the timetable is now measured in ninety-day deliverables, and that the thing being decided is not a rule about machines but a rule about which of their governments they are permitted to petition.

What I understood only later is that a briefing is never only a briefing. Every explanation carries a frame: what counts as a risk, what counts as normal industry practice, which comparisons are apt, what a reasonable timeline looks like. I chose all of that. Nobody in the room had an independent means of checking it, because there was no independent means available to them. The state's understanding of the industry, in those rooms, was the understanding presented. I did not think of that as capture. I thought of it as being helpful, and most of the time I was. But the asymmetry was total, I benefited from it, and I never once suggested that the committee ought to hire someone of its own to tell it whether I was right. So, when I write in this chapter that the state's knowledge of a technology should not be supplied by the party being regulated, understand that I am describing a room I was in, on the side that had the advantage.

II.

Assemble the pressures. Each is defensible standing alone. Technology moves faster than legislatures, so agencies were given rulemaking authority. Agencies are slow, so executive orders became the instrument of choice. Executive orders are contested, so funding conditions became the enforcement mechanism. Fifty different state rules impose real compliance costs, so uniformity became a legitimate business ask. Local opposition can be parochial, so state law preempted municipalities. Benefits fraud is real, so detection was automated. Automated detection is cheap, so caseloads were raised. Higher caseloads make human review expensive, so review became exceptional. Backlogs are cruel to claimants, so agencies bought tools promising speed. Speed metrics are measurable and accuracy metrics are not, so speed became the reported number. Legislatures lack technical staff, so they consulted industry. Campaigns need to reach voters, so persuasion became individually targeted. Individually targeted persuasion has no public version, so there is nothing to rebut. And every one of those steps was taken by someone reasonable, under real pressure, for a stated reason. The result is not a coup. It is a compression. The interval between a decision being contemplated and a decision being irreversible has shortened to something less than the interval required to notice, organize, and object — and the second interval is the one democracy actually runs on.

III.

Before blaming the machine, be honest about the ground it landed on. Representative government was already struggling to legislate at the pace of anything. Consider the velocity gap in its plainest form. In the 2025 legislative sessions, artificial intelligence legislation was introduced in all fifty states, Puerto Rico, the Virgin Islands, and the District of Columbia; thirty-eight states adopted or enacted roughly one hundred measures.176 Independent tallies of the National Conference of State Legislatures' bill database put the number of bills introduced well above a thousand, though the Conference itself publishes no aggregate and the counts vary by who is counting.177 In the same Congress, the federal statute book acquired one narrow AI-adjacent criminal law: the TAKE IT DOWN Act, Public Law 119-12, enacted the nineteenth of May, 2025, addressing non-consensual intimate imagery.178 The leading vehicle for a general federal framework, the SANDBOX Act, was introduced on the tenth of September, 2025, read twice, referred to committee, and has had no recorded action since.179 Meanwhile the institution whose job is to tell Congress how things work is getting smaller. The Government Accountability Office's own budget testimony to the House Appropriations Committee in March 2026 reports 3,556 full-time equivalent staff in fiscal 2025, 3,353 enacted for fiscal 2026, and 3,210 requested for fiscal 2027 — a 4.2 percent reduction from the prior year and a 10.2 percent reduction since the end of fiscal 2024, requested alongside a 5.9 percent increase in dollars.180 The Office notes that it issued its first science and technology trend paper in fiscal 2025 and a technology assessment on generative AI, and that it participates in a Legislative Branch AI Working Group. It does not publish a staffing figure for its science and technology arm, and I could not obtain one.181 Against that: in 2025, lobbyists reported representing 774 organizations on artificial intelligence issues and filed more than 3,500 reports mentioning the subject — increases of 423 and 505 percent — according to OpenSecrets data reported in the trade and advocacy press, with total federal lobbying reaching a record $5.08 billion.182 I have not been able to fetch OpenSecrets' own tabulation and so I attribute the figures to the outlets reporting them. Even discounted, the juxtaposition holds: 774 organizations working the question on one side of the table, and a congressional support agency asking for 143 fewer people on the other. Now the complication, because the honest version of this chapter is not a lament about a democracy that no longer functions. It functioned. Repeatedly. Visibly. Within the last eighteen months.

The Senate struck the state-law moratorium ninety-nine to one.183 Texas enacted an AI governance act by votes of 146 to 3 in the House and 31 to 0 in the Senate — and included in it an affirmative obligation on government agencies to tell a person when they are dealing with an AI system, "even if it would be obvious to a reasonable consumer."184 A federal judge in the District of Massachusetts held on the twenty-fifth of June, 2026, that Sections 2 and 3 of Executive Order 14399, which would have built a federal apparatus over state voter rolls, were ultra vires and unconstitutional violations of the separation of powers, granted the plaintiff states summary judgment, and enjoined named cabinet officers by name.185 A federal judge in the Eastern District of California struck down California's political deepfake statute on First Amendment grounds, holding that a law punishing synthetic election speech according to whether it harmed or benefited a candidate was "the essence of viewpoint discrimination" and that "a mandatory disclaimer for parody or satire would kill the joke."186 In Lysander, New York, three hundred and fifty residents turned up to a town meeting and the board voted unanimously for a six-month data-center moratorium.187 Roughly fourteen states saw data-center ban or moratorium bills in the 2026 sessions; trackers disagree substantially on how many died, but Maine's was vetoed by Governor Mills on the twenty-fourth of April, 2026, and New York's passed the legislature on the fourth of June.188 So the institutions are not inert. Courts moved. A state legislature imposed a disclosure duty on itself. A town in Onondaga County stopped a data center by showing up. What I take from that is narrower than a reassurance, and more useful. Look at what the wins have in common. Almost every one is either a veto — a court voiding, a legislature striking, a town pausing — or a disclosure duty with no adjudicative machinery behind it. Very few are affirmative constructions: a register that is maintained, an appeal that is heard, a capacity that is funded, a proceeding whose record the public can read. Democracy is currently demonstrating that it can still stop things. It has not yet demonstrated that it can administer this, which is the harder and more necessary competence, and the one that decays if unused.

IV.

Here is what is actually happening, named. I set out the full record — the orders, the dockets, the subpoenaed documents, the price lists — in the casebook at the back of this book, because it runs to some thousands of words and a reader who wants to check me should be able to do it without reading a chapter twice. What follows is the anatomy. The compression of the interval. The relevant metric is not how fast technology develops. It is the ratio between the time required to deploy a system and the time required to obtain consent for it. Louisiana supplies the cleanest measurement I have found: on the seventeenth of December, 2025, the Public Service Commission voted four to one to waive the requirements of its own General Order for capacity additions meeting five conditions, and to target a Commission vote within eight months of application. Four of the five conditions are things a large applicant satisfies by signing documents. One is a letter from the state's economic development office. None relates to what the public knows, or when. The amendment's informal name, in the Commission's own minutes, is the "Lightning Amendment."189 Preemption from above, and from below. Executive Order 14365 is not a rule about artificial intelligence. It is a rule about federalism executed through the appropriations power: states whose AI laws the Commerce Department deems onerous become ineligible for certain broadband funds. It also exempts state procurement and state government use of these systems from its preemptive reach — which is to say the federal government moved to preempt states' regulation of what companies may do to citizens while expressly leaving alone states' regulation of what states themselves may do.190 The same logic runs downhill. A state that has decided to attract data-center investment has every incentive to remove the parish and the town from the decision; Louisiana's data-center sales-tax

exemption is a state-level grant of value that no local body voted for and none can withdraw.191 Local authority matters here because data centers consume water, land, and electricity — three things that are locally sited, locally metered, and locally suffered — and the venue for that grievance has been moved to a body that meets in the capital.

The official who is no longer there. This is the mechanism that answers the chapter's question with the word nowhere. Michigan's automated unemployment system determined fraud without human involvement; the state's Auditor General reviewed twenty-two thousand of those determinations and found that ninety-three percent did not involve fraud.192 The class action settled for twenty million dollars in January 2024 — individual checks of roughly sixteen hundred dollars for having been accused of a felony by a computer.193 The casebook sets out four more: Australia's income-averaging scheme, which a Royal Commission examined and which now carries a settlement sum of $548.5 million; Idaho's undisclosed Medicaid formula, in litigation for thirteen years and still active; Arkansas's attendant-care algorithm, which cut care hours by an average of forty-three percent; and the multi-state benefits-eligibility contractor whose systems are now the subject of a federal ruling that a state violated the Constitution. The common structure is identical in every one. A system produced an adverse determination. The person could not learn the basis. The agency could not or would not explain it. The vendor was not a party the citizen could petition. And the remedy, when it came, came years later, through a class action, in money. Denial by inference, in the largest program we have. The best-documented case in the record exists because a Senate subcommittee subpoenaed the documents, and the number I cannot stop thinking about is a written record of a choice. One insurer's inpatient predictive model, configured to maximize savings, produced roughly $3.6 million in net savings at a 1.5 percent mistake rate; configured to maximize automatic approvals, it produced a net loss of more than $400,000 at a 2.1 percent mistake rate. The difference in error was six-tenths of a percentage point. The company chose.194 Six-tenths of a point is invisible in any dashboard and decisive in the life of every person on the wrong side of it.195 The casebook carries the denial-rate series, the automation findings, and the two cautions that matter: the Subcommittee expressly did not find uniform algorithmic denial across the industry, and the statistics everyone quotes from the private litigation are complaint allegations, denied by the defendants.

Registers that exist, and are thinner than their reputation. The United Kingdom's mandatory algorithmic transparency register held one hundred forty-three records when I checked it on the twenty-ninth of August, 2026 — real institutional progress, and I do not want to be sour about it.196 Against that, the United States federal inventory reports 3,611 self-reported use cases, of which 445 are designated high-impact.197 The second number is larger and less useful, because a self-reported inventory entry is a description an agency wrote about itself and a register record is a document the public can hold the agency to. The European Union's high-risk database — a legal commitment under Article 71 — has had its operative date pushed to December 2027. The casebook gives the counts and the deferral votes. A right to an explanation that keeps receding as you approach it. In European law the right is stronger than most people think, and one sentence from the Court of Justice belongs on the wall of every agency in America: "the mere communication of an algorithm does not constitute a sufficiently concise and intelligible explanation," and where a controller claims trade secrecy it must give the disputed information to the supervisory authority or the court to balance.198 That disposes of both standard evasions at once. Then watch it recede. Colorado's comprehensive act never took effect, was repealed and reenacted in May 2026, and now gives a consumer a plain-language description of the technology's role within thirty days of an adverse outcome, enforceable only by the Attorney General after sixty days' notice and an opportunity to cure.199 Texas requires agencies to announce that you are dealing with a machine. In no

American jurisdiction I can find does a person denied a benefit by an automated system have a self-executing right to an explanation, human review, and a timely appeal, enforceable by that person, in a forum that can order the determination reversed.

Persuasion below the threshold of rebuttal. Two days before the 2024 New Hampshire primary, 9,581 telephone calls carried an AI-generated imitation of the President's voice, made with commercial software by a third party paid one hundred fifty dollars. The Federal Communications Commission imposed a six-million-dollar forfeiture.200 In June 2025 a jury acquitted the operative on all twenty-two counts.201 I do not know what to do with that acquittal, and saying so is more honest than pretending. Around that case the law has largely declined to act, and the casebook records the declining. But every instrument in that record regulates the artifact: the cloned voice, the altered video, the undisclosed synthetic image. None regulates the addressing. The mechanism that most threatens public reason is not a fake video of a candidate. It is ten thousand true-enough messages, each tailored to one person's inferred anxieties, delivered without a public version, so that there is no shared claim for anyone to answer. A lie in a broadcast can be refuted on the same channel. A different appeal to each of ten thousand voters cannot be refuted at all, because there is no it to refute. The roll. In March 2026 an executive order sought to build a federal apparatus for verifying citizenship on state voter rolls, anchored on a Homeland Security verification service that its own operating agency says "does not determine an individual's eligibility to register to vote, eligibility to stay on a voting roll, or eligibility to vote." Twenty-three states and the District of Columbia sued, and in June 2026 a federal judge held the operative sections legally void as ultra vires and as violations of the separation of powers, enjoining named cabinet secretaries from creating "a new federal program to superintend and control Plaintiff States' maintenance of voter rolls" and giving them seven days to notify their own employees in writing, under penalty of contempt.202 I include this because it is the strongest verified thing in this chapter and it cuts in the direction of hope. The system worked. A court read an order, found it lawless, and told cabinet secretaries by name to send a memo saying so.

The state renting the machine. In August 2025 the General Services Administration began supplying federal agencies with commercial AI tools at no cost, and then came the price list: enterprise offerings at one dollar, at forty-seven cents, at twenty-five cents, most of them expiring within a single fiscal cycle.203 What bothers me is not the discount. It is that the state's operating dependence is being established at prices that cannot be a market and on terms that expire in weeks, by which time the agencies will have built workflows, retrained staff, and retired the alternatives. That is not procurement. It is the customary structure of a dependency, and the government has entered it as the dependent party. The same asymmetry appears in who is doing the deciding: the Government Accountability Office could not determine whether the people running the federal automation push had completed ethics training, because the Executive Office of the President did not answer its requests.204 Procurement transparency that points inward. The federal government does have a serious AI acquisition policy, and it deserves credit. It requires solicitations to state whether a use is high-impact, requires contract terms on data rights and vendor lock-in, requires testing against validation data "that should not be accessible to the vendor," and permanently prohibits training commercial models on non-public agency inputs absent explicit consent.205 It is well drafted. Two things about it are visible in the text. Every transparency obligation it creates runs from the vendor to the agency; not one runs from the agency to the public. And the repository of acquisition tools it directs is expressly available only to executive branch agencies. The government built itself a right to know what it is buying and did not extend that right to the people it is buying it to use on.

V.

The concerns, stated plainly enough to read aloud at a hearing.

VI.

Take it together and the cumulative issue is not that government is using computers. Government has used computers since the punch card, and much of what it does with them is unremarkable and good. Nobody should want a human being retyping a change of address. The cumulative issue is a change in the location of the decision, and with it the disappearance of the addressable party. Democratic accountability has never depended on citizens understanding how the state works. It has depended on there being someone to blame. A name, an office, a term of years, a hearing, a docket, a vote. The whole apparatus is machinery for finding the person responsible and imposing a cost on them.

What is being dismantled is not comprehension. It is addressability. The determination was made by a system. The system was configured by a vendor. The vendor was selected under a procurement whose terms are not public. The procurement was authorized by a policy that runs vendor-to-agency. The agency's authority came from an executive instrument that preempted the legislature. The legislature lacked the staff to evaluate it and relied on the industry to explain it. And at the end of that chain stands a person holding a denial letter that says the system determined they are ineligible, with a telephone number that reaches a queue. Every link in that chain is individually lawful. Together they constitute a government that acts upon people who cannot locate anyone to argue with. And that is a different failure from tyranny, and in some ways harder, because tyranny has an address.

VII.

Objections I take seriously. "Government is too slow and people suffer from the slowness." This is the strongest objection in the chapter and I will not wave it off. Benefits backlogs are not an abstraction; they are evictions and untreated illness and repossessed cars. Nevada's unemployment appeals system faced a backlog of more than forty thousand cases, and the state says its AI-assisted drafting reduced turnaround substantially while referees retained full authority.206 If that is true it is a real good, delivered to people with no other advocate. My honest position: I want that speed and I want it audited, and no independent evaluation of those claims has been published, which is the whole of my complaint. "A patchwork of fifty state laws is a genuine burden, and uniformity is a legitimate goal." True, and I say so without hedging. A company operating in fifty jurisdictions under fifty definitions of an automated decision faces costs that are real and largely deadweight. The correct answer is a federal statute passed by Congress, which is exactly the instrument that has not been produced. What is objectionable is not uniformity. It is uniformity obtained by conditioning grants rather than by legislating, because that method sets a national rule without a national vote and leaves the citizen no forum at all — not the state one that was preempted, and not the federal one that never convened.

"You are demanding a level of process that will make government unaffordable." Partly. A public register, a written explanation within a deadline, human review, and a real appeal all cost money, and I am asking for an expenditure. But compare it to what the current arrangement costs: a $20 million Michigan settlement, a $548.5 million Australian settlement, thirteen years of Idaho litigation, a $125 million overcharging settlement in another chapter of this book. Due process is not the expensive option. It is the option we are currently paying

for retroactively, at litigation rates, with interest, to lawyers rather than to claimants. “Publishing a register or an explanation lets people game the system.” Sometimes, and for genuine fraud-detection systems this concern is legitimate and I would accept redaction of specific detection thresholds on a stated and reviewable ground. But notice how far that argument has been stretched. Idaho withheld the formula for calculating disability budgets — not a fraud tripwire, an allocation rule. There is no gaming to be done with an allocation rule; there is only the question of whether the arithmetic is correct, which is precisely what secrecy prevents anyone from checking. The European court’s answer is the right one: disclose to the tribunal even where you will not disclose to the world, and let a judge balance it. “Requiring disclosure of political advertising drives platforms out and reduces speech.” This is not hypothetical and I have to concede it. Europe passed a transparency regulation and the largest platform stopped carrying political ads in the Union entirely.207 Whether that is a loss depends on what you think paid political micro-targeting contributes, and I hold a minority view: I think its disappearance is closer to a gain than a loss, because organic political speech remains lawful and only its purchased amplification stopped. But a rule whose predictable effect is to remove a category of speech rather than to illuminate it has not achieved what it set out to achieve, and a smaller campaign without a platform’s paid reach is disadvantaged against an incumbent with earned media. I do not have a clean answer. I would rather have the argument in the open than pretend the tradeoff is not there. “Courts handled it — the moratorium was struck, the voter-roll order was voided, the deepfake statute fell.” They did, and I put every one of those in Section III precisely because they are real. But look at the reliance profile. Each required a plaintiff with standing, a well-resourced litigant or a coalition of state attorneys general, a judge willing to act within weeks, and an appellate posture that held. The Massachusetts order required twenty-three states and the District of Columbia to bring it. A person denied home care by an algorithm does not have twenty-three states. Litigation is the remedy of the well-organized, and a system whose only functioning accountability mechanism is litigation has quietly restricted accountability to those who can afford to sue.

VIII.

Three questions I would ask about any system a government uses to make or materially inform a decision about a person. They are not technical. Anyone can ask them at a hearing, and the shape of the answer tells you what you are dealing with. Who can witness it? A public register is the minimum condition of being seen. Not the source code, not the weights, not the training data — the existence of the system, its purpose, the decisions it touches, who operates it, who built it, and how to contest its output. A hundred and forty-three published records against thirty-six hundred self-reported use cases is not a transparency gap; it is a difference in kind between a register and an inventory. An inventory is what an institution says about itself. A register is a document the public can hold it to. A system nobody outside the agency knows exists cannot be argued with, cannot be evaluated, cannot be litigated, and cannot be voted about, which makes secrecy about the mere existence of a decision system the most efficient available method of removing a subject from democratic reach. You do not need to suppress the debate. You need only decline to disclose that there is something to debate. Who carries the burden if it is wrong? Right now, entirely the person. When Michigan’s system found fraud where there was none in ninety-three percent of a sampled twenty-two thousand determinations, the state’s cost was a settlement a decade later and the claimant’s cost was the accusation, the penalty, the garnishment, and the years. When a model was configured to maximize savings rather than approvals at a six-tenths-of-a-point cost in accuracy, the institution booked the savings and the patients absorbed the errors. When an agency cannot say who was trained to run its automation, the agency suffers nothing and the governed carry the entire risk of its

incompetence. Look for the asymmetry every time: what happens to the agency when the system is wrong, and what happens to the person. If the answer to the first is nothing, the system's error rate will not improve, however sincere everyone involved is, because nothing in the arrangement is asking it to.

Can it be undone? Appeal rights are the only mechanism by which authority returns to the people it was exercised over. Everything else — notice, disclosure, explanation, a register — is preparation for this. An explanation you cannot act on is a courtesy. A register that leads to no forum is a catalogue. The question is whether there exists a named human being with the authority to reverse the determination, a deadline by which they must decide, an obligation to give reasons in writing, and a further body that can compel them. Measure the American landscape against that: a thirty-day plain-language description with no private right of action, a chatbot disclosure enforceable only by an attorney general, a European right whose companion obligations were deferred to 2027. Almost everywhere, the reversal mechanism is a lawsuit, which means the answer to can it be undone is yes, if you can afford five years. A system that fails all three — unregistered, costless to its operator when wrong, and unappealable — is not administration. It is rule, and the fact that it is conducted by software procured at a dollar rather than decreed by a person in a palace should be no comfort whatsoever.

IX.

What I would insist on. Conditions, with a name attached and a consequence for failure.

1. A public register of every automated system a government uses to make or materially inform a decision about a person. One register per jurisdiction — federal, state, county, municipal — searchable, free, with a plain-language entry for each system: what it does, which decisions it touches, who operates it, who supplied it, when it entered service, what it was tested against, what its measured error rate is and on what population, and where to contest its output. Registration precedes use. An unregistered system's determinations are void. The United Kingdom has shown this is achievable at national scale; its register held a hundred and forty-three records when I checked it, and the standard covers tools with significant influence on decisions with public effect or that interact directly with the public.208 The difference from a self-descriptive repository is not technical. It is whether the entry creates an obligation. Who: the head of each agency, personally named in the entry; the state chief information officer; the clerk of each municipal body. If not: no determination made by an unregistered system has legal effect, and no public funds may be disbursed under the contract that supplied it.

2. A statutory right to an explanation, human review, and a timely appeal — enforceable by the person, not only by an attorney general. On any adverse determination materially informed by an automated system: written notice that a system was used; a plain-language explanation of the procedure and principles actually applied to this person's data, sufficient to show which facts mattered and how a different fact would have changed the result; a named human reviewer with authority to reverse; a decision deadline; written reasons; and an appeal to a body that can compel. The European court has already drawn the line correctly in both directions — publishing the algorithm does not satisfy the duty, and trade secrecy does not defeat it, because the material goes to the tribunal even when it does not go to the world.209 Colorado's thirty-day plain-language description is the right shape and the wrong strength: a cure period and no private right of action mean the duty is real for a regulator and theoretical for a claimant.210 The person affected must be able to enforce it, or it is not a right. Who: the deciding agency's named appeals officer; the vendor as a compellable third party; the reviewing court. If not: the adverse determination is vacated by operation of law and the prior status restored pending a compliant redetermination, with the cost of the delay borne by the agency.

3. Independently funded legislative and municipal technical capacity, on multi-year appropriations. Every legislature that regulates these systems needs its own people to evaluate them — not consultants seconded from industry, staff. At the federal level that means restoring and expanding the technology assessment function rather than shrinking the agency that houses it; the Government Accountability Office is requesting 3,210 full-time equivalents for fiscal 2027 against 3,556 in fiscal 2025, a 10.2 percent reduction since fiscal 2024, and publishes no separate figure for its science and technology arm.211 At the state level it means a standing office with subpoena support. At the municipal level it means a shared regional resource, because a parish with two staff cannot hire a machine learning engineer and should not have to. And the appropriation must run multiple years, because a capacity that can be zeroed in one budget cycle is a capacity the regulated party can wait out. Set beside 774 organizations lobbying the subject in a single year, this is the cheapest structural reform in this chapter.212 Who: the appropriating body; the presiding officer of each chamber; the parish or county administrator. If not: no bill regulating or authorizing an automated system advances to a floor vote without a written technical assessment prepared by staff who are not employed, funded, or seconded by any party with an interest in the outcome.

4. Disclosure of automated political persuasion: funder, targeting, and synthetic content, on the face of the message. Every paid political communication carries, in the message itself and in a public repository, who paid, how much, which audience criteria selected the recipient, and whether any element was generated or materially altered by a machine. The repository is public, free, machine-readable, retained through the next full election cycle, and available at the precinct level. The First Amendment constrains the American version severely, and a state statute punishing synthetic election speech according to whether it helped or harmed a candidate has already been struck down as viewpoint discrimination and compelled speech — correctly, in my reading, and the parody holding especially.213 So I frame this narrowly: a disclosure duty about funding and targeting, which is conduct, not viewpoint, and which the Federal Election Commission itself has already said is governed by technology-neutral rules it declined to elaborate.214 The synthetic-content element should be a neutral labeling requirement applied without regard to whom the content favors, or it will not survive, and should not. Who: the purchaser of record; the platform's named compliance officer; the state elections board. If not: the communication is an unlawful expenditure, the platform is jointly liable for accepting it, and the repository gap is itself the violation.

5. Protected local authority over siting, water, and land use, which no state incentive may override. Where a facility consumes water from a local aquifer, occupies local land, or draws load that local ratepayers will amortize, the local body retains a genuine, non-advisory decision — with the applicant's figures on the public record before the vote, a hearing at which residents may be heard, and no state economic-development instrument capable of waiving it. This is the provision that Louisiana's Lightning Amendment runs directly against: five conditions an applicant can satisfy with signatures and a letter from the economic development office, and an eight-month target for a Commission vote, with the competing large-load rulemaking deferred.215 The eight months is not the problem. The problem is that the clock was set by the applicant's requirements and nobody set a clock for the public's. Who: the parish, county, or municipal governing body by recorded vote; the state public service commission; the attorney general on petition. If not: the state incentive is unavailable for the project, the siting approval is void, and the proceeding reopens with the withheld figures in the file.

I have reviewed enough governance documents to hear the objection: this is a great deal of machinery for a county with three commissioners and a part-time attorney. So take the two that cannot be given up. Number one, the register — because a system whose existence is not disclosed cannot be voted about, litigated, evaluated, or argued with, and every other provision here is downstream of knowing the thing is there. And number two, the appeal — because a determination that cannot be reversed by the person it was made about is not administration of a law. It is the exercise of power on a subject, and we have an older word for that arrangement than any of the words in this chapter.

X.

I am not asking the state to move slowly out of nostalgia. I have no affection for the fourteen-month permitting review, the docket that dies in committee, the hearing scheduled at two in the afternoon on a Tuesday when everyone affected is at work. Slowness is not a virtue and I have spent a good part of my career being impatient with it, usually correctly. I am asking the state to move at the speed at which consent can actually be obtained. Those are different requests, and the difference is the whole of this chapter. Consent has a minimum duration. It takes time for a proposal to become known, for the affected to find each other, for someone to read the filing, for a question to be formed, for an answer to be demanded, for the answer to be tested, for an objection to be lodged in a place where it must be addressed. You cannot compress that below a certain floor and still have obtained anything. Below the floor, what you get is not fast consent. It is a decision with a consultation stapled to it. Everything I have described here is a method of operating below the floor. The eight-month target that starts when the applicant is ready. The ninety-day deliverable that preempts the legislature that would have been the forum. The determination issued in milliseconds and appealable in years. The tailored message with no public version to answer. The system whose existence is not on any register, procured under terms the public cannot read, from a vendor the citizen cannot petition, priced at one dollar for a term that ends next month. None of it is illegal. Most of it was done by people who would be genuinely offended by this chapter, and some of them would have a point, and one of them used to be me. But here is the thing I keep returning to. Everything else in this book assumes the existence of a functioning venue. A place where a claim can be made, heard, tested, and decided; where the losing side knows it lost and why; and where the decision can be revisited when the facts change. Every protection I will ask for later is, in the end, a request addressed to that venue. The venue is currently being closed from three directions at once: from above by preemption, from below by automation, and from the side by a persuasion apparatus that has no public version and therefore admits no reply. What I want is small and difficult. I want there to continue to be a room. With a name on the door, a person in the chair, a file on the table that I am entitled to read, an hour at which I may speak, and an obligation on someone to answer me in writing. Not a portal. Not a queue. Not a comment period whose docket closed before the notice reached me. A room. Who decided this, and where do I go to argue with them. If that question has an answer, the rest of this book is a set of arguments to be had. If it does not, the rest of this book is a wish list addressed to no one. — D.M.D.

Chapter Five: The Emergency

Urgency as a Technique of Government

Urgency as a Technique of Government

Who gets to decide that there is no time to ask me?

I.

On the twenty-third of May, 2025, the Secretary of Energy ordered a coal plant on the eastern shore of Lake Michigan not to close. The J.H. Campbell Generating Complex sits at West Olive, Michigan. Its owner, Consumers Energy, had planned to retire it on the thirty-first of May. The Midcontinent Independent System Operator had reviewed the retirement and approved it. The Department of Energy's own description of the plant notes that it "was originally slated to shut down on May 31, 2025, 15 years before the end of its scheduled design life." Order No. 202-25-3 directed MISO, working with Consumers Energy, to keep the plant "available to operate," to minimize "any potential generation shortfall that could lead to unnecessary power outages," and to "take every step to employ economic dispatch to minimize costs to ratepayers."216 The order ran ninety days, because ninety days is what the statute allows. On the twentieth of August it was extended to the nineteenth of November. On the eighteenth of November, to the seventeenth of February, 2026. On the seventeenth of February, to the eighteenth of May. On the eighteenth of May, to the sixteenth of August. On the fourteenth of August, effective the seventeenth, to the fourteenth of November, 2026.217 Six orders. Eighteen months. As I write this in the last days of August 2026, the plant is running under the sixth. I want to point at the calendar rather than the coal, because the coal is a separate argument and people of good faith land in different places on it. The calendar is not. An emergency, in the ordinary meaning of the word and in the logic of every statute that uses it, is a condition that arrives without warning and does not persist. This one arrives every ninety days, on schedule, two or three days before the previous instrument lapses, and has done so for a year and a half without interruption. Nobody has had to argue, in any forum where a resident of Ottawa County could be heard, that the emergency is continuing. The renewal is the argument. The existence of the previous order is the strongest available evidence for the next one, and at no point in the sequence must somebody stand up and make the case from the beginning to a person with the power to say no.

I should admit something before going further. I have declared emergencies. Not national ones — corporate ones, which are the same instrument at a smaller scale. As a chief executive I have called a "code red", stood up a war room, suspended the normal approval path because something had been promised or something had broken or a competitor had moved. Most of those calls were defensible and a few were probably necessary. But I know what the instrument does, because I used it, and what it does is not primarily to create speed. What it does is suspend the obligation to persuade anyone. That is the whole appeal. Under normal conditions I had to convince a board, a finance function, a legal review, and several people whose objections I would have had to answer. Under an emergency I had to convince no one, and the people who would have objected understood that objecting had become a kind of disloyalty. The speed was a byproduct. The product was the silence. And here is the part I did not see clearly at the time: emergencies in an organization do not end because the condition resolves. They end because somebody with standing insists on ending them, in a meeting, on a date, against resistance. Absent that person and that date, the emergency simply becomes how the place is run. I have

watched a ninety-day freeze become the operating model of a division. I have watched an exception process become the process. It is not that emergency powers are abused by bad people. It is that they are used, sincerely and lawfully, by ordinary people who find that the instrument works, and that no feature of the instrument requires anyone to give it back.

II.

Start with what the statute actually says, because the statute is more permissive than most people assume and less permissive than it is being used. Section 202(c) of the Federal Power Act, codified at 16 U.S.C. §824a(c), lets the Department of Energy "require by order such temporary connections of facilities and such generation, delivery, interchange, or transmission of electric energy as in its judgment will best meet the emergency and serve the public interest." The predicate is a determination "that an emergency exists by reason of a sudden increase in the demand for electric energy, or a shortage of electric energy or of facilities for the generation or transmission of electric energy, or of fuel or water for generating facilities, or other causes."218 Read the last three words again. Every carefully enumerated trigger in that sentence is followed by a clause that makes enumeration optional.

Congress did impose a limit. Under amendments made in 2015, an order that may result in a conflict with an environmental requirement issues for a period of ninety days, is renewable in additional ninety-day periods, must be in effect "only during hours necessary to meet the emergency and serve the public interest," and must "to the maximum extent practicable" minimize "any adverse environmental impacts."218 So there is a sunset. It is ninety days. It is also renewable, the statute does not say how many times, does not require that a renewal be justified on a fresh record, and does not create any forum in which a renewal can be contested by the people who pay for it or breathe near it. A sunset with an unlimited renewal clause and no reopened record is not a sunset. It is a subscription. The volume tells you the rest. Counting the orders and lettered extensions that the Department itself lists on its own pages, there were nineteen distinct designations in 2025 and fifty-five in 2026 through the middle of August — seventy-four in twenty months.219220 For comparison, the Congressional Research Service, reviewing the entire history of the authority, identifies twenty-six discrete events for which §202(c) orders were issued between 2000 and February 2026, and twenty-nine uses before 2000, of which twenty-two were during the Second World War.221 A single event, Winter Storm Fern, accounted for twenty orders.221 What is new is the standing order, renewed on a quarterly rhythm, keeping a scheduled retirement from happening. There are at least six of these running concurrently. Eddystone Units 3 and 4 in Pennsylvania, dual-fuel oil and gas, ordered on the thirtieth of May, 2025 and extended five times through the twentieth of November, 2026.222 Centralia Unit 2 in Washington, ordered in December 2025 and extended twice.223 Schahfer Units 17 and 18 and F.B. Culley Unit 2 in Indiana, and Craig Station Unit 1 in Colorado, all ordered in the last ten days of December 2025 and all extended twice since.224 Wagner Unit 4 in Anne Arundel County, Maryland, first ordered to run past its "environmental run-hour limitations" in July 2025 and running under an August 2026 extension as I write.225 And a second new kind, more revealing than the retirement deferrals: the permit-override dispatch order. In late January 2026, during a cold event, the Department directed all units in the PJM footprint to maximum output "notwithstanding air quality or other permit limitations or fuel shortages during the pendency of this emergency."226 In June and July 2026 it issued two more to Duke Energy authorizing maximum output "notwithstanding air quality or other permit limitations arising under federal, state, or local law."227

Notice what those words accomplish. A single administrative instrument, issued without notice, comment, or hearing, sets aside conditions in permits that were themselves issued after notice, comment, and hearings, and does so at the federal, state, and local level simultaneously. Every one of those permit conditions is the residue of an argument somebody won. Some of them were won by people who spent years on it. They are all switched off by a sentence. I want to be fair about the Department's stated grounds, and I have to be careful because I was unable to obtain the underlying orders individually. According to a trade-press log of the orders, the original Campbell order cited the North American Electric Reliability Corporation's 2025 Summer Reliability Assessment warning that MISO faced an "elevated risk of operating reserve shortfalls," along with the retirement of roughly 2,700 megawatts of Michigan coal since 2020; the extensions cited MISO grid alerts on at least forty of the sixty-nine days between the eleventh of June and the eighteenth of August, 2025, an Energy Emergency Alert on the twenty-third of June, and survey projections of a 1.4 gigawatt capacity deficit in 2027 growing to 8.2 gigawatts by 2030.228 Those are not trivial figures and I will not pretend they are. But one fact about these orders is contested and belongs in the open. The same trade-press log, summarizing a Congressional Research Service report, states that for the May 2025 retirement-deferral orders "the relevant grid operators had neither requested DOE action nor identified reliability risks specifically tied to those plant retirements at the time they approved them," and notes that PJM publicly supported the Eddystone order while MISO had approved the Campbell retirement having found no reliability problem.229 I could not verify that characterization against the CRS product itself, and I flag it as secondary. If it is right, it matters enormously: the emergency was identified by the party with the power to declare it and not by the parties responsible for keeping the lights on.

III.

Now the deeper mechanism, which is not the order but the declaration underneath it, and not the declaration but the definition inside the declaration. Executive Order 14156, signed on the twentieth of January, 2025 and published at 90 Federal Register 8433, declares that "the United States' insufficient energy production, transportation, refining, and generation constitutes an unusual and extraordinary threat to our Nation's economy, national security, and foreign policy. In light of these findings, I hereby declare a national emergency."230 Section 8 defines what "energy" means for its purposes: "crude oil, natural gas, lease condensates, natural gas liquids, refined petroleum products, uranium, coal, biofuels, geothermal heat, the kinetic movement of flowing water, and critical minerals."231

Eleven categories. Solar is not among them. Wind is not among them. Storage is not among them. Efficiency is not among them. I am not making a climate argument here — that argument comes later in this book and it can stand on its own. I am making a structural one. The emergency is defined by the same instrument that grants the powers, and therefore the definition is where the whole contest actually happens. An emergency in eleven named categories that exclude the fastest-growing sources of new generation on the American grid is not a neutral technical condition. It is an industrial policy wearing an emergency's clothes, and the reason it can wear them is that nobody outside the drafting room gets a vote on the definition. The order then goes shopping. Section 3 directs agencies to "identify and use all relevant lawful emergency and other authorities available to them to expedite the completion" of infrastructure and energy projects, with subsection (b) singling out "the West Coast of the United States, Northeast of the United States, and Alaska."232 Section 4 directs agencies to use "to the fullest extent possible" the emergency treatment provisions of the Army Corps permitting regime under section 404 of the Clean Water Act, section 10 of the Rivers and Harbors Act, and section 103 of the

Marine Protection, Research, and Sanctuaries Act.233 Section 5 directs agencies to use, "to the maximum extent permissible under applicable law," the Endangered Species Act's consultation regulation at 50 C.F.R. §402.05.234 Section 7(b) invokes section 301 of the National Emergencies Act to make available the military construction authority of 10 U.S.C. §2808 to the Secretary of the Army.235 Section 6 convenes the Endangered Species Act Committee — the God Squad, whose function is to exempt federal actions from the Act's substantive protections — with an initial eligibility determination within twenty days and resolution within a hundred and forty days.236 And Section 9(c) provides that the order "is not intended to, and does not, create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States."237 An instrument that directs the maximum available use of every environmental waiver on the books, and that by its own terms gives no one standing to complain about how it is used. On the twelfth of January, 2026, a notice at 91 Federal Register 1667 continued the emergency for one year under the National Emergencies Act.238 Three months into the emergency, a separate order had directed the Department of Energy to "streamline, systemize, and expedite" the issuance of exactly the kind of emergency orders described in the last section.239

On the twentieth of March, 2025, Executive Order 14241 invoked Titles III and VII of the Defense Production Act to increase domestic mineral production. Under section 303, a presidential determination is normally required, and projects exceeding fifty million dollars cumulatively normally require authorization from Congress. The order waived those requirements — not by asking Congress, and not by making the section 303 determination, but by relying on the existence of the national energy emergency declared in January.240 Read that slowly. One emergency declaration, made by the executive under one statute, was used to switch off a different statute's requirement that Congress consent. The consent requirement in section 303 is the mechanism by which the elected legislature retains a say over large industrial commitments of public money. It was removed without a vote, by reference to a condition that the same office had declared fourteen weeks earlier, and the removal was itself never voted on. Thirteen months later, on the twentieth of April, 2026, five section 303 determinations issued on a single day: on coal supply chains and baseload power generation capacity, at 91 Federal Register 21927;241 on large-scale energy infrastructure, at 91 Federal Register 21929, which found that without section 303 action "United States industry cannot reasonably be expected to provide these capabilities … in a timely manner due to financing risks, regulatory delays, and market barriers";242 on grid infrastructure, at 91 Federal Register 21931, which brought transformers, transmission lines, substations, high-voltage circuit breakers, power control electronics, protective relay systems, capacitor banks, and electrical core steel within the definition of resources "essential to the national defense";243 and two more on petroleum and on natural gas and liquefied natural gas capacity. A further determination in July 2026, under section 101, found that recoverable critical minerals and materials "are industrial resources necessary to promote the national defense."244 Look at the phrase doing the work in the second: regulatory delays. Regulatory delay is not an act of a foreign adversary. It is the operation of domestic law, including law that exists to give citizens a hearing. When delay itself becomes a finding of national-defense necessity, the emergency has stopped being a response to an external event and become a standing objection to procedure as such. One honest note. I went looking for a Defense Production Act invocation aimed specifically at artificial intelligence or semiconductor capacity in this period and did not find one. The semiconductor trade action of January 2026 proceeded under a tariff authority, not the Production Act. I mention the absence because it cuts against the strongest version of my own argument.

But I did find this. Executive Order 14409, published on the second of June, 2026, directs a classified benchmarking process to determine the threshold at which a model is designated a "covered frontier model."245 The threshold is the entire regulation. Everything else — what obligations attach, who is covered, what is reported — depends on where that line falls. Putting the line inside a classification means the definition that determines the scope of a public rule is itself unavailable to the public it binds. That is definitional capture in its purest available form, and it did not require an emergency at all.

IV.

Which brings me to the development that most changed how I think about this subject, and it is the reason this chapter is not simply a chapter about emergency powers. The acceleration has stopped needing the word. Executive Order 14318, signed on the twenty-third of July, 2025, is titled "Accelerating Federal Permitting of Data Center Infrastructure." It defines a "Data Center Project" as a facility requiring more than a hundred megawatts of new load dedicated to artificial intelligence inference, training, simulation, or synthetic data generation, and a "Qualifying Project" as one with at least five hundred million dollars of capital expenditure, or more than a hundred megawatts of incremental load, or one that "protects national security," or one designated by the Secretary of Defense, the Interior, Commerce, or Energy. Section 5(b) directs the Council on Environmental Quality to establish new categorical exclusions and provides that agencies "shall rely on any sufficient basis to do so as each such agency determines."246 Section 5(c) provides that federal financial assistance amounting to less than half of a project's total cost is presumed not to be a "major Federal action" requiring environmental review at all. Section 7(a) directs the Environmental Protection Agency to expedite these projects by "developing or modifying regulations" under the Clean Air Act, the Clean Water Act, the Superfund statute, and the Toxic Substances Control Act. Section 8(a) directs a programmatic Endangered Species Act consultation covering "common construction activities … over the next 10 years." Section 9 directs Interior and Energy to offer federal sites, and the Army to lease land on military installations.247 The word "emergency" does not appear anywhere in the order. I checked several times because I did not believe it. There is no declaration, no invocation of the National Emergencies Act, no finding of unusual and extraordinary threat. The order simply performs, as ordinary executive management, the substance of what an emergency declaration would have been used to obtain.

This is the ratchet's final position. The first stage is an emergency that justifies a waiver. The second stage is the waiver's renewal on a schedule. The third stage is the waiver without the emergency — because once the accelerated path has been built, staffed, and normalized, the declaration becomes surplus. You do not need to suspend the rules if you have rewritten them. On the twenty-third of April, 2025, the Department of the Interior adopted a document titled "Alternative Arrangements for Compliance with the National Environmental Policy Act amid the National Energy Emergency," authorized by the Council on Environmental Quality that same day. Under it, an environmental assessment is to be completed "within approximately 14 days," with no pre-decisional public comment required, and an environmental impact statement "within approximately 28 days," with no draft statement required.248 I have rarely seen a sentence do so much damage so quietly as no pre-decisional public comment required. The comment period is the only point in the entire federal environmental process at which a person who lives somewhere is entitled to put a fact into a file that a decision-maker must consider. Remove it and the assessment becomes a document the agency writes to itself about a decision it has already made. And on the eighth of January, 2026, the Council on Environmental Quality published a final rule at 91 Federal Register 618 removing all of its regulations implementing the National

Environmental Policy Act — the whole of 40 C.F.R. parts 1500 through 1508, the framework that had governed federal environmental review since the 1970s. The Council's stated rationale: "In the absence of an executive order delegating rulemaking authority to CEQ, the agency lacks independent statutory authority to maintain its NEPA implementing regulations and binding agencies in their implementation of NEPA."249 Six months earlier, the Department of Energy had issued an interim final rule at 90 Federal Register 29676 removing its own NEPA implementing regulations from the Code of Federal Regulations and moving the procedures into a guidance document that is, in the Department's words, "intended to be non-binding guidance."250 There is a serious legal argument behind the Council's rescission and I am not equipped to adjudicate it. What I can observe is the effect: the binding, judicially enforceable, publicly promulgated rulebook has been replaced, at the government-wide level, by nothing, and at the agency level, in at least one case, by guidance that binds no one. A rule you can sue about is a right. A guidance document you cannot sue about is a preference.

One last instrument, because it is where I expected to find the most and found the least. Federal procurement law contains its own emergency, and it requires no declaration by anyone. Under 41 U.S.C. §3304(a)(2), 10 U.S.C. §3204(a)(2), and Federal Acquisition Regulation 6.302-2, an agency may forgo full and open competition where there is an "unusual and compelling urgency" — subject to requesting offers "from as many potential sources as is practicable under the circumstances," a contract term generally capped at one year, and publication of the justification within thirty days.251 Obligations under that exception ran to $40.8 billion in fiscal 2021 at the pandemic peak, roughly $4.0 billion in fiscal 2025, and $18.4 billion in the first half of fiscal 2026 — of which $16.9 billion, ninety-two percent, was the Department of Homeland Security.252 I expected to find artificial intelligence and data-center acquisitions in that number. I did not. The Congressional Research Service's June 2026 review of urgency-exception use offers, as its examples from this period, media-buying services and a swimming-pool lining. I am recording that because it is what the record shows, and because a book that only reports the findings that flatter its thesis is not worth the paper.

V.

Concerns, then, set out plainly.

VI.

Now the question this book keeps returning to. When review is waived, who carries what was in it? Start with money, because money is the easiest part of the answer to verify and the hardest to argue with. For the Campbell plant, Consumers Energy reported spending roughly $401 million to keep the plant online through the thirty-first of March, 2026, recovered about $221 million through wholesale energy sales, and is seeking approximately $180 million from ratepayers.253 TransAlta filed at the Federal Energy Regulatory Commission in April 2026 seeking $19.9 million in recovery for Centralia, with an estimated further $23 million in repair costs if additional orders issue.254 I want to be scrupulous. Those company figures trace to filings with the Securities and Exchange Commission and with FERC, but I read them through trade reporting rather than pulling the filings themselves, and the reader should discount accordingly. Larger program-wide totals are in circulation — $235 million, roughly $1.5 million a day — but every one is an advocacy-group estimate, and a book that demands verifiable numbers from institutions has to hold itself to the same rule. But the money is the least of it, and here the emergency instrument has quietly crossed a line that ought to be visible from a distance.

On the twenty-fourth of January, 2026, the Department issued an order authorizing the Texas grid operator "to direct backup generation resources at data centers and other large load customers to operate as a last resort before declaring an Energy Emergency Alert (EEA) 3 or during an EEA 3."257 In May 2026 it issued a comparable order authorizing PJM "to deploy backup generation resources at data centers and other major facilities."258 There is a defensible reading of these. If a data center has diesel generators sitting idle and the grid is minutes from shedding firm load, using them is obviously better than a blackout. But look at what the instrument has become. A national emergency declared over insufficient energy production is now the legal basis on which a federal department reaches into privately owned equipment inside privately owned facilities built to serve a small number of very large customers. The facilities create load that contributes to the shortfall, and the emergency that the shortfall justifies is then used to conscript their equipment. Texas has legislated in the same direction, requiring the grid operator to develop a service to procure demand reductions from large loads of at least seventy-five megawatts "to be deployed in the event of an anticipated emergency condition," and providing that the operator may direct large loads with on-site backup generation "to either deploy the customer's on-site backup generating facilities or curtail load."259 I do not think this is sinister. I think it is what happens when a category of customer becomes large enough that the public emergency apparatus and the private facility's operations can no longer be described separately. Then there is the clock, which is the part a citizen actually collides with. In Louisiana, the Public Service Commission adopted a directive in December 2025 that, in the words of the Commission's own filing, "provided for certain waivers of the Commission's Market Based Mechanism Order … for a proposed capacity addition where certain factors are demonstrated," "with the objective of the Commission being able to vote within 8 months of an application."260 In Pennsylvania, the Commission ordered that interconnection studies "will be completed within six (6) months from the date that an application for interconnection study is accepted by the Company as complete," with a fifty percent study-fee refund for each ninety-day delay.261 A South Carolina bill would put the smallest tier of data centers on "an expedited review with a hearing being held within forty-five days of the data center's completed application being received by the commission, and a decision made within seventy-five days," and would not require an environmental review statement for that tier at all — though I could not confirm that the bill was enacted and treat it as text only.262

I have no objection to a government that decides things in eight months. My objection is to where the clock starts and whose convenience it serves. A target of eight months from the applicant's completed filing is a commitment made to the applicant. Nothing in it commits anything to the person who has to find out that the filing exists, obtain the figures in it, understand the load and water and rate implications, and get an objection into a docket that is already running. The applicant arrives with counsel, engineers, and a filing team on retainer. The resident arrives with an evening. And it is not the whole picture. New York established a temporary moratorium on data centers by executive order and opened a proceeding in February 2026.263 A Colorado large-load bill was postponed indefinitely on a nine-to-nothing committee vote in May 2026.264 Legislatures are contesting this. That is the good news in the chapter and I do not want to bury it. One more thing about who is on the other side of the table. In May 2026 the Wisconsin commission decided an application that, in the commission's own words, "sought approval of a single Electric Service Agreement to serve one customer."265 A single agreement, for one customer, at a public commission, under an emergency-adjacent posture of urgency about load growth. I am not alleging anything improper; individual contract rates for very large industrial customers have existed for decades. I record it because it is the clearest available illustration of the asymmetry this chapter is about. On one side of the docket, one company with one facility and one agreement. On the

other side, everyone who will pay for the transmission, breathe the emissions of whatever runs to serve it, and live with the outcome — represented, if at all, by whoever read the notice.

VII.

There is a second asymmetry in emergency government, and it is the one that reaches ordinary people most directly. Waiver is fast. Repair is slow. And the repair is usually built out of the same materials as the failure. The pandemic unemployment programs are the cleanest case in the American record. Because the numbers are genuinely contested, I have set out the competing estimates, their methods, and their stated limits in the casebook at the back of this book. The short version is that the Government Accountability Office estimated fraud at between one hundred and one hundred thirty-five billion dollars, roughly eleven to fifteen percent of benefits paid, while saying plainly that "The full extent of UI fraud during the pandemic will likely never be known with certainty" and declining to state a ninety-five percent confidence level because statistical intervals do not capture the uncertainty involved.266 Set the emergency's two speeds side by side. The money went out in weeks. Four to six years later, between seven and twelve cents on the dollar has come back, depending on which denominator you accept — and the denominator itself is disputed, because the recovery rate reported by the responsible department excludes waived overpayments from the calculation.267 By early 2023, twenty-four of fifty-three state workforce agencies had hired one of ten identity-verification contractors using facial recognition to screen claimants. The Department of Labor's own Inspector General set out the mechanism of harm in a single sentence I have not been able to stop thinking about: "a false negative would mean an eligible claimant would be identified as a non-legitimate claimant for benefits." Claimants who cannot verify may be "denied benefits and must start an adjudication process." The Inspector General warned that "without comprehensive guidance, SWAs are at risk of using technology that discriminates against claimants entitled to receive UI benefits," citing federal testing that found false-negative rates ranging "from below 0.5 percent to above 10 percent," with higher false-positive rates for Asian and African American faces. Of the contracts reviewed, seventy-five percent did not address whether identification was one-to-one or one-to-many, sixty-three percent lacked data-storage terms, and fifty-four percent lacked disposal terms.268 Follow the whole arc, because the arc is the argument. An emergency justified paying out with verification suspended. The suspension produced losses on a scale nobody defends. The losses justified a verification system built at emergency speed with the same absence of review. That system, by the government's own account, denies benefits to eligible people at a rate that varies by the color of their face. At every stage the cost of the error was borne by the claimant: first by the honest claimant who waited, then by the honest claimant who was denied, and never at any point by the office that chose the tool. Redress, built under the same emergency logic, generates its own new class of victims — and they are the least able to litigate about it.

VIII.

So what does an emergency that actually ends look like? I set out three worked examples — a British statute drafted with its own death in it, the European rule on urgent deployment, and the American federal picture — in the casebook at the back of this book. Here is what they establish.

The United Kingdom's Coronavirus Act 2020 expired by its own terms after two years, capped each extension at six months, and required a Minister every six months to come to the floor of the House of Commons and move, in a form the statute itself dictates: "That the temporary provisions of the Coronavirus Act 2020 should

not yet expire." If the motion is rejected the provisions lapse within twenty-one days.269 Read that wording again, because the drafting is superb and the superiority is entirely in the default. The government is not asked to justify ending the powers. It is required to ask permission to keep them. Two of the temporary provisions did become permanent, by primary legislation, debated and voted, with a title and a date. Which is exactly right, and is the point. If a society decides after two years of experience that a crisis measure should be kept, that is a legitimate democratic outcome. What is not legitimate is a measure that becomes permanent by nobody ever having to decide. The European Union applies the same discipline to the technology directly, and applies it in hours rather than years. A high-risk system may be authorized early for exceptional reasons of public security or the protection of life, but only "for a limited period while the necessary conformity assessment procedures are being carried out" — and then the sentence that gives the structure teeth: "If the authorisation referred to in paragraph 1 is refused, the use of the high-risk AI system shall be stopped with immediate effect and all the results and outputs of such use shall be immediately discarded."270 The same architecture governs live facial recognition in public spaces: authorization must be requested "at the latest within 24 hours. If such authorisation is rejected, the use shall be stopped with immediate effect and all the data, as well as the results and outputs of that use shall be immediately discarded and deleted."271 Twenty-four hours. Not ninety days, renewable. One sourcing limitation matters: the Official Journal text was not retrievable when I checked, and that article language comes from a faithful reproduction rather than from the Journal itself.

The American federal picture, examined honestly, has less of this than I expected. Of disaster declarations, the Congressional Research Service states with unusual bluntness that they "do not expire, nor do they have set durations established in statute or regulation," and that "generally, federal officials do not unilaterally terminate" them.272 The National Emergencies Act does contain the closest American analogue to the British six-month motion — each House "shall meet to consider a vote on a joint resolution to determine whether that emergency shall be terminated" every six months.273 I tried to establish whether those meetings occur and could not. The parallel duty to report expenditures every six months produced a finding I would put to any committee: the Service was "unable to identify recorded instances" of expenditure reports for emergencies invoking standby authorities other than the international economic powers statute, while noting that the absence of records "does not necessarily mean that the expenditure reports were not submitted," there being "no single repository for reports submitted to Congress."274 I looked for an expenditure report for the energy emergency of January 2025. I did not find one. I cannot tell you it was not filed. I can tell you that a citizen who wants to know what an emergency in its second year has cost has nowhere obvious to look, and that this is itself a finding.

IX.

And then, on the twentieth of February, 2026, something happened that I did not expect and that changes the shape of what can be asked for. The Supreme Court decided Learning Resources, Inc. v. Trump, together with a consolidated case, and held: "IEEPA does not authorize the President to impose tariffs."275 The measures at issue were tariffs imposed under the International Emergency Economic Powers Act by a series of proclamations and orders in early 2025 — twenty-five percent on most Canadian and Mexican imports, ten percent on most Chinese imports, and "reciprocal" tariffs of at least ten percent "on all imports from all trading partners."275 The reasoning matters more than the result. The Court observed that the statute's list of specific powers contains no mention of tariffs or duties, and that "had Congress intended to convey the distinct and extraordinary power to impose tariffs, it would have done so expressly, as it consistently has in other tariff

statutes." It noted that "in IEEPA's half century of existence, no President has invoked the statute to impose any tariffs." And it said this: "There is no exception to the major questions doctrine for emergency statutes."

I have spent this chapter describing a structure whose common feature is the assumption that emergency framing suspends the ordinary requirement that major exercises of power rest on a clear grant from Congress. That assumption has now been rejected by the Supreme Court in plain words, in a case about the most aggressive reading of the broadest available emergency statute. Now the hedges, and there are four. It is a decision about tariffs. Nothing in it holds that the energy emergency is unlawful, that the section 202(c) orders exceed authority, or that a data-center permitting order is defective. Its reach into those questions is a matter of argument, not of holding. The Court was fractured. The Chief Justice announced the judgment and delivered the opinion for six justices as to some parts, but the reasoning I have just quoted appears in parts joined by only three; four separate concurrences and two dissents issued.276 A principle stated in a plurality portion of a divided opinion is a strong argument, not a settled rule. It is a veto, not a construction. The Court told the executive what it may not do under a statute. It did not build anything: no end condition, no accounting requirement, no reopened record, no forum for the person downwind. And the case that would test the argument directly has not been decided. Michigan and other states petitioned the D.C. Circuit to set aside the Campbell order; the consolidated case was argued on the fifteenth of May, 2026 and remains undecided as I write, more than three months later.277 Which is the timing problem in its purest form. The plant was ordered to stay open in May 2025. The challenge was argued a year later. The decision has not issued. And the emergency has been renewed five times in the interval, each renewal a fresh fact on the ground, each one making the eventual ruling more nearly moot. A ninety-day instrument litigated on a fifteen-month calendar cannot be corrected by litigation. It can only be outlasted.

X.

Objections I take seriously. "Emergencies are real, and you are describing a country that cannot respond to one." True, and I should not be read as denying it. Grid emergencies kill people; so do hurricanes and pandemics. Rapid executive action under pre-authorized standby powers is a legitimate feature of a functioning state, and I would not remove a single one of the acute-event orders described above. I am not asking for the first ninety days. I am asking that the ninety-first require an argument. "The reliability warnings were genuine, and hindsight is cheap." Also fair. The reserve-margin assessments, the grid alerts, the capacity-deficit projections are real documents produced by serious people, and the load growth driving them is real too.228 If the choice is between an old coal plant running for one more summer and rolling blackouts in July, I know which I would pick. But the choice is being presented as a binary by the only party entitled to describe it, in a proceeding with no opposing side, and the alternatives that would have avoided the binary — the transmission that was not built, the demand-side program that was not procured, the interconnection queue that was not cleared — are precisely the things that a permanent emergency posture crowds out. An emergency justified by a shortage that predictable planning would have prevented is not an emergency. It is an outcome. "The renewals are lawful. Congress wrote a renewable authority." Correct, and almost nothing in this chapter is an allegation of illegality. Congress wrote a ninety-day order renewable in ninety-day periods and did not cap the renewals, and the Department is doing what the statute permits. That is the problem I am describing. The gap is legislative, which means it is fixable by legislation. "You are demanding process for its own sake, and process is what got us into a nine-year interconnection queue." Partly true, and the objection I feel the pull of most. I have watched review used as an instrument of obstruction by people who had no intention of ever saying yes, and a

proceeding with unlimited entry points is a proceeding with no decision. I am not asking for indefinite review, unlimited standing, or a veto for every objector. A six-month clock is fine. A one-year clock is fine. What is not fine is a clock that begins when the applicant is ready and ends before the public has been told. Louisiana's eight months and Pennsylvania's six are not too short. They are asymmetrically allocated. "Your cost numbers are too small — the real ratepayer burden is far higher." Possibly so. I have declined to use the advocacy estimates in circulation because they are estimates by parties to the dispute and I would not accept the equivalent from the other side.256 If a government audit produces a program-wide figure, I will use that, and I expect it will not flatter the program. "You had a claim wrong." I did, and it belongs in the chapter rather than in a private correction. In an earlier draft I had written that Texas's 2025 large-load statute offered expedited interconnection in exchange for a flexible-load commitment of at least ten years. When I went back to the enrolled text and the legislative analysis to source it, no such provision existed — not the expedited processing, not the ten-year commitment, not the trade between them. I had absorbed it from a summary, and it was wrong. I also nearly described a Louisiana executive order on economic-development coordination as a data-center fast-track; the order does not mention data centers or energy loads at all.

I put those in because this chapter is an argument about the discipline of having to justify a claim on a record, and the argument would be worthless coming from someone unwilling to do it. Two withdrawn claims cost me an afternoon and improved the chapter. A withdrawn permit costs somebody a project, which is exactly why the party seeking the permit prefers a process in which nobody has time to check.

XI.

Three questions I would ask about any emergency instrument, in any forum where it can be raised. None of them requires legal training. Who can witness it? Emergency proceedings are precisely where outside observation is removed, and that is what the waiver is for. For the orders in this chapter, a member of the public could have seen the order, after it issued, on a web page, with no docket in which to respond. For the fourteen-day environmental assessment, no pre-decisional comment is required at all.248 For the classified benchmark that will determine which models are regulated, the answer is nobody by design.245 An emergency that cannot be witnessed cannot be evaluated, and an emergency that cannot be evaluated cannot be ended, because ending it requires somebody to show that the condition has passed. Who carries the burden if the waiver was wrong? Look for the asymmetry every time. If the retirement deferrals turn out to have been unnecessary, the Department bears nothing; the utility recovers its costs through rates; and the ratepayers of five states and the people living nearest six power plants carry the entire difference.253–255

If the facial-recognition screen produces a false negative, the agency's cost is an adjudication file and the claimant's cost is rent.268 Whenever the answer to what happens to the institution is nothing, the error rate will not improve, because nothing in the arrangement asks it to. How does it end? Not "when do we hope it ends" — what event, stated in advance, terminates it; who is obliged to declare that the event has occurred; what happens automatically if they do not; and what public accounting is required at expiry. A renewable ninety-day order with no cap has no answer.218 A Stafford declaration has no answer in statute at all.272 A national emergency in its second year has an annual continuation notice, which is the opposite of an answer.238 The British statute had an answer, and used it.278 The European regulation's answer is twenty-four hours and the discarding of everything the system produced.271 An instrument that fails all three — unwitnessed, costless to

its issuer if wrong, and with no stated ending — is not emergency management. It is a change in the form of government, adopted one quarter at a time.

XII.

What I would insist on. Five conditions, with a name attached and a consequence for failure.

1. No environmental, civil-rights, labor, or public-participation law suspended on grounds of technological emergency or international competition. Not narrowed, not expedited-with-safeguards — not suspended. These are the bodies of law whose entire function is to protect people who are not parties to the transaction, which is precisely why an accelerated process reaches for them first. Competitive pressure from another country is not an emergency; it is the ordinary condition of an open economy, and it has no end condition, which is what makes it the most dangerous justification in current use. If domestic law is genuinely an obstacle to a national objective, the answer is to change the law by legislation, in public, with a vote, and to accept the possibility of losing that vote. Who: the head of each agency issuing the waiver, by name in the instrument; the Council on Environmental Quality; the committee of jurisdiction. If not: the waiver is void and the approval it enabled is vacated, with the applicant on notice from the date of the waiver that its authorization is contingent.

2. Every emergency authority carries a stated end condition, a sunset date, and a required public accounting at expiry. The end condition must be an event, described in advance in the instrument, that a person outside the issuing agency could recognize — not "until the Secretary determines," which describes discretion rather than an ending. The sunset must be absolute, not renewable, so that continuation requires a fresh instrument on a fresh record. At expiry, a public accounting: what was waived, what was built, what it cost, who paid, and what harm was identified. The British model is the one to copy, and copy the default: a motion in the legislature every six months, in the words the statute prescribes, asking permission for the powers to continue, with automatic expiry if the motion is not carried.269 Where the statute already requires periodic consideration and an expenditure report, as the National Emergencies Act does, the requirement should be enforceable and the report should be published in one findable place.273274 Who: the declaring authority; the presiding officer of each chamber, on whom the scheduling duty falls; the Comptroller General for the accounting. If not: the authority lapses by operation of law on the sunset date, no funds may be obligated under it thereafter, and actions taken after lapse are void.

3. The party declaring the emergency may not be the party that benefits from the waiver, and the beneficiary must be named on the face of the instrument. Where an emergency instrument confers an identifiable private benefit — a plant kept running, a permit accelerated, a facility connected, an agreement approved for a single customer — the instrument must name the beneficiary, state the value of the benefit, and be issued by an authority that does not gain from the outcome. A commission deciding "a single Electric Service Agreement to serve one customer" under a general posture of urgency about load growth is not a neutral proceeding unless somebody is representing the other side of the table.265 Who: the issuing official personally; the agency's designated ethics official; the state public service commission's independent consumer advocate, who must be funded and appear. If not: the instrument is voidable on petition by any affected ratepayer or resident without further showing of injury, and the beneficiary bears the cost of the reopened proceeding.

4. A renewal is a new decision. Beyond a first term, continuation requires: a fresh factual record; publication of that record before the renewal issues; a public comment window of stated length; a written response to the

substance of the objections filed against the previous instrument; and a finding, on the new record, that the condition still exists and that the alternatives developed since the last renewal remain inadequate. The five-times-renewed order would have had to answer, on the record, in November 2025, why the transmission and demand-side alternatives available since May had not been pursued. It never had to answer that, and so, predictably, they were not. Who: the issuing department; the grid operator or equivalent technical body, on the record and in its own voice rather than the department's; the affected state utility commissions as parties of right. If not: the renewal does not issue, and the underlying authority expires at the end of its stated term.

5. Retroactive review with real remedies where a waiver caused harm. Every emergency waiver is a bet made on the public's behalf without the public's consent, and the minimum decency is that we settle up. Within a stated period after expiry, an independent body examines what the waiver permitted and what it cost, with authority to order remediation, to reallocate costs away from ratepayers and onto the beneficiary, and to void approvals obtained through review since shown inadequate. The model already exists in European law: where an urgent authorization is refused, use stops immediately and "all the results and outputs of such use shall be immediately discarded."270 109 Who: the Comptroller General or a state auditor with subpoena authority; the reviewing court on petition; the commission that approved the cost recovery. If not: costs recovered under the waiver are refunded with interest, and the approval is reopened with the withheld material in the file. I know how much machinery that is, so take the two that cannot be given up. Number two, the stated ending — because everything else in this chapter follows from the absence of one. An emergency with a real end condition and an accounting at expiry can be tolerated even if it is wrong, because it terminates and we find out what it cost. An emergency without one is not a measure at all. It is a change in the constitution of the thing, adopted without amendment. And number four, that a renewal is a new decision — because it is the cheapest reform in this book and the most immediately effective. It requires no new agency, no new right of action, no new appropriation. It requires one sentence in a statute providing that the ninety-first day is argued from the beginning.

XIII.

I said at the outset that I have used this instrument, and I want to close there rather than on the legislation, because the legislation will not be what decides this. What I remember about running things under emergency conditions is how good it felt, and how much I liked the version of myself that operated that way. Decisive. Unblocked. Free of the people who wanted to talk about it. And what I remember about the aftermath is that the decisions made in those periods were not better than the ones made slowly. Some were worse, in the specific way that decisions are worse when nobody was permitted to tell you what was wrong with them. The emergency did not improve my judgment. It removed the mechanism by which my judgment got corrected, and then it flattered me for the absence. So I would put it as plainly as I can. A decision made fast enough to escape scrutiny has not been made democratically, whatever its merits. It may be right. It may be necessary. It may be the very thing I would have voted for. But it was not decided; it was announced, and the announcement was arranged so that the announcing was all that could occur. Emergency coordination and permanent sovereignty are different things. The first is a capacity every competent state needs: the ability to act in hours when hours are what there is. The second is a claim of authority that does not have to be renewed by consent, and there is a word for it, and the word is not emergency. The whole art of citizenship in this period is refusing to let the first become the second. Not by opposing speed, but by insisting, in every forum available, on the one question that distinguishes them. Who gets to decide that there is no time to ask me? If somebody can be made to answer that

in writing, with a date attached, in a room where I am entitled to be, then whatever else is happening, we are still governing ourselves. If the question has no addressee, then the emergency has already ended, in the only sense that matters, and it ended by becoming permanent. — D.M.D.

Chapter Six: Force and Autonomous Harm

Policing, Borders, Weapons, and the Accountability Gap

Policing, Borders, Weapons, and the Accountability Gap

Which decisions about a person should never be made by a machine?

I.

On the sixteenth of February, 2023, Porcha Woodruff was arrested at her home in Detroit and charged with carjacking and robbery. She was eight months pregnant. She told the officers they had the wrong person. She was held, processed, and released. The charges were later dropped, and in August 2023, she sued. On the fifth of August, 2025, a federal judge granted summary judgment against her.279 I start with the case that lost, because it is the one that shows you the machinery. Two weeks before the arrest, an investigator ran a still image from a service-station surveillance video through a facial recognition system. The image was not high definition. The person in it wore a hat that covered the forehead, the eyebrows, and the hairline — conditions the court acknowledged "could have affected the accuracy of the facial recognition search." The system returned a list of seventy-three candidates. The investigator narrowed the list to one candidate: Porcha Woodruff. The method by which he did so does not appear anywhere in the record.280 Read that again, because it is the whole chapter in a single fact. A machine produced seventy-three names. A human selected one. And when a federal court came to examine how a pregnant woman ended up in a holding cell, the step at which the seventy-three became the one was simply not documented. The court then did something worth stating precisely. It declined to decide whether facial recognition had tainted the identification at all. Civil-liberties organizations had filed as friends of the court arguing exactly that, and the plaintiff expressly did not adopt their arguments, so the court wrote: "As stated earlier, Plaintiff explicitly rejected these arguments, and the Court therefore declines to consider the issue." And then it added, in the alternative, that even if the arrest warrant had lacked probable cause because of the facial recognition, the officer would likely be protected by qualified immunity, because no existing case had given her "fair warning" that using the technology this way was unlawful.281 So: a machine that produced a list. A human step that left no record. A court that did not reach the technology. And an immunity doctrine that protects the officer precisely because no court has yet reached the technology. The loop is closed.

I should admit something. I have witnessed and participated in building systems that made determinations about people. Not policing systems — access decisions, eligibility decisions, fraud flags, scores that governed whether someone got the thing they had asked for. And what I remember is the order in which we built the pieces. We built the model first. Then the interface. Then the throughput monitoring, because throughput was what the business cared about. Then, quite late and usually because someone in a support function insisted, we built the escalation path — the way a person could say this is wrong about me and reach a human with the authority to change it. Its purpose, if I am honest about how we discussed it internally, was to handle the volume of complaints rather than to correct the volume of errors. A system optimized for the first will absorb complaints and change nothing. A system optimized for the second is expensive, because it has to be able to tell you that the model you paid for is wrong. And I remember what a human in the loop meant on our product roadmaps. It did not mean a person exercising judgment. It meant a person clicking to approve, at a rate that was measured, under an efficiency expectation that made careful review a performance problem. I do not think we

were unusual. I think we were ordinary, and I think that is worse. That is why the seventy-three matters more to me than any argument about accuracy rates. The human step was there. It was documented nowhere and constrained by nothing, and it was the step on which the entire legal defensibility of the arrest depended.

II.

Start with what the machine actually produces, because almost every public argument about facial recognition misdescribes it. A facial recognition search does not identify a person. It ranks a gallery. Given a probe image, the system scores every face in the database it has access to and returns the closest matches in order of similarity, and that list — the candidate list — is the output. Seventy-three, in the Woodruff case. The conversion of a ranked list into a suspect is performed by a human being, and it is at that conversion, not at the algorithm, that the accountability disappears.113

The National Institute of Standards and Technology has run the most systematic public evaluation of facial recognition in existence. Its 2019 report on demographic effects tested one hundred eighty-nine algorithms from ninety-nine developers against 18.27 million images of 8.49 million people, and its findings are stated flatly: "Across demographics, false-positive rates often vary by factors of 10 to beyond 100 times." False negatives "tend to be more algorithm-specific, and vary often by factors below 3." And on the specifics: "We found false positives to be between 2 and 5 times higher in women than men, the multiple varying with algorithm, country of origin and age," and "In domestic mugshots, false negatives are higher in Asian and American Indian individuals, with error rates above those in white and black faces."282 A later NIST working document put the spread more starkly, reporting within-group false-positive variation up to a factor of several hundred while false-negative variation stayed near a factor of three — and this document exists publicly only as an initial public draft. I could not locate a final version, nor any newer numbered NIST report on demographic differentials. The most current government numbers live in continuously updated web tables whose demographic sections were last revised in March of 2025.283284

A false positive here does not mean the system announces a wrong name. It means the wrong face is on the list — ranked, plausible, in the pile of seventy-three, waiting for a human being to pick it out for reasons that will not be written down. The one court I found that has taken this seriously took it seriously procedurally. In 2023 a New Jersey appellate court reversed a trial judge and ordered the State to hand over eleven categories of discovery about a facial recognition search that had produced the defendant, including false match rate and false non-match rate data. The court's reasoning is the most honest judicial sentence I found in this entire body of research: "Given FRT's novelty, no one, including us, can reasonably conclude without the discovery whether the evidence is exculpatory or 'merely potentially useful evidence.'" And: "The evidence sought here is directly tied to the defense's ability to test the reliability of the FRT."285 A judge saying, on the record, that the court cannot evaluate the thing that produced the defendant. It is not a ruling about accuracy. It is a ruling about darkness. I could not find a government tally of arrests attributable to facial recognition errors. Not a national one, not an agency one, not a state one. The only running count I located is maintained by an advocacy organization, which as of April 2026 described more than a dozen such arrests.286 There is no federal register of these events. Nobody is required to report one. The true number is unknown to the public and, as far as I can determine from the documents, unknown to the agencies as well.

III.

There is one place in the United States where a detailed, enforceable rulebook for police facial recognition exists in writing. It exists because one man was arrested on his own front lawn in front of his daughters and sued for three years. Robert Williams was arrested by Detroit police in January of 2019 after a facial recognition search. He filed suit in April of 2021. On the twenty-eighth of June, 2024, a federal court entered a stipulated order of dismissal, attaching a settlement agreement, with the court retaining jurisdiction to enforce it for four years.287 "Members shall not use Facial Recognition to surveil the public through any camera or video device." "Facial Recognition technology may only be used on a still image of an individual, including still images captured from video." "A request for an arrest warrant, or an arrest, shall not be made solely on the basis of an investigative lead developed through Facial Recognition technology in combination with a lineup identification." And the clause that shows the deepest understanding of the failure mode: "The photographic lineup shall not contain an image derived from facial recognition."288 That last clause is the sharpest thing in the document. If a search produces a candidate, and you then put that candidate's photograph into a six-pack lineup and ask a witness to pick, the witness is not independently corroborating the machine. The witness is being handed the machine's answer and asked to ratify it, and the ratification will then be described in the warrant affidavit as an eyewitness identification. The settlement forbids it. Nothing in federal law does. The agreement also requires a retrospective audit of every facial recognition request logged from the twenty-second of February, 2017 forward, traced through to any resulting warrant or arrest, and a four-year training program.289 That rulebook was in force when the Woodruff case was decided. The reform arrived through private litigation, on behalf of one plaintiff, binding one department, in one city. It did not arrive through the legislature of Michigan or of the United States, neither of which has enacted anything on the subject. This is the pattern to watch for through the rest of the chapter. The rules that exist are overwhelmingly the residue of individual lawsuits, individual audits, and individual budget crises. They are not the product of anyone deciding in advance which decisions about a person a machine should not make.

IV.

Prediction is the older and larger story, and we can now look at how three programs ended. A CITIZEN'S MANIFESTO FOR THE AGE OF AI 115 Chicago ran a system that assigned people a numerical score for their likelihood of involvement in gun violence. In January of 2020 the City's Office of Inspector General published an advisory on it. As of July 2018, 399,412 individuals had a score under one version of the model; as of March 2019, 313,513 had a risk tier under its successor.290291 Of the 398,684 people scored under one version, 16.3 percent were identified as confirmed gang members — meaning roughly five in six of the people the system had scored were not.292 The Inspector General's finding on quality was four words long: "Risk scores and tiers were unreliable." Until the successor model launched in January of 2019, the scores "had not been updated since August 2016, resulting in scores which were static for nearly two and a half years."293 For two and a half years, a police department in a major American city held a risk score on four hundred thousand people that could not go down no matter what any of them did. Because the update job was not running. The program cost $3.8 million in federal grants.294 The department decommissioned the model on the first of November, 2019.295 Los Angeles ran two related things: a place-based prediction product and a person-based "Chronic Offender" designation. Its Office of Inspector General reviewed both in March of 2019, and the sentence that matters is this: "the majority of people identified as Chronic Offenders had few, if any, actual contacts with the police."296 The database had been suspended in August of 2018 holding 637 individuals, of

whom 234 were marked active and 403 inactive.297 On the fifteenth of April, 2020, the department discontinued the prediction product. Its own memorandum cited the cost of "an annual professional services agreement" and "The City's financial crisis, coupled with impact of the COVID-19 pandemic," and concluded that "the Department will immediately discontinue the use of PredPol and its associated reports."298 It ended because of the price and the pandemic. Not because the audit persuaded anyone. DAVID M. DOYLE 116 The third program ended in a courtroom. A Florida sheriff's office ran an intelligence-led policing program that generated a list of people expected to commit future crimes and then subjected them and their families to repeated home visits, code-enforcement citations, and arrests. A federal judge, denying most of a motion to dismiss in August of 2021, described the program as using a "crude computer algorithm to identify individuals who may commit unspecified future crimes," noted that targeted people received no notice that they were on the list, and found that the plaintiffs had stated claims under the First Amendment and procedural due process.299 The case settled in December of 2024, and the plaintiffs' counsel reported that the sheriff agreed the program would not resume and acknowledged constitutional violations — I flag that this account comes from a party to the litigation rather than from a court order, because I was unable to obtain the signed agreement itself.300 One ended after an audit, one over a line item, one in litigation. I could not locate a single American state statute or municipal ordinance, in force in 2026, that prohibits predictive policing outright. The only in-force instrument I found that names the practice at all is a Boston ordinance that subjects "software designed to integrate or analyze data from Surveillance Technology, including surveillance target tracking and predictive policing software" to a council approval and oversight process — regulation, not prohibition.301302

So we have a technology that three separate public bodies examined and found unreliable, unlawful, or not worth the money, and that remains perfectly legal to buy tomorrow in nearly every jurisdiction in the country.

V.

Concerns, then, set out plainly.

VI.

Now the border, where the least is owed to the person and the most is done to them. The Department of Homeland Security publishes an annual inventory of its artificial intelligence use cases. It is a spreadsheet, it is public, it runs to nearly two hundred rows and seventy-two columns per row, and it is the single most useful document I found in the research for this chapter — not because it is critical of anything, but because it is an agency's own structured self-description, field by field.303 Row DHS-2408 describes a model called the Hurricane Score. In the agency's words, it is "a quasi-binomial, binary classification machine learning (ML) model" that "models the potential risk (1-5) that a noncitizen who is released from detention with the requirement to check in with ICE through monitoring will abscond from the program, with a higher number indicating a higher risk of absconding." It operates inside the alternatives-to-detention program, helping determine how closely a person is supervised. The stated purpose is candid: the agency "would like to explore technology that could help lead to time savings and eliminate human error due to overburdened officers."304 Then the fields. Is the use case rights-impacting? "Yes." Presumed rights-impacting category: "Law Enforcement."305 Which demographic features does the model use? "Sex/Gender, Age, Other" — and, under the field asking what "Other" means: "Country of Origin." What steps has the agency taken to detect and mitigate significant disparities in the model's performance across demographic groups? The entry is one word.

"No." Has the Senior Agency Official for Privacy assessed the privacy risks associated with this use case? "No." Does this use case have an authority to operate? "No."306 The agency records the risk itself, correctly: "There is a risk that the Absconder Model may include false positives and model bias."307 And then the field I have not been able to stop thinking about. The inventory asks whether affected individuals are notified that artificial intelligence is being used. The answer reads: "N/A - individuals are not interacting with the AI for this use case."308 The person is not interacting with the system. That is true, in the sense the field means it: the noncitizen does not log in, does not type anything, does not see a screen. They are the subject of the computation, not its user. Because the form's category of "interaction" is built around the user rather than the subject, the person about whom a rights-impacting score is generated falls outside the category that would have triggered a duty to tell them. That is the accountability gap in one cell of a spreadsheet. Not a conspiracy. A category error, formalized, and then filled in accurately. Human oversight is described as "Intermittent and Manually Updated," and the agency's own gloss on the human role is: "H-Score adds one piece of info to a larger profile for decision-making by a human."309 One piece of info in a larger profile, for a decision made by a human. Which is exactly what facial recognition is in the Woodruff case, and exactly what the risk score is in the sentencing cases I come to next. The machine's defenders are right that it does not decide. And in each instance the record of what the human did with it is thin, unstandardized, or absent, so the claim cannot be checked.118 The most recent inventory I could confirm was released in December of 2024; I could not confirm publication of a 2025 or 2026 inventory as of the end of August 2026.310 The best public window into federal enforcement automation is about twenty months stale. Now the complication. The Department's own Inspector General audited the immigration risk-assessment process and found that in fiscal 2023, officers failed to conduct a risk assessment for a third of detained noncitizens, and in fiscal 2022 for 43 percent.311 So the same system is simultaneously applied to people who are never told about it and not applied to a third of the population it is supposed to cover. A serious critic will say the problem here is inconsistency rather than automation, and that a properly and universally applied instrument would be fairer than the officer discretion it replaced. I answer that later. One more thing about the border. A 2025 privacy impact assessment for a departmental mobile application describes a "Submit Intent to Depart" function; explains that facial comparison is performed one-to-one so that "the exit is recorded as a biometrically confirmed exit"; notes that the system of record retains such records "for 75 years"; and, in the pattern common to these documents, concludes several risk discussions with the sentence "This risk is fully mitigated."312 Seventy-five years is not a data-retention decision. It is a life. And the federal picture on facial recognition generally, from the Government Accountability Office: of forty-two federal agencies employing law enforcement officers surveyed, twenty reported owning or using facial recognition systems; six reported using the technology on images of the protests following the killing of George Floyd; three on images of the events at the Capitol on the sixth of January, 2021.313 Two years later, the same office found that six agencies had cumulatively conducted about sixty thousand searches during periods when no training requirement was in place, and that at the Federal Bureau of Investigation "only 10 staff completed facial recognition training of 196 staff that accessed the service."314 Sixty thousand searches. Ten trained people out of a hundred and ninety-six with access. This is not a story about a dangerous technology in careful hands.

VII.

The courts are the place where the argument is hardest for me, because in the courts the tool is more honest than the system around it, and because one of the instruments in question sets people free.119

Start with the case everyone cites and few quote. In 2016 the Wisconsin Supreme Court held that "if used properly, observing the limitations and cautions set forth herein, a circuit court's consideration of a COMPAS risk assessment at sentencing does not violate a defendant's right to due process."315 That sentence is normally where the citation stops. The limitations are the interesting part: a score may not be the determinative factor in deciding whether an offender can be safely supervised in the community; the court must explain the factors that independently support the sentence; any presentence report containing a score must carry specified cautions; and "risk scores may not be used: (1) to determine whether an offender is incarcerated; or (2) to determine the severity of the sentence."316 The decision is not a green light. It creates a right exercised by objection, which means its practical reach is a function of how well-resourced the defense is, and does not protect the person who takes the first offer. Now the federal instrument, and the concession I owe. The First Step Act's risk tool determines eligibility for earned time credits, meaning a favorable score can shorten the time a person spends in custody. This is not a purely carceral instrument. It releases people. The Department of Justice has done the validation work and published it. Its most recent review reports that the tool "remains a strong and valid predictor of general and violent recidivism at the one-, two-, and three-year follow-up periods, with Area Under the Curve (AUC) statistics ranging from .745 to .776."317 And in the same document: an acknowledgment that the instruments overpredict recidivism risk for Black, Hispanic, and Asian men and women relative to white individuals, alongside the concession that the agency has not yet worked out how to fulfil the statutory mandate that such disparities be reduced. Credit is due and rare. The government built the tool, validated it, measured its racial performance, published the finding, and acknowledged a statutory duty it has not yet discharged. That is what institutional honesty looks like. Which is why the next fact is the one I would put to a committee. That review covers the 2019 cohort and appeared in August of 2024. As of the day I write this, at the end of August 2026, I could not locate any subsequent review or revalidation, nor any annual report under the statute's reporting provision later than June of 2024.319 I am not asserting that no review was conducted; I am asserting that after searching the responsible agencies' own publication pages I could not find one, and that a duty whose discharge cannot be confirmed by a member of the public is not functioning as a duty. Periodic-review obligations are the most common accountability mechanism in this field, and the easiest to let lapse, because nothing happens when they do.

One more court-side finding, and it is the single most transferable provision I encountered anywhere in the research. Idaho, in 2019, required pretrial risk assessment tools to be "transparent," to be "open to public inspection, auditing, and testing," and entitled a party to "review all calculations and data used to calculate the defendant's own risk score." Then it added the clause no other jurisdiction I examined has: "No builder or user of a pretrial risk assessment tool may assert trade secret or other intellectual property protections" in order to quash discovery of those materials.320 One sentence. It costs nothing. It abolishes, in one American state, the objection that has done more than any other to keep these instruments unexamined. I could not confirm from a fetched code page its current codified status in 2026, and I say so rather than assert it — but the drafting exists, it has existed for seven years, and it is available to copy. And the corresponding negative: I could not find any state statute prohibiting the use of pretrial risk assessment tools, and I could not verify from primary sources the current status of the validation requirements a large state is often said to have enacted.321

VIII.

Somewhere behind every system in this chapter there is a company, and whether the company can be reached is the question on which most of the rest depends. The best-documented instance in American law comes not

from policing but from unemployment insurance. Michigan built a system that determined fraud automatically. Two plaintiffs were told they owed $19,910 and $8,302; the state intercepted the first man's state and federal tax refunds; the class was in the range of twenty-six thousand people; the case settled for twenty million dollars eight years later.322323 Two private contractors had built the adjudication component, and they argued they could not be sued under the federal civil-rights statute because they were not the state. In January of 2022 a federal judge rejected that argument: "The state worked hand in hand with CSG to develop and implement the robo-fraud-adjudication system," and "This nexus between the state and the challenged action far exceeds 'mere cooperation,' contract, authorization, restriction, or funding."324 That is a ruling on whether the vendors may be sued, not a judgment that they are liable. But it establishes the principle on facts that recur everywhere: a state that could not have built the thing, a vendor that designed the adjudication logic, and a person whose tax refund was taken by a determination no human being had made.

The underlying constitutional standard is older than any of this technology: "'[T]he root requirement' of the Due Process Clause [is] 'that an individual be given an opportunity for a hearing before he is deprived of any significant property interest.'"325 Before. Not after the refund is taken and the person has spent eighteen months proving a negative. I could not locate a single statute, federal or state, that assigns liability to the maker of a software system for a harm caused by that system in an enforcement context.326 Not one. The doctrine exists because a district judge in Michigan worked it out from constitutional first principles. It is available to the next plaintiff who can afford five years of litigation and is not available to anyone else. The European Union built two routes to the same destination and then closed one, in the open, with the reason given. It kept product liability, and its revised directive states that "product" "includes electricity, digital manufacturing files, raw materials and software" — software is a product — and then addresses the evidentiary problem that actually defeats these cases. A court "shall presume" defectiveness or causation where the claimant "faces excessive difficulties, in particular due to technical or scientific complexity" in proving them, subject to rebuttal, and the later-defect defence is unavailable where the defect lies in software within the manufacturer's control.327328329 What it abandoned was the dedicated instrument on liability for harm caused by artificial intelligence, listed for withdrawal in February of 2025 with the stated reason "No foreseeable agreement," and formally withdrawn that October.330331 The jurisdiction that has gone furthest kept strict liability for defective products, including software, and abandoned the instrument that would have eased proof for AI-caused harm generally. That is a workable theory when a machine crushes someone and an awkward one when a ranked list of seventy-three names ends with a pregnant woman in a cell.

IX.

Which brings me to weapons, and to the most consequential absence I found in any document in two years of research on this book. DAVID M. DOYLE 122 The governing United States policy on autonomy in weapon systems is a Department of Defense directive, most recently reissued on the twenty-fifth of January, 2023. It defines its subject cleanly: an autonomous weapon system is "A weapon system that, once activated, can select and engage targets without further intervention by an operator."332 Its central substantive requirement is a single sentence: "Autonomous and semi-autonomous weapon systems will be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force."333 Its procedural requirement is a two-stage senior review, with four exempt categories.334 And there is a waiver: "In cases of urgent military need," a Deputy Secretary of Defense waiver of those requirements may be requested.335 Urgent need, again. I searched the text of that directive for the word "prohibit." It appears zero times. I

searched it for the word "meaningful." It appears zero times.336 The governing American instrument on autonomous weapons forbids nothing. It does not prohibit any category of autonomous weapon, including one designed to select and engage human beings. Here is the government's own explanation, from a working paper the United States submitted to the international expert process in 2018. The directive's human-judgment standard, it says, "reflect[s] a deliberate decision to permit weapons that are programmed to make 'decisions' that relate to targeting."337 A deliberate decision to permit. The absence of a prohibition is the policy, stated as such, in a document the United States filed with an international body. The same working paper explains why it uses "judgment" rather than "control": "'Human judgment over the use of force' is distinct from human control over the weapon. For example, an operator might be able to exercise meaningful control over every aspect of a weapon system, but if the operator is only reflexively pressing a button to approve strikes recommended by the weapon system, the operator would be exercising little, if any, judgment over the use of force."338 That is correct. It is the criticism I made of my own products earlier in this chapter, made better, by the government, eight years ago. The United States identified the hollowness of nominal human oversight before most of its critics did, and then drew from that insight the conclusion that the standard should be flexible rather than enforceable. The directive has not been amended since January of 2023.339 In January of 2026 the department issued an artificial intelligence strategy naming seven pace-setting projects.340 It is an acceleration document, and it leaves the 2023 review architecture where it was.

Meanwhile the flagship autonomy program illustrates the oversight problem without needing any critic's gloss. Congress's own research service reports that the initiative was unveiled in August of 2023 to field thousands of uncrewed systems by August of 2025, and that "Although Replicator was intended to field thousands of uncrewed systems by summer 2025, one former defense official noted that DOD had fielded only hundreds by that time." On money: "Little information is available publicly about Replicator's potential total cost." On oversight: "Some Members of Congress have stated that it has been difficult for them to obtain information about Replicator," and "Lack of information on Replicator, either in the classified or unclassified realms, has the potential to raise doubts about whether DOD has adequately analyzed the initiative."341 And when the department's own Inspector General evaluated the program in September of 2025, the public page for that report carries one sentence where the findings would be: "This report contains classified information and no redacted version is available."342 I am not asserting that the classification is improper. I have no way to know, and the subject is one where classification is often exactly right. I am recording what the citizen can see: a program of unknown total cost, behind schedule by an order of magnitude on its own stated goal, which members of the legislature say they cannot get information about, and whose independent evaluation exists in a form no member of the public may read. That is the entire evidentiary position available to a person who wants to form a view.

X.

The international process is where the prohibition was supposed to come from, and its record over the last three years is a study in how a deadline dies. In October of 2023 the Secretary-General of the United Nations and the President of the International Committee of the Red Cross issued a joint call: "The autonomous targeting of humans by machines is a moral line that we must not cross. Machines with the power and discretion to take lives without human involvement should be prohibited by international law." And it set a date: conclude negotiations of a new legally binding instrument "by 2026."343

Then watch two lines move in opposite directions. The votes get stronger: the General Assembly's first resolution was adopted 152 to 4 with 11 abstentions;344 the next, 166 to 3 with 15;345 the most recent, 164 to 6 with 7.346 The language gets weaker. The 2024 text asked the expert group to fulfil its mandate "as soon as possible, preferably by the end of 2025."345 The 2025 text asks states "to work towards completing the set of elements for an instrument being developed within the mandate of the Group of Governmental Experts …, with a view to future negotiations."346 In one year the objective moved from concluding a mandate by a date to working towards completing a set of elements with a view to future negotiations. The sentence got longer and the verb got softer, and a hundred and sixty-four countries voted for it. I am writing on the twenty-ninth of August, 2026. Five days ago the Secretary-General and the President of the Red Cross issued a renewed call whose main verb is the plainest evidence that the 2026 deadline was missed: "Negotiations must begin now to urgently adopt a legally binding instrument to regulate autonomous weapon systems with clear prohibitions and restrictions."347 I cannot show you a case. There is no official document from any government, any United Nations body, any court, or any inspectorate that makes an affirmative factual finding that an autonomous weapon selected a human target and killed that person without human intervention. The claim most frequently cited traces to a single paragraph in a 2021 report of a United Nations panel of experts on Libya. It says retreating forces "were subsequently hunted down and remotely engaged by the unmanned combat aerial vehicles or the lethal autonomous weapons systems." The word doing the work is "or." The panel attributes no casualty and no death to an autonomous engagement. The sourcing note for the surrounding passage reads, in its entirety, "Confidential source." And the annex cited in support is an infographic.348 That paragraph has been reported worldwide as the first autonomous killing of a human being by a machine. It does not say that. The strongest recent official material I could obtain is a 2024 press release in which United Nations independent experts deplored the "purported use of artificial intelligence" in target selection and wrote "If proven true, the shocking revelations …" That is a document written conditionally by rapporteurs in their individual capacity. It is not a finding of fact and I will not present it as one.349

So my position is this. The prohibition I am about to argue for is not justified by a body count. It is justified by a structure: an instrument that permits by design, a review process with exemptions and an urgency waiver, an evaluation the public cannot read, a program whose cost is not public, and an international process whose deadline passed without anyone being required to explain why. In that structure, the absence of a documented case is not evidence that nothing has happened. It is evidence that nothing would be documented if it had. What is on the table, and has been for five years, is the Red Cross's two-tier proposal, draftable as it stands. Prohibit unpredictable systems — "autonomous weapon systems that are designed or used in a manner such that their effects cannot be sufficiently understood, predicted and explained." Prohibit anti-personnel autonomy — "autonomous weapon systems that are designed or used to apply force against persons." Regulate everything else through limits on target type, on duration, geographic scope and scale of use, and through "requirements for human–machine interaction, notably to ensure effective human supervision, and timely intervention and deactivation."350 Two prohibitions and a regulatory tier. Five years on the table.

XI.

There is exactly one place where a government has written down the sentence "these things are prohibited," and it is worth reading both for what it does and for where it stops. The European Union's artificial intelligence regulation contains a list of prohibited practices that has applied since the second of February, 2025.351 Among them, the one that speaks directly to the first half of this chapter: it prohibits the use of "an AI system for

making risk assessments of natural persons in order to assess or predict the risk of a natural person committing a criminal offence, based solely on the profiling of a natural person or on assessing their personality traits and characteristics." And then the carve-out: "this prohibition shall not apply to AI systems used to support the human assessment of the involvement of a person in a criminal activity, which is already based on objective and verifiable facts directly linked to a criminal activity."352 Profile-based prediction of who will offend: banned. Analytic support for an assessment already grounded in objective and verifiable facts about a specific criminal act: permitted. Whatever one thinks of the boundary, somebody drew it, in a statute, in advance. The same list prohibits creating or expanding facial recognition databases "through the untargeted scraping of facial images from the internet or CCTV footage," and prohibits inferring emotions in workplaces and educational institutions outside medical and safety uses.353

The American counterparts exist, and they are subnational. Maine prohibits state and local agencies from obtaining, retaining, accessing or using facial surveillance; requires logs of every permitted request; provides that "Facial surveillance data does not, without other evidence, establish probable cause justifying arrest, search or seizure"; and makes data collected in violation "unlawfully obtained," subject to deletion and "inadmissible in evidence in any proceeding."354 Virginia permits local law-enforcement use but bars real-time tracking of an identified individual in public, requires an accuracy score of at least ninety-eight percent true positives, and provides that "A match made through facial recognition technology shall not be included in an affidavit to establish probable cause."355 San Francisco has prohibited its departments from obtaining, retaining, accessing or using facial recognition technology since 2019.356 No federal statute exists. A 2023 Senate bill would have barred federal agencies from acquiring, possessing or using biometric surveillance absent specific congressional authorization; it was read twice, referred to committee, and never enacted.357 There is no enacted federal law restricting facial recognition or predictive policing. And now the hole. The European regulation does not apply to systems used exclusively for military, defence or national security purposes, and provides that it does not affect member state competences concerning national security.358 I have to flag the sourcing: I obtained the article text from a widely used verbatim reproduction rather than from the official journal, and it should be checked against the official text before anyone relies on it in a proceeding. The Council of Europe's framework convention on artificial intelligence, opened for signature in September of 2024, goes further. A party "shall not be required to apply this Convention to activities within the lifecycle of artificial intelligence systems related to the protection of its national security interests." And then: "Matters relating to national defence do not fall within the scope of this Convention."359 I could not confirm from a primary source how many states have ratified that convention or whether it has entered into force, and I would treat any confident secondhand claim about it with suspicion.360 Predictive policing of a European citizen by a European police force: prohibited since February of 2025. The same computation, performed by the same company, for a defence ministry: outside the regulation entirely.

This is not hypocrisy, and I am not going to call it that. National-security carve-outs are in almost every human-rights instrument ever drafted, for reasons that are not frivolous, and a treaty that tried to bind defence ministries would probably have no signatories at all. But the citizen should understand exactly what has been achieved: the world has produced its first prohibitions on automated determinations about persons, and it has drawn them so that the domain where a machine may select a human being for lethal force is the one domain they do not reach.

XII.

Objections, including the ones that damaged this chapter while I was writing it. The strongest program in my argument ended for the wrong reason. I used the Los Angeles case as evidence that audits work. The department's own memorandum says the prediction contract ended because of procurement cost and the city's pandemic-era financial crisis.298 The audit had been published thirteen months earlier and the program continued after it. My response is narrower than my original claim: audits do not reliably stop programs, but they are the reason we know what the programs were, and the Chicago decommissioning followed one directly.295 Transparency is a precondition for accountability, not a substitute for it. A third of the people were never scored, which is an argument against me. The immigration inspectorate found that officers conducted no risk assessment at all for thirty-three percent of detained noncitizens in one fiscal year and forty-three percent in the year before.311 Somebody can fairly say: your complaint is that a model with no privacy review and no disparity analysis is applied to people who are not told about it, but the audit shows the real failure is that it is not applied consistently. What I do not accept is the implied conclusion. The remedy for inconsistent use of an unreviewed instrument is not universal use of an unreviewed instrument. It is review — of the kind the same agency's own form records as not having been performed.306 I have no dead body for the weapons argument. No official finding anywhere establishes that an autonomous weapon has killed a human being without human intervention. The most-cited source does not say what it is said to say.348 My argument rests on structure rather than incident, and a reader who requires an incident should regard this section as unproven.

Two of my examples turned out not to be about algorithms at all. I had intended to use the Department of Justice's 2015 investigation of a Missouri municipality as an instance of automated enforcement, because it documents a court issuing over nine thousand arrest warrants in a single year over parking infractions, traffic tickets and code violations, and refusing to issue the compliance letters that would restore a driver's licence until the entire fine was paid.361 It is also not about an algorithm. The practices it describes are routinized and administrative, and I removed the example rather than let the reader infer a machine that was not there. Likewise, I could find no official finding about an algorithmic fines-and-fees or licence-suspension system, and no state audit that quantifies the false-positive rate of the Michigan fraud system, only the control failures.362 Three recent facial-recognition arrests are missing from this chapter. During 2025 and 2026 at least three further wrongful-arrest cases were reported — one in New York, one in St. Louis involving a commercial recognition service, one further case in Detroit. I could not obtain a primary court record for any of them. Their absence understates the current count rather than overstating it. And the tool I complain about sets people free. The federal risk instrument governs eligibility for earned time credits. A favourable score shortens custody. The Department published its own finding that the instrument overpredicts risk for Black, Hispanic and Asian people while continuing to use it for a purpose that releases people.317 A reform that abolished the tool without replacing the credit mechanism would keep people in prison. That is a real cost and I would not pay it. XIII. Three questions, the same three I have asked of every subject in this book, applied here where they bite hardest. Who can witness it? In the Woodruff record, no one could witness the step that mattered: the narrowing of seventy-three candidates to one is simply not in the file.280 In the immigration inventory, the person being scored is recorded as not interacting with the system, so no notice is owed and there is nothing to witness by design.308 Enforcement systems are the least observable and the most consequential objects in modern government, and that combination is not accidental — observability is what a machine-mediated process most easily dispenses with, because the machine does not need to be told what it did.

Who carries the burden if it was wrong? Porcha Woodruff carried it: arrested at eight months pregnant, charges dropped, suit lost on summary judgment, and an immunity doctrine holding that the officer had no fair warning because no court had ruled.279281 The four hundred thousand scored Chicagoans carried it, invisibly, for the two and a half years the scores could not move.293 Against that: no vendor in the United States faces statutory liability for an enforcement harm. Whenever the answer to what happens to the institution is nothing, the error rate will not improve. How is it returned? This is the question that decides why prohibition, and not mitigation, is the right instrument in this chapter alone. Everywhere else in this book I have argued for procedure: notice, comment, appeal, sunset, accounting. Procedure works when the harm is reversible. A day in a cell at eight months pregnant cannot be returned. A person killed by a munition that selected them cannot be returned. Liberty and life are the two goods in the whole of governance for which there is no remedial mechanism, only compensation, and compensation is not return. Where the harm cannot be undone, the correct instrument is a rule about what may not be built, not a rule about how carefully it must be operated. The mechanisms that do point toward return are small and procedural: the retrospective audit of every facial recognition request logged since 2017;289 the holding that a risk score may not determine whether a person is incarcerated.316 Each gives somebody a way to get something back. None of them can give back a life.

XIV.

What I would insist on. Five conditions, each with a name attached and a consequence for failure.

1. A named human of record, who decided, and who can be reached. For every determination that touches a person's liberty — arrest, detention, supervision level, custody classification, deportation, denial of a benefit that keeps a household housed — there must be a natural person identified in the file as having made the determination, with the authority to reverse it, and reachable by the person affected. Not the agency. Not "the system determined." A name. And the file must record what the automated component produced and what the human did with it, including, where the system returns a ranked list, the full list, the rank of the selected candidate, and the stated basis for the narrowing. If not: the determination is void and the evidence derived from it is inadmissible, on the model of the state statute that already provides exactly that remedy for unlawfully obtained facial surveillance data.354

2. No trade-secret objection inside a criminal process. Where an automated system has contributed to the identification, charging, detention, sentencing, or supervision of a person, that person is entitled to the system's documentation, its validation studies, its error rates disaggregated by demographic group, and the calculations applied to them personally — and no builder or user may assert trade secret or intellectual property protection to resist that disclosure. One American state enacted it in 2019 in a single clause.320 It costs the public nothing. It costs a vendor the ability to sell an instrument it cannot defend, which is the point. If not: the system's output is excluded, and the party that relied on it bears the cost of proceeding without it.

3. Independent pre-deployment audit and published, disaggregated error rates, with automatic suspension when the report is late. Before any predictive or identifying system is used in policing, courts, immigration, or benefits enforcement, an independent body evaluates it and publishes the evaluation. Thereafter, the operator publishes error rates broken down by demographic group at a stated interval. The federal risk instrument's own most recent published review found that it overpredicts risk by race and said the statutory duty to reduce that disparity had not yet been met methodologically, and then no successor review appeared for two years.318319

Attach the consequence and the review happens. If not: use suspends automatically on the day the report is late, and determinations made after that date are void.

4. Liability follows the harm through the supply chain, by statute. The maker of a system that produces an enforcement harm is answerable for it, jointly with the deploying agency, and neither may point at the other. A vendor that designed the adjudication logic was "entwined" with the state and could be sued.324 Make it statutory, and borrow the European evidentiary architecture: a duty to disclose relevant evidence once a plausible claim is presented, and a presumption of defect or causation where the claimant faces "excessive difficulties, in particular due to technical or scientific complexity."328 If not: the deploying agency is liable for the whole of the harm and may not seek contribution.

5. Prohibition, not mitigation, where a machine would select a human being for force. Two rules, adopting the two-tier structure that has been on the table since 2021: no system whose effects "cannot be sufficiently understood, predicted and explained" may be deployed, and no system "designed or used to apply force against persons" without a human decision may be deployed at all.350 Domestically, that means no autonomous discharge of force by any police, corrections, or border agency, with no urgency waiver. Internationally, it means the instrument the Secretary-General and the Red Cross have now asked for three times, and it means closing the carve-outs: an artificial-intelligence prohibition that exempts military and national-security use, and a human-rights convention that places national defence outside its scope by definition, do not reach the case the prohibition is for.358359 Who: the legislature, by statute, for domestic agencies; the review conference in November, for the international instrument; the senior officials named in the existing review directive, who currently have authority to approve such systems and no obligation to refuse any of them.334 If not: deployment is unlawful, the systems are not procurable with public funds, and the officials who approved them are personally accountable rather than institutionally shielded. Take the two that cannot be given up. Number one, the named human of record with the ranked list in the file — because it is the only requirement that repairs the specific failure this chapter opened with. The seventy-three-to-one step is simply not written down, anywhere, in any system I examined. It is the cheapest fix in this book and the one that would change the most cases. A form field and a signature. And number five, the prohibition on machine selection of human targets — because it concerns a harm no procedure can repair, and because the window in which it can be adopted before the systems are widespread is, in the Red Cross's own words, "rapidly shrinking."

XV.

Nobody in the room ever proposed that a machine should decide. What happened is that we built something that produced an output, and then we designed a workflow in which a human being disposed of that output at a rate that made independent judgment impossible, and then we described the arrangement, accurately in every particular and falsely as a whole, as human oversight. The machine did not take the decision. The decision evaporated. It was distributed across a model, a queue, a throughput target, and a person who had eleven seconds, and afterwards there was no one who had made it and nothing that recorded its having been made.

That is the accountability gap, and it is why this chapter is not about robots. The gap is not the absence of a human; it is the absence of a human whose judgment was actually exercised and actually recorded, and who can be found afterwards by the person the judgment was about. Capability is not permission. That a thing can be computed does not establish that it may be decided by computation, and the question of which decisions require a human being is not a technical question that engineers will answer as the systems improve. It is a

political question, answered by statute or by default, and answering it by default means answering it in favour of whoever has already deployed. Some decisions about a human being require another human being to make them, to defend them, and to answer for them. Not because humans are more accurate — often we are not. Because accountability is a relationship between persons, and it cannot be held by a thing that has no interests, no standing, and nothing to lose. When a machine is wrong about you, there is no one to be wrong. Which decisions about a person should never be made by a machine? Ask it about a specific system, in a specific building, and something useful happens: somebody has to name a decision, and then somebody has to say who makes it instead, and then their name goes in a file. That is all accountability has ever been. Not a principle — a name, a date, and a place where the person affected is entitled to stand. — D.M.D.

What Will Requires

1. A named human being with authority to explain, revise, and repair every consequential automated decision.

2. Timely notice and a real appeal before avoidable harm becomes irreversible.

3. Emergency authority that is limited, documented, reviewable, and temporary.

4. Public participation before contracts, permits, rules, and deployments are final.

5. Accessible records, open meetings, and plain-language explanations of public action.

6. Civic institutions designed for ordinary people with work, caregiving, disability, and limited time.

A Practice for the Reader

Identify one public body that decides something consequential in your community. Learn when it meets, where it publishes its agenda, and what authority it actually holds. Before its next meeting, choose one agenda item and ask: Who can change this decision? What record supports it? What would a real appeal look like?

Coda: The Power to Act Must Have a World to Protect

Will matters because people must be able to shape the conditions of their own lives. But action without attention to consequence can reproduce the very forms of domination it seeks to resist.

A public may have a hearing while its water is depleted. A worker may retain formal rights while the physical systems of housing, energy, and health become increasingly unlivable. Democracy depends on the material world within which people live.

The next question is therefore larger: What, exactly, are we acting to preserve? Technology is never only virtual. It has a body. It draws upon land, water, energy, labor, public infrastructure, and places where its costs are left behind.

NEXT: PART THREE — FORM

THE PHYSICAL BODY OF THE CLOUD

Energy, Water, Land, and the Geography of Who Pays

TECHNOLOGY MUST ENLARGE HUMAN AGENCY, NOT REPLACE IT.

We live in an age of extraordinary capability. But capability is not wisdom. A system does not become just because it is efficient. A prediction does not become legitimate because it is accurate. A decision does not become democratic because it is automated.

Artificial intelligence is becoming part of the conditions under which people learn, work, receive care, seek housing, obtain benefits, communicate, and participate in public life. It is also built into the physical world: land, water, energy, infrastructure, labor, and communities asked to absorb costs in the name of progress.

This foundational manifesto establishes the six conditions of freedom: Perception, Will, Form, Exchange, Threshold, and What Citizens Do. Together, they form a civic blueprint for truth, democratic agency, ecological stewardship, economic fairness, human dignity, and accountable public action.

THE SIX CONDITIONS OF FREEDOM

PERCEPTION • WILL • FORM • EXCHANGE • THRESHOLD • WHAT CITIZENS DO

"A workable Earth is an Earth on which ordinary people can breathe clean air, drink safe water, find meaningful work, participate in decisions that affect their lives, and pass forward a world no less livable than the one they inherited."

