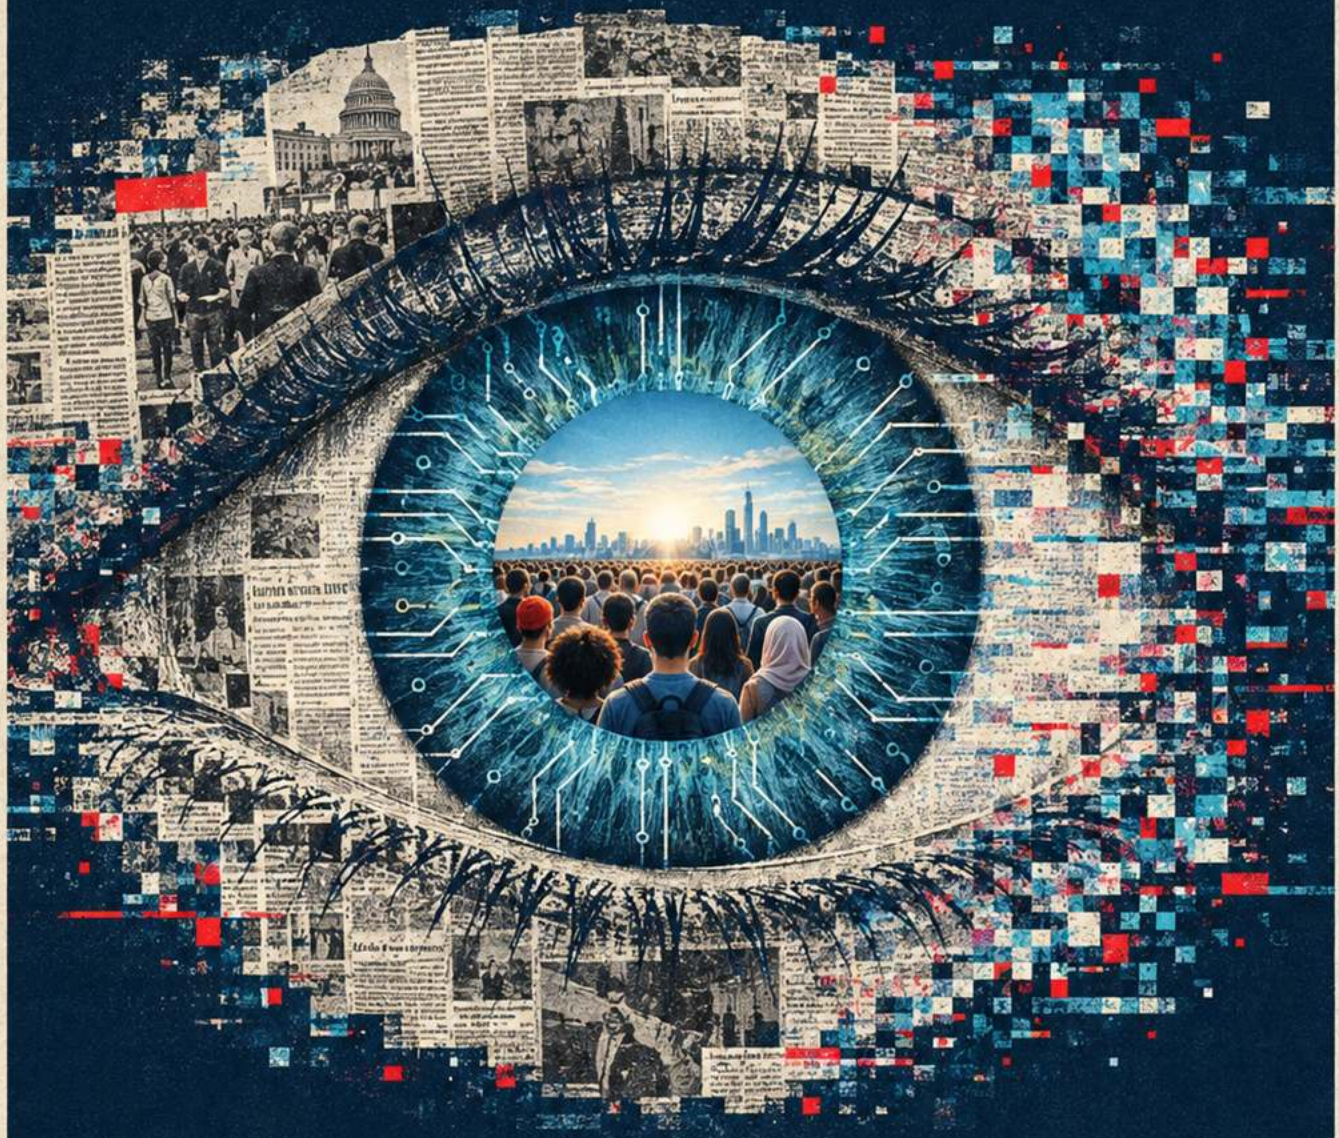
A CITIZEN'S MANIFESTO FOR THE AGE OF AI

PART ONE — PERCEPTION



THE INFORMATION COMMONS

Truth, attention, and the machinery of synthetic consensus

A CITIZEN'S MANIFESTO FOR THE AGE OF AI

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

PART ONE — PERCEPTION

THE INFORMATION COMMONS

Truth, Attention, and the Machinery of Synthetic Consensus

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.

This volume begins with Perception because no other freedom can be exercised without a public capacity to establish what is happening. A people cannot govern what it cannot see, consent to what it cannot understand, contest what it cannot verify, or learn from what it cannot remember.

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 One

The information commons is the public capacity to find the record, test the evidence, identify the source, preserve the account, and distinguish what is known from what is merely claimed.

The chapters in this book ask four questions of anything that claims to inform you: What is the record? Who holds it? Will it be there next year? And can I show it to somebody who doubts me?

Perception is sound to the extent that a citizen who does not have to explain why can obtain the record, check it against something outside the party that produced it, and still find it there a year later.

PART ONE — PERCEPTION

Keeping communication, inquiry, and memory open

Its corruption: surveillance, saturation, and the loss of a shared factual floor.

Contents

Chapter One — The Information Commons: Truth, Attention, and the Machinery of Synthetic Consensus

Chapter Two — Education: Learning in the Presence of a Machine That Will Do the Work for You

Chapter Three — Memory and the Record: Who Keeps History, and Whether It Stays Correctable

Part 1 Opening

Perception is first because it is first in time. Nothing else in this manifesto is available to a person who cannot establish what is being done to them.

A society cannot consent to what it cannot see. It cannot contest what it cannot verify. And it cannot learn from what it cannot remember.

We now live in a condition in which the volume of the record rises while its evidentiary value can fall; where a smooth answer may replace the document that would allow a citizen to check it; and where surveillance and saturation produce the same result—a public unable to establish what happened.

This part does not ask readers to become permanently suspicious of everything. Skepticism alone can become another form of helplessness. It asks for a stronger civic discipline: preserve the record, identify the witness, verify the claim, and keep the evidence available for the next person who needs it.

The right to see is not a luxury for specialists. It is the beginning of citizenship.

On the fourth of July, 1966, a statute was approved whose stated object was "to clarify and protect the right of the public to information." It took effect one year later, on another fourth of July. Its central move was not the disclosure it compelled but the presumption it reversed: a citizen asking a federal agency for a record no longer had to establish standing, demonstrate a proper purpose, or explain the interest served. The burden shifted to the government to say why not.6 Fifty-nine years later, in the fiscal year just reported, the federal government received 1,707,197 requests for records under that statute — a record, and a 13.7 percent increase over the year before. Agencies processed a record 1,635,055.7 Whatever else is failing, the appetite to look is not. Then look at the queue. At the close of the same fiscal year, 339,671 requests were sitting unanswered across the government — 27 percent more than the year before.

The backlog of administrative appeals, the remedy for a wrongful denial, stood at 11,652, an increase of 116.5 percent in one year.8 More people than ever are asking to see, and the line to see is lengthening faster than the asking. Perception is first in this book because it is first in time. Nothing else in these eighteen chapters is available to a person who cannot establish what is being done to them. A society cannot consent to what it cannot see. It cannot contest what it cannot verify. And it cannot learn from what it cannot remember. Those three sentences are the three chapters that follow: the information commons, the schooling that produces people able to use it, and the record that has to survive long enough to be checked. I am not making a romantic argument about transparency. Perception is not sufficient. A society can be extremely well informed and entirely unable to act — that is the failure Part Two exists to address.

But perception is prior. Every remedy in this book presumes that somebody can find out what happened. Here is what is new, and it is not censorship. The record is dissolving on its own. Of a random sample of just under a million webpages drawn from each year between 2013 and 2023, a quarter were no longer reachable as of October 2023. Of the pages sampled from 2013, 38 percent were gone a decade later. Even among pages captured in 2021, about one in five had vanished within two years. 9 Nobody decided this. There was no order, no suppression, no official who could be named and asked. Pages were deleted in the ordinary course of housekeeping, domains lapsed, a company reorganized its site, a contract ended. The same study sampled roughly five hundred thousand pages from government websites across federal, state and local levels: 21 percent contained at least one link that no longer resolved.10 A statutory right of access to a document that has ceased to exist is not a diminished right.

It is a formality. And the people who used to do the looking on your behalf are being subtracted. In its October 2025 count, Northwestern's Medill School put the number of American counties with no local news source at 213, up from 206 the year before. In another 1,524 counties, one source remains. Taken together, roughly 50 million Americans have limited or no access to local news.11 Newspaper closures over the year came to 136 — better than two a week. Twenty years ago there were about 150 such counties and about 37 million people living in them.12 The statute I opened with confers a right that somebody has to exercise. It is executed by a person who notices that a permit was filed, who knows which agency holds the file, who has the standing habit of asking, and who has somewhere to publish the answer. In 213 counties there is now no such person at all, and in 1,524 there is one.

So why does any of this belong in a book about artificial intelligence? Because three things change at once, and each attacks a different part of the same capacity. The volume of the record rises while its evidentiary value falls. When producing fluent, confident, plausible text costs nothing, the existence of a document stops being evidence that anyone did the work the document describes. The pile grows and the pile means less. The

intermediary stops showing you the record. A citizen who asks a question and receives a smooth paragraph in reply has been served — but has not been shown anything. The step where you could check the claim against the underlying document becomes optional, then unusual, then unavailable, and the loss is invisible because the answer got easier. And the two corruptions named at the head of this part turn out to be the same corruption.

Surveillance watches you too closely; saturation drowns you in material. Both end in a citizen who cannot establish what happened — one because the observation runs only in one direction, the other because everything is observable and nothing is verifiable. A person who cannot establish what happened cannot consent, cannot object on the record, and cannot prove later that they objected.

The chapters that follow ask four questions of anything that claims to inform you: What is the record? Who holds it? Will it be there next year? And can I show it to somebody who doubts me? Here is the standard I would hold a society's perception to, stated as one test: perception is sound to the extent that a citizen who does not have to explain why can obtain the record, check it against something outside the party that produced it, and still find it there a year later.

Obtainable. Independently checkable. Durable. An access right with a three-hundred-thousand-request queue fails the first. An answer generated by the same system whose conduct is in question fails the second. A citation to a page that will not resolve in eighteen months fails the third. None of the three is a technology problem. Each is a decision somebody made, or failed to make, and can be made differently. That is what the three chapters are for. — D.M.D.

Chapter One: The Information Commons

Truth, Attention, and the Machinery of Synthetic Consensus

How do I know what is true when producing convincing falsehood has become free, and the people whose job was to check have been laid off? I. The room, and who was in it On August 12, 2026, the Louisiana Public Service Commission voted three to one to let Meta keep key records about its Richland Parish data center out of public view, reversing an administrative law judge's ruling in Docket U-37425.13 I know that because two newsrooms were paying attention: a public radio station and a nonprofit statehouse outlet. Not a wire service. Not a network. Not an algorithmically assembled feed. Reporters who read dockets. Hold that fact in your hand for a moment, because it is the whole chapter. A decision that will shape the electric bills, the water table, and the tax base of a region for thirty years was made in a room that could have held forty people. The number of professional journalists in the United States whose job includes reading utility dockets is small and shrinking. The number of synthetic websites capable of publishing a confident, fluent, entirely fabricated account of that same meeting within ninety seconds is, as of this writing, at least 3,749, in sixteen languages.17 Consider what verification actually costs. The Richland Parish matter, at the point I looked at it, involved a utility's request to build generating capacity for a single private customer, an intervenor motion for investigation, an administrative law judge's ruling on confidentiality, a commission vote reversing it, competing figures for the monthly effect on residential bills, and a registered water withdrawal that differed from the expected withdrawal by an order of magnitude. To hold an informed opinion, a citizen must locate the docket, understand that a "certificate of public convenience and necessity" is the instrument that matters, read filings written for lawyers by lawyers, distinguish a registered maximum from an expected average, and notice which of the numbers in a press release is a projection. That is a week of unpaid evenings for someone with a full-time job. A reporter on a beat does it in a day, because she has done it forty times and knows which paragraph to read first. That is what a newsroom is: not a moral authority, but accumulated procedural fluency, paid for. When it goes, the knowledge does not transfer to the public. It evaporates. And the vacuum it leaves is not neutral space; it is immediately filled, because filling it is now nearly free, and because the parties with the most at stake in the docket are the parties best funded to fill it.

There is also a quieter loss. A local reporter is a person who can be recognized in a grocery store by someone whose house is next to the substation. That proximity is a form of accountability that no distributed system replicates. A synthetic account of a public meeting cannot be asked, at the hardware store, why it left something out. I want to be precise about what I am not arguing. I am not arguing that people have become stupid, or that they were once wise and are now degraded. I am not arguing that a golden age of shared truth existed and has been stolen. The American information commons was always contested, often captured, frequently racist in what it chose to cover and whom it chose to believe. I am arguing something narrower and more mechanical: the institutions that once absorbed the cost of verification are being dismantled at the same moment that the cost of fabrication has fallen to approximately zero, and no one has been made responsible for the gap. That is not a cultural problem. It is a governance problem, and governance problems have addresses. II. The compounding risk Every other chapter depends on this one. If a community cannot establish what is being proposed, at what cost, on whose land, with what alternatives, then education policy, energy policy, labor policy, and human rights become matters of assertion rather than deliberation. A citizen who cannot verify cannot

consent, and a citizen who cannot consent is a subject with extra steps. The corruption of the information commons does not merely add a risk to the list; it removes the instrument by which every other risk could have been measured. III. What is actually happening The marginal cost of persuasion has gone to zero. For most of the modern era, producing something that looked like journalism required capital: presses, transmitters, salaries, an office, a lawyer. That capital requirement functioned, imperfectly and often unjustly, as a filter. It is gone. NewsGuard's tracking center counted 3,749 sites it classifies as AI content farms — publishing at scale with little or no human oversight — across sixteen languages, as of its June 2026 update, up from roughly 2,000 eight months earlier.17 The significant fact is the trajectory, and the fact that the count is maintained by a private company because no public body maintains one.

Fabrication is now cheaper than reporting, and enforcement is slower than both. In January 2024, an AI-cloned voice of a sitting president told New Hampshire Democrats not to vote in their primary. The Federal Communications Commission proposed and then imposed a $6,000,000 forfeiture against the political operative who commissioned the calls, and settled with the transmitting carrier for a $1,000,000 civil penalty plus a compliance plan.18 Then a Belknap County jury acquitted him on all counts in June 2025.19 I raise this not to relitigate a verdict I did not sit through, but because it is the clearest available demonstration of the enforcement gap: the technical act took an afternoon and a few hundred dollars, the regulatory response took eight months and produced a fine, the criminal response took seventeen months and produced nothing, and a related civil judgment produced an award of $22,500 that the defendant publicly refused to pay.19 If you are deciding whether to try this in the next election cycle, that is the arithmetic you are looking at. Manufactured consensus is an old crime with a new engine. Before any of this was automated, the Federal Communications Commission's 2017 net-neutrality proceeding received more than 22 million public comments, of which the New York Attorney General later confirmed nearly 18 million were fake: more than 8.5 million impersonating real people, 9.3 million using fabricated identities, plus more than half a million forged letters to Congress. Three lead-generation firms paid more than $4.4 million collectively in penalties and disgorgement.20 That operation required buying identities, running campaigns, and coordinating vendors. It cost money and left a paper trail, which is how it was caught. The same result can now be produced by one person with a laptop, in fluent and individually varied prose, at a marginal cost near zero and with no vendor to subpoena. If the channel becomes computationally indistinguishable from a flood, agencies will discount all of it, and the honest commenter loses standing along with the fraudulent one. The abuse does not merely add noise, it destroys the value of the signal.

The rules are a patchwork, and the federal center is empty. Somewhere between thirty and thirty-three states now have laws addressing synthetic media in elections, depending on whether you count a statute struck down in court and one that sunset — Public Citizen's tracker, the most complete public accounting I could find, says thirty-three in its text and lists thirty-four rows in its table.21 Federally: the Federal Election Commission declined to open a rulemaking in September 2024, adopting instead an interpretive rule stating that its existing fraudulent-misrepresentation regulation applies regardless of the technology used. The FCC proposed disclosure rules for AI-generated content in broadcast political advertising in July 2024 and, as of this writing, has not adopted them.23 The European Union's Article 50 transparency obligations became applicable on August 2, 2026.27 So the operative situation for an American voter is this: the disclosure regime protecting you is either a state law that may not survive review, or a European regulation you have no standing to invoke. The newsroom that used to attend the hearing is gone. Medill's 2025 census of local news counted 213 counties with no local news source at all, 1,524 counties with only one, 136 newspaper closures in a single year, and roughly 50 million Americans with limited or no access to local news — against about 150 news-desert counties and 37 million

affected residents twenty years earlier.13 Newspaper employment fell another seven percent in that year alone, and more than three-quarters of newspaper jobs have disappeared since 2005.13 The net number of newspapers lost is genuinely contested; reputable accounts of the same report give figures from roughly 2,800 to nearly 3,500, depending on whether one counts gross closures or net change, and Medill's own release does not state the number.14 Where I cannot resolve a discrepancy, I will say so rather than pick the larger number. The measure that matters more than titles is bodies. The Local Journalist Index puts the national average at 7.8 local journalist equivalents per 100,000 residents in its 2026 edition, down from 8.2 the prior year and from roughly 40 per 100,000 in 2002 — a decline on the order of eighty percent.15 Louisiana sits below that national average.15 The outlets I could verify producing original reporting on Public Service Commission dockets — the proceedings that decide who pays for the data centers — are a nonprofit statehouse newsroom, a New Orleans nonprofit, a national nonprofit energy newsroom that co-publishes with both, a public radio station, and one metro daily in the northern part of the state.16 That is the entire apparatus standing between a multi-gigawatt buildout and an unexamined record. Public information infrastructure is being withdrawn at the same time. Federal funding for public broadcasting was rescinded in July 2025 — more than a billion dollars, including advance appropriations that had been approved on a bipartisan basis since 1976 — and the Corporation for Public Broadcasting's board voted in January 2026 to dissolve the organization.31 The attempted elimination of the Institute of Museum and Library Services was blocked, but only through litigation: a temporary restraining order, a preliminary injunction, a permanent injunction in November 2025, and a binding settlement in April 2026.32 In rural counties, the public radio station and the public library are frequently the last two institutions that hold a verified local record. One of them was defunded. The other survived because it was sued for, not because it was protected. The ranking machinery optimizes for engagement, which is not comprehension. This is often asserted; it has now been tested. In a pre-registered field experiment that re-ranked participants' real feeds, engagement-based ranking amplified emotionally charged, partisan, and out-group-hostile political content relative to a chronological baseline, made users feel worse about their political opponents, and did not match users' own stated preferences even as it successfully maximized their revealed clicks.24 A separate study in Science found that changing feed composition measurably changes partisan animosity.34 The system is not giving people what they want. It is giving them what they cannot look away from, and it is doing it to material about who should govern. The liar's dividend is real, and it is worst in exactly the medium that governance runs on. Five pre-registered experiments with more than fifteen thousand American adults found that politicians who falsely dismissed true reports of scandal as "fake news" or fabrications gained support — pooled effects around a fifth of a standard deviation — and that this tactic outperformed both denial and apology. The effect was strong for text-based scandals and largely disappeared for video.25 The domain where false denial works best is documents: memos, filings, transcripts, contracts, dockets. The material of self-government is the material most vulnerable to the claim that it was faked. Provenance exists as a standard and not as a duty. The Coalition for Content Provenance and Authenticity has a mature technical specification, version 2.4, dated April 2026, with a conformance program and broad industry membership.26 The corresponding international standard, ISO/CD 22144, remains under development and unpublished; anyone telling you Content Credentials "is an ISO standard" is ahead of the record.26 The EU's Article 50 will require machine-readable marking of synthetic output; the United States requires nothing. We have the seatbelt. We did not pass the law.

The window into the machinery was closed. CrowdTangle, the tool researchers and journalists used to observe how content actually spread on the largest platforms, was shut down in August 2024 and replaced by a library open to vetted academics through a university consortium and largely closed to working journalists.29 The

European Union has since created a legal right of researcher access — a delegated regulation adopted July 1, 2025 and published in the Official Journal that October — and researchers still report refusals, including a platform denial in October 2025 on the grounds that the requesting team had not adequately proved its credentials; a survey of specialists rated the overall data-access experience six out of ten, an improvement from 4.6 the year before.30 In July 2026 the Commission accepted a remediation plan from one platform requiring it to provide effective free access, stop contractually forbidding the collection of public data, and submit to an independent audit.30 Every bit of that leverage is European. An American researcher studying an American election on an American platform has no equivalent instrument. And the machine is now learning from its own exhaust. A 2024 Nature paper described "model collapse": train successive generative models on data produced by earlier models and the distribution's tails vanish, the outputs converge toward a low-variance mush.28 That paper received a formal author correction in 2025, and its inevitability claim is contested by peer-reviewed conference work arguing that collapse is an artifact of replacing real data rather than accumulating it alongside synthetic data; another paper documents eight incompatible definitions of the term in the literature.28 I am reporting a live scientific dispute, not a settled finding. But the governance implication survives the dispute intact: the value of authenticated human-generated records — archives, court transcripts, local newspapers of record, library holdings — goes up, not down, precisely when we are defunding the institutions that hold them.

IV. Concerns

V. The cumulative issue

Taken one at a time, each of these is arguable. A layoff is a business decision. A shuttered research tool is a product decision. A declined rulemaking is a jurisdictional judgment. A rescission is a budget priority. A voluntary standard is industry self-governance working as intended.

Taken together, they describe a single condition: the cost of producing a convincing false account of a public proceeding has collapsed, while every institution that existed to produce a verified account of that same proceeding has been defunded, laid off, litigated against, locked out, or left voluntary. No one decided this. That is what makes it dangerous. There was no meeting at which the American public deliberated whether to trade its verification capacity for cheaper content, and no body currently holds a mandate to notice the aggregate. The utility commission reviews the rate case. The state legislature considers the deepfake bill. The appropriations committee handles the rescission. The platform manages the API. Each is doing its job, and the sum of their jobs is a country that can be told anything about a hearing no one attended. This is the same structure as the load on the grid in Chapter Seven and the water in the aquifer: individually permissible decisions that together exceed a limit no one is charged with monitoring. The information commons is the first case in this book because it determines whether the others can be argued at all.

VI. Objections I take seriously

"You are describing a speech problem, and the cure is worse than the disease." Provenance marking and synthetic-media disclosure are not restrictions on what may be said. They are requirements to state how a thing was made, of the same species as a byline, a dateline, a nutrition label, or the "paid for by" line that has survived a century of First Amendment scrutiny. I would hold the line at process, not content: no public body decides

what is true; it decides only that the maker of an artifact must not conceal that a machine made it. Where a proposal starts sorting claims into permitted and forbidden, I am against it, and I would rather have a badly informed public than a licensed one. "Local newspapers were never the neutral watchdogs you are mourning." Correct, and the record is worse than that. American local papers ran lynching apologetics, redlining boosterism, and decades of company-town stenography. I am not asking for their restoration as moral authorities. I am asking for the restoration of a much more modest function: a paid person, physically present, taking notes that can later be checked against a transcript. That function is compatible with the newspapers having been wrong about a great deal, and its absence is not corrected by anyone's good intentions.

"Provenance mandates will entrench the largest companies." This is the strongest objection, and it is partly right. Any compliance regime advantages whoever can afford compliance. Two answers. First, the specification in question is an open one, developed in a multi-stakeholder coalition, not a proprietary format licensed by an incumbent, and the duty should attach at the point of generation and distribution rather than to every downstream user of a file.26 Second, the alternative to a rule is not an open field; it is the present arrangement, in which the largest firms already run private provenance and moderation systems on terms they set alone and change without notice. Unwritten rules do not disadvantage the powerful. "AI will expand local coverage, not shrink it." It might. A model that reads every filing in a rate case and flags the paragraph where the cost allocation changed is worth a great deal to an understaffed newsroom, and outlets are already using tools like these. But note what that requires: a newsroom, a person to direct the tool, and a person willing to sign the resulting story. The tools multiply capacity where capacity exists. Where the last reporter has been laid off, there is nothing to multiply, and a machine that generates plausible local coverage without a local witness is not a substitute for the witness — it is the most efficient possible way to fill the vacuum the witness left. "This is all just the internet; every generation panics." Some of it is. Moral panic about new media is a genre, and I have no interest in joining it. So I am trying to restrict myself to claims that carry a number and a source: how many counties have no news source, how many journalists per hundred thousand residents, how many synthetic news sites, how large a fine, whose access was revoked on what date, which agency declined to write which rule. Where the numbers are contested, this chapter says so. If those numbers are wrong, the argument should fail, and I would rather be corrected in public than be persuasive.

VII. The three tests

I said in the introduction that I would apply three questions to every deployment of consequence. Here they are, applied. The first test — who is watching, and can they still see? Not whether oversight formally exists, but whether any person or institution is actually positioned to observe and report. In 213 counties, no one. In 1,524 counties, one outlet, unbacked.13 For platform distribution, no journalist and only vetted academics, through a consortium, with refusals still common.29 30 For the Richland Parish records, an administrative law judge said the public could see them and the commission said it could not.33 This test is failing at the center, and it is failing quietly, because the absence of a watcher generates no headline by definition.

The second test — who carries the burden of proof, and who carries the cost of being wrong? Today the citizen carries both. You are asked to authenticate a video, evaluate a claim about a docket, and detect a synthetic voice, using unpaid time, against systems built by organizations with far more resources than you have. The burden belongs on the party with the capability: the maker of a synthetic artifact should mark it, the deployer of an automated ranking system should disclose its objective, the utility seeking a rate increase should publish its

assumptions. The third test — what returns to the place it was taken from? An attention economy extracts hours, data, civic bandwidth from a community and returns advertising. A data center extracts power, water, land, and tax abatement from a parish and returns some jobs and a payment schedule. A reasonable question is what portion of the value comes back as capacity: not sentiment, not a donation, but a reporter who can attend a meeting, a librarian who can hold a record, an archive that survives a change of administration. Right now the return on the information commons is close to zero, and in the public-media case it is negative.

VIII. What I would insist on

These are floors, not aspirations. Each names who must act. Mark what the machine made. Any system that generates or materially alters audio, image, video, or text for public distribution attaches durable, machine-readable provenance, using the existing open specification rather than a proprietary substitute. Who: the model provider and the distributing platform, by statute. If not: the content may not be monetized or algorithmically promoted, and the platform loses any safe-harbor defense as to that item. Disclose synthetic political communication, federally, before the next cycle. Voice, likeness, or fabricated depiction of a candidate or official in paid or broadcast political communication carries a clear, contemporaneous disclosure. Who: Congress, and the FCC and FEC within existing authority. If not: the communication is an unlawful expenditure, with penalties that scale to the spend rather than to a fixed schedule that large actors treat as a fee. Make the penalty faster than the offense. Statutory expedited injunctive relief for demonstrably fabricated depictions of identifiable persons in the final sixty days before an election, with a preservation duty and a public docket. Who: state legislatures and Congress. If not: the remedy arrives after the vote, which is the same as no remedy.

Publish the objective function. Any platform above a defined scale publicly states what its recommender optimizes for, discloses material changes, and submits to independent audit of the gap between stated and revealed preference. Who: the platform, verified by auditors with no financial interest in the finding. If not: no eligibility for public contracts, and the disclosure duty is enforced as a consumer-protection matter. Guarantee researcher and journalist access by law, not by permission. A statutory right of vetted access to public-content distribution data for accredited researchers and working journalists, with a defined response window, an appeal, and a prohibition on contract terms forbidding the collection of public data. Who: Congress; state attorneys general as to deceptive-practice claims. If not: refusal is itself the violation. Fund the watcher. Local news and public media are treated as civic infrastructure, funded through mechanisms structurally insulated from the officials being covered — refundable subscription and payroll credits, competitive public-media appropriations restored on a multi-year advance basis, and library and archive funding protected from single-administration reversal. Who: Congress, state legislatures, and municipal governments. If not: accept that most local decisions will be made unobserved, and stop describing the result as public consent. Preserve the authenticated human record. Statutory funding and legal protection for public archives, court records, library holdings, and newspapers of record, including mandatory deposit of digital public records with independent institutions. Who: state archivists, Congress, and the courts. If not: the only durable copy of the record is held by the party with an interest in its contents. No secret record where the public pays the bill. Where a project's costs are recovered from ratepayers or taxpayers, the load, water withdrawal, contract terms, and cost allocation are public. Commercial sensitivity may justify redaction of a specific figure; it does not justify sealing the category. Who: utility commissions and permitting bodies. If not: the application is incomplete and may not be approved. Requirements one through five feed the disclosure and independent-audit non-negotiables in the manifesto that

closes this book; six and seven are the affirmative obligation that makes the rest enforceable; eight is the cumulative-review principle applied to information rather than to megawatts.

IX. The stance

I am not asking for a return to a media environment that never existed. I am asking for something narrower: that we stop treating the capacity to know what our own governments are doing as a byproduct of an advertising market.

The machines are not the adversary. A model that summarizes a two-thousand-page rate filing accurately is a gift to a citizen who has a job and three hours a week. I use these tools daily; this book was written with them and says so. What I refuse is the arrangement in which the same technical capability that could equip a citizen is deployed instead to flood the channel, while the institutions that would have taught that citizen what to check are dismantled for reasons entirely unrelated to technology and entirely convenient to it. So the instruction at the end of this chapter is not "be skeptical." Skepticism, unaccompanied, produces the liar's dividend — a population that disbelieves everything is precisely the population in which a true document can be waved away as a fabrication.25 The instruction is: fund and use a watcher. Subscribe to the outlet that reads the dockets in your state; if there is none, that absence is your first civic fact. Learn the name of the body that approves the thing being built near you, and read one filing. Keep a copy. The commission voted three to one on a Wednesday in August, and it is only in this book because somebody was in the room. — D.M.D.

Chapter Two: Education

Learning in the Presence of a Machine That Will Do the Work for You

Learning in the Presence of a Machine That Will Do the Work for You

If a machine can write the essay, what exactly was the essay for? I. I want to begin with an admission, because the rest of this chapter is not honest without it. Parts of this book were drafted with the help of a machine. I gave it instructions. It produced paragraphs. I cut most of them, rewrote the rest, chased down every number myself, and threw away whole sections that read fluently and meant nothing. The final text is mine in the sense that matters — I chose every claim in it and I will answer for every one — but I did not produce it the way I produced the papers I wrote at nineteen, alone at a desk at two in the morning, sick with the knowledge that no one was coming to help. Here is what I noticed. The machine was useful to me in exact proportion to what I already knew. Where I had judgment, it accelerated me. Where I did not, it produced confident, well-formed, plausible garbage, and I could not always tell. The ability to evaluate its output came from the two-in-the-morning years. It came from the friction. Which raises the question this chapter exists to answer: if the friction is what built the capacity to use the tool safely, and the tool removes the friction, what happens to the child who gets the tool first? That is not a rhetorical flourish. It is an operational question that roughly fifty million American families are now answering by default, one homework assignment at a time. A ninth-grade teacher assigns five paragraphs on why the Roman Republic fell. Nobody involved believes the world needs another fourteen-year-old's account of the fall of the Roman Republic. The essay is not the product. The essay is the instrument. It is a device for forcing a mind to do eight specific things it will not do on its own: to select from more material than fits; to notice that two sources disagree; to decide which one to believe and to be able to say why; to arrange causes in an order that implies a claim; to discover in the third paragraph that the claim in the first paragraph was wrong; to fix it; to feel the difference between a sentence that says something and a sentence that merely sounds like it does; and to finish, and to own the thing that is finished. The essay was for the eight things. A machine can now produce the artifact without any of them occurring in the student. And the assessment regime in almost every American school still grades the artifact.

That is the whole problem. Everything else in this chapter is consequence. Now assemble the pressures, each reasonable on its own terms. A district is under budget pressure and has been for years. A vendor arrives with a product that promises personalization at scale — the thing every underfunded district has wanted since the invention of the thirty-child classroom. Students already have the tools, on their phones, faster than the adults around them. Teachers, who are drowning, discover that the same tools give back several hours a week. Nobody has trained them. The assessment system still measures a product the machine can now generate. Parents cannot see what is happening. The state has issued guidance that is voluntary. The federal agency that would have written the rules has been cut roughly in half, and the one federal document that told schools specifically how to avoid hurting children with these systems has been formally withdrawn. Take each on its own and I can defend it in a sentence. Assemble them and you get an outcome no one selected: a generation fluent in prompting and unpracticed in thinking, sorted early by systems their parents cannot inspect, on records that will outlive their childhoods. I do not think anyone in this chain is a villain. I think the chain does not have anyone in it whose job is to look at the whole thing. III. Before the argument about machines, the ground they landed on. American reading scores were falling before any student ever opened a chatbot. On the 2024 National Assessment of Educational Progress, average reading scores for national public school students were 214.3 in

grade four and 256.7 in grade eight, down from 219.4 and 262.0 in 2019.35 The more useful number is not the average but the floor: the share of students scoring below NAEP Basic — meaning they cannot reliably do the partial, fundamental things the assessment describes — rose in grade four reading from 34.8 percent in 2019 to 40.7 percent in 2024, and in grade eight reading from 28.5 percent to 34.2 percent.36 The National Assessment Governing Board described roughly four in ten fourth graders below Basic in reading, the largest share since 2002, and about a third of eighth graders, the largest ever recorded.

The most recent federal release, in June 2026, offered the first genuinely mixed news in years: nine-year-olds gained four points in both reading and mathematics against 2022, with the gains concentrated at the tenth and twenty-fifth percentiles — among the students who were furthest behind. Thirteen-year-olds did not move at all, and both ages remain below where they were in 2020.38 In the same release: 14 percent of students said they read for fun almost every day, down three points from 2020 and thirteen points from 2012.39 And the children are not in the building. Chronic absenteeism — missing at least ten percent of the school year — rose from 14.8 percent in 2018–19 to 28.3 percent in 2021–22, a 13.5-point jump amounting to roughly six and a half million additional students, across forty states and the District of Columbia covering 93 percent of American students.40 The federal data series stops at 28 percent for 2022–23.41 A think-tank tracker compiling state releases since then puts the figure at 23.5 percent for 2024 — better, and still half again the pre-pandemic rate.42 I flag the source because it matters: there is no current federal number. We are navigating on a private tracker. Hold those figures next to the promises being made for machine tutoring, and notice what the promises are being asked to fix: not a system operating at 95 percent that needs a boost, but one where four in ten nine-year-olds cannot read at the assessment's basic level and one child in four is chronically not there. One state moved. Louisiana's fourth-grade NAEP reading average went from 209.9 in 2019 to 212.3 in 2022 to 216.0 in 2024 — above the national average, and up from fiftieth in the country to sixteenth.43 The state's department of education reported it was one of only two states whose fourth graders exceeded their 2019 scores in both reading and mathematics, and that it led the nation in fourth-grade reading growth across two consecutive cycles.44 In July 2026 the state reported its highest-ever statewide results on its own assessments: 36 percent of students in grades three through twelve at Mastery or above, with performance improving or holding in twenty-eight of thirty-one tested grades and courses.45 I live an hour from that state line, and I have watched this happen with something close to disbelief. Here is what did it: phonics. Screening every child early. Tutoring the ones who screen low. Training teachers in a specific method and then checking whether they used it. Retaining third graders who could not read, which was politically expensive. Telling parents the truth about their own child's scores.

Not one element of that was technological. Not one element was novel. It was a decade of adults doing unglamorous, legible, human things, in public, and being accountable for the result. I want that on the record before a single vendor is allowed to claim the credit, because the one documented turnaround in American reading in my lifetime was achieved by the cheapest and least exciting means available. I note one gap honestly: I could not confirm from a primary legislative source the bill number of the 2019 Louisiana literacy legislation that anchors this sequence, so I describe the reform without it rather than cite a number I cannot verify. IV. What is actually happening now, mechanism by mechanism. Adoption preceded pedagogy, and it is not close. In a Pew Research Center survey fielded in late 2025, 64 percent of American teens aged thirteen to seventeen reported using AI chatbots, and 54 percent had used them for help with schoolwork; 10 percent said all or most of their schoolwork involved chatbot help, and another 21 percent said some.46 On the teacher side, a Gallup survey of 2,232 public K-12 teachers fielded in spring 2025 found 60 percent had used an AI tool for work that

school year and 32 percent used one weekly or more.47 And the number that tells you what kind of system this is: 18 percent of teachers reported receiving any formal guidance from administrators on AI use. Across ten AI-related tasks measured, 34 percent had received no guidance at all and 48 percent only informal guidance; 69 percent had no guidance on using AI for one-on-one instruction or tutoring, and 58 percent none on grading and feedback.48 The tool is in the hands of the majority of students and a majority of teachers, and four in five teachers have been told nothing. Whatever is happening in American classrooms right now is not a policy. It is a weather pattern. I should note, because I promised to be careful with numbers: Pew's earlier finding that 26 percent of teens had used ChatGPT for schoolwork asked about one product in 2024, while the 54 percent figure asks about any chatbot in 2025.49 Those are two different questions and I will not draw a line through them, however tempting the slope. The research does not say what either camp claims it says.

In a study of roughly a thousand high school mathematics students published in PNAS, students given access to a standard chatbot during practice improved their practice performance by 48 percent, and students given a version with tutoring safeguards improved by 127 percent. Then the access was removed for the exam. The students who had used the unguarded chatbot scored 17 percent worse than students who never had access at all. The safeguarded version largely eliminated that harm.50 The tool did not have a fixed effect. Its design determined the sign of the effect. At the other end: a randomized crossover trial in a Harvard introductory physics course, published in Scientific Reports, found that students working with a purpose-built AI tutor showed median learning gains more than double those of the same students in well-run in-class active learning, with a regression effect size of 0.63 and quantile estimates ranging from 0.73 to 1.3 standard deviations.51 That is a real result, in a real course, against a genuinely good control condition — active learning, not lecture. And here is the caution the field owes you. The single most-cited meta-analysis of ChatGPT's effect on learning, reporting a large positive effect across fifty-one studies, was retracted in April 2026, the journal citing discrepancies in the meta-analysis that undermined confidence in its validity and conclusions.52 I will not quote its number, and neither should anyone else. A non-retracted meta-analysis of thirty-seven studies reports a smaller but real effect, Hedges' g of 0.577.53 On the alarmist side, the same discipline is required. The widely circulated MIT Media Lab EEG study — the "your brain on ChatGPT" study — reported reduced brain connectivity, the lowest sense of essay ownership, and the worst ability to quote one's own writing among participants using a language model.54 It also had fifty-four participants, of whom eighteen returned for the fourth session, was drawn from one geographic area, and has not been peer-reviewed. The authors themselves write that "all the conclusions are to be treated with caution and as preliminary."55 Anyone using it to claim that chatbots damage children's brains is misrepresenting the authors. The honest synthesis: the effect of a machine on learning depends on the design of the machine and the design of the assessment, and almost not at all on the enthusiasm of the person buying it. A tutor built to withhold the answer and demand the reasoning helps. A tutor built to produce the answer, assessed by a test that grades the answer, harms — and harms invisibly, because the practice scores go up while the capacity goes down. Every district in America is choosing between those two things right now, and most of them are choosing on the basis of a sales demonstration.

Cheating was never the emergency. A survey of 4,354 high school students across six schools in four states, fielded in 2024 and published in a peer-reviewed journal in 2026, found 72 percent reported at least one academically dishonest behavior in the prior month, and about 24 percent reported using AI or a digital device as an unauthorized aid.56 The earlier before-and-after work at the same schools found overall cheating rates of roughly 60 to 70 percent both before and after ChatGPT's release.57 The authors' own conclusion is that overall cheating prevalence has not changed. Students were already outsourcing the artifact, because the artifact

was what we graded. The machine did not create the incentive. It industrialized the supply. And the response — buy detection, buy monitoring — treated a curricular design failure as a policing problem, which is how schools ended up spending the next three years building surveillance instead of rebuilding assessment. Detection cannot bear the weight that has been put on it. OpenAI built a classifier to identify AI-written text and launched it in January 2023. Its own page now reads: "As of July 20, 2023, the AI classifier is no longer available due to its low rate of accuracy." In evaluation it correctly identified 26 percent of AI-written text while falsely labeling 9 percent of human text as machine-written.58 The company that built the generator could not build the detector. The dominant commercial detector reports, by its own account, a document-level false-positive rate under one percent for documents with at least twenty percent AI writing, and a sentence-level false-positive rate of about four percent; the vendor advises that the indicator be used "to initiate a conversation, not to draw a conclusion."59 I have seen the figure "98 percent accuracy" attached to this product repeatedly; I could not find it in the vendor's own documentation, so I do not repeat it as a fact. Then the finding that should have ended the practice on the spot. A study in Patterns tested seven detectors and found that over half of non-native English writing samples were misclassified as AI-generated, while accuracy on native samples remained near perfect — and that the detectors were trivially defeated by prompting the model to use more literary vocabulary.60 The tool is most likely to falsely accuse the student with the smallest vocabulary and the least standing to fight back, and least likely to catch the student sophisticated enough to ask for better prose. A one-percent institutional false-positive rate is not a statistic when it lands on your child. It is an accusation of dishonesty, in a file, at fifteen.

Surveillance filled the vacuum that pedagogy left. A peer-reviewed content analysis of the school-monitoring vendor market, published in the Journal of Medical Internet Research in July 2025, examined fourteen companies. Eighty-six percent monitor students twenty-four hours a day, seven days a week. Seventy-one percent use AI for automated flagging. Ninety-three percent monitor school-issued devices, and 36 percent claim to monitor students' personally owned phones and computers. Twenty-nine percent generate a student "risk score." Pricing was unavailable for thirteen of the fourteen, and six of the fourteen market help applying for government grants to fund the surveillance.61 Twenty-nine percent generate a risk score on a child. It is a number, produced by a private model, about a minor, held by a vendor, visible to a school, and — in most states — not disclosed to the parent, not explained to the student, and not appealable by anyone. The harms are documented. In survey research by the Center for Democracy and Technology, 29 percent of LGBTQ+ students reported that they or someone they know had been outed by monitoring technology.62 In related CDT work, 44 percent of teachers reported that a student had been contacted by law enforcement as a result of monitoring.63 And when two United States senators investigated four of the largest monitoring vendors in 2022, they found that 43 percent of surveyed teachers said their schools used monitoring to identify student discipline violations, that three of the four companies did not directly notify students or guardians, and that none of the four had analyzed their products for discriminatory bias.64 None of them had looked. That is not an oversight. That is a business model that survives only because nobody with power ever asked. I should also say what I could not find, since the absence is itself informative. I could not locate a Government Accountability Office study specifically evaluating whether student monitoring software prevents the harms it is sold to prevent; the GAO's K-12 work in this area addresses cybersecurity and data breaches.65 There is a widely repeated claim that a major research institution found "scant evidence" of efficacy, and I could not verify it against the original, so I leave it out. We are three years into universal deployment on tens of millions of minors and the efficacy literature is thin enough that I cannot cite it cleanly. That is the finding.

The records are permanent, and the market holding them is not accountable. In December 2024, an intruder accessed the student information system of PowerSchool, a vendor whose products sit under a very large share of American school districts. The company learned of the incident on December 28, notified customers on January 7, and disclosed in May that a threat actor was contacting individual districts attempting to extort them with the same data.66 The Texas Attorney General sued, alleging that more than 880,000 Texas schoolchildren and teachers were affected and that the stolen data included names, addresses, Social Security numbers, medical details, disability and special-education records, and bus stop information.67 The perpetrator, nineteen years old, was sentenced to four years with more than $14 million in restitution; reporting on the sentencing put the total scale at more than sixty million students and ten million teachers and noted that PowerSchool paid a ransom of $2.85 million in bitcoin.68 The company has never published a total number of affected individuals. Bus stop information. On children. For sale. And the background rate: 24 percent of teachers in the CDT survey said their school had experienced a large-scale data breach in the 2023–24 school year, while the share of teachers receiving substantive privacy training fell from 69 percent to 56 percent between 2019–20 and 2023–24.69 The vendor bears one contract cycle; the child bears the record. In March 2024, the Los Angeles Unified School District unveiled "Ed," an AI chatbot for students, under a contract worth up to $6 million over five years, of which roughly half had been paid. On June 14, 2024 — twelve weeks later — the district unplugged it, the same day the vendor told its employees it was virtually shutting down. A former employee said seven of eight chatbot requests were processed through overseas servers.70 The vendor's founder was arrested that November and accused of defrauding investors of nearly $10 million; prosecutors allege she claimed $3.7 million in revenue from ninety-two customers in a year when actual revenue was about $11,000. The company filed for bankruptcy. By 2026 the FBI had searched district headquarters and the superintendent's home, and the superintendent was placed on indefinite paid administrative leave.71 Set aside the fraud. Ask the question that concerns your child: where did the student data go? A bankrupt company's servers are an asset in a proceeding. The reporters who asked that question in 2024 did not get a clean answer.72 Two tiers, hardening into architecture.

At one end, private schools built around machine instruction, where the model is two hours of screen-based core academics in the morning and then workshops, small groups, and adults for the rest of the day. Tuition at one such campus is $50,000 for the coming year, with published tuition across the network's cities running from roughly $40,000 to $75,000.73 At the other end, a tuition-free public charter approved by Arizona's charter board in December 2024, serving grades four through eight, with the same two-hour academic engine and a marketing claim of "2.8 times faster learning" measured on the vendor's own chosen assessment — a claim I could find no independent verification of.74 Notice what the expensive version buys. It does not buy a better algorithm; it is broadly the same engine. It buys the adults in the afternoon. The wealthy school purchases AI and small classes and expert teachers and a human being who notices when a child goes quiet. The poor district is being sold AI instead of them. And the floor keeps dropping. The Affordable Connectivity Program, which subsidized home internet for more than twenty-three million households — about one in six in the country — ended on June 1, 2024 for lack of congressional funding; 68 percent of enrolled households had reported inconsistent or no connectivity before enrolling.75 Federal data for 2023 show 2.3 percent of children aged three to eighteen with no home internet access and 5.4 percent whose only access is through a smartphone.76 Five percent sounds small until you convert it: that is roughly three and a half million children being assigned homework on a platform they can only reach on a phone. The teacher is being restructured, not assisted. The time savings are real: teachers using AI at least weekly reported saving 5.9 hours a week, about six weeks over a

school year.77 I take that seriously; teacher workload is a genuine emergency and I will not sneer at six weeks. But a district that has not trained its teachers, has not changed its assessments, and is under budget pressure will not experience those six weeks as relief. It will experience them as capacity — and capacity, in a budget conversation, is an argument for a larger class or one fewer position. The path from "AI saves teachers time" to "AI reduces the number of teachers we need" runs through exactly one decision, made by a school board, usually in a meeting no one attends. The deeper risk is what the scripted vendor product does to the teacher's own expertise. A teacher who delivers a system's prescribed sequence for five years, and whose judgment is overridden by the platform when it disagrees, is a teacher who has been quietly converted from a professional into a delivery mechanism. That is how you lose the only person in the room who can tell that a child is faking comprehension.

The federal adult left the room — and took the safety guidance with them. In April 2025 the President signed an executive order, "Advancing Artificial Intelligence Education for American Youth," creating a White House task force, a Presidential AI Challenge, and deadlines for the Education Secretary to issue guidance on using existing grant funds for AI.78 I have no objection to any of that in principle. Here is what happened in the same period. In March 2025 the Department of Education announced a reduction in force that took its headcount from 4,133 to 2,183, followed by an executive order directed at closing the Department; a district court enjoined the action in May and the Supreme Court stayed that injunction in July 2025, with three justices dissenting.79 The litigation is still live; a judge granted the plaintiffs partial discovery in February 2026.80 Through 2026, trade and government press reported successive transfers of programs and staff to the Labor, State, and Health and Human Services departments, including reporting that more than $20 billion a year in K-12 funding would be administered by the Labor Department under an interagency agreement.81 I attribute those to the press because no primary agency document confirming them was available to me. And this. In November 2024 the Department's Office for Civil Rights published a document called Avoiding the Discriminatory Use of Artificial Intelligence, with worked examples — an AI tool that falsely flagged the only two English Learner students in a class; risk-prediction software that used two-parent-household status as an input. That page now carries a notice that the document "has been formally rescinded by the Department and remains available on the web for historical purposes only."82 Read that sequence in order. The federal government told schools to adopt AI, cut in half the agency that would supervise it, and withdrew the single federal document that explained how these systems discriminate against children. That is not deregulation. That is issuing a permit and closing the inspections office. What is left is a patchwork. A nonprofit tracker counted thirty-five states plus Puerto Rico with official K-12 AI guidance as of August 2026 — there is no federal tally, which is itself the point.83 Louisiana's, issued in August 2024, is a reasonable document: it calls itself the first edition of a living document and sets four tiers of classroom use, from AI-Empowered to AI-Prohibited.84 It is also guidance. It does not bind a vendor, does not require an error rate, and does not give a parent anything to hold.

The case that has not been filed yet. More than half of licensed special-education teachers report using AI in some way to help develop individualized education programs, and 15 percent report using it to write an IEP or 504 plan in full — nearly double the prior year, according to reporting on survey research I was not able to retrieve at the source, so I flag it as secondhand.85 Now recall that the federal guidance on discriminatory AI has been rescinded, and that the only legal instrument protecting these children is a plan whose narrative a machine may have written. I searched for a documented case — an adjudicated due-process decision, a civil-rights finding, a published audit — turning on an algorithmic special-education determination. I could not find one.86 It does not mean no child has been misclassified. It means there is no mechanism by which a parent

would ever learn that a model contributed to the decision, and therefore no mechanism by which the harm becomes a case. The absence of litigation is not evidence of safety. It is evidence of opacity. V. The concerns, stated plainly enough to read aloud at a school board meeting. VI. Take the whole thing together and the cumulative issue is not cheating, not screens, not even privacy, though privacy will produce the first lawsuits. The cumulative issue is that we are automating the production of the artifact while continuing to grade the artifact, and simultaneously dismantling every mechanism by which an adult could notice. The assessment does not detect the hollowing. The detection tools do not work and misfire on the vulnerable. The monitoring apparatus watches behavior instead of learning. The teacher has not been trained. The parent cannot see the model. The state guidance is voluntary. The federal civil-rights document has been withdrawn. The vendor's exposure ends with the contract; the child's record does not end at all. A system with no functioning feedback loop does not fail loudly. It drifts. And the drift will not be visible in the test scores for years, because the practice scores will look better — that is precisely what the PNAS study found before the crutch was taken away.87 We will get the good news first and the bill later, and the children who pay it will be the ones whose parents had the least capacity to inspect anything.

There is a second loss, harder to measure and in my view larger. The capacities the essay was for — reading a source skeptically, noticing a contradiction, holding a position under pressure, saying I don't believe that and being able to say why — are the equipment of citizenship. Every other chapter in this book depends on a population that has them. A generation that has never once had to build an argument alone will not be able to evaluate the machine-generated claims that arrive during their first election, their first mortgage, their first diagnosis, or their first public hearing about a data center in their parish. The education chapter is not one chapter among eighteen. It is the load-bearing wall. VII. Objections I take seriously. "The tutor genuinely works, and you are arguing against a proven intervention." Partly right, and I concede it. The Harvard trial is good evidence that a well-designed AI tutor can outperform strong in-class active learning.88 My argument is not that machine tutoring is ineffective. It is that the same technology produced a 17 percent deficit in a different design, and that the difference lay in whether the system withheld the answer and demanded reasoning.89 I am against buying an unspecified tool with an unpublished error rate and no independent evaluation, and calling it the tutor from the study. "Teachers are drowning and you are standing between them and six weeks of their lives." This is the objection that gives me the most trouble, because it is true that teachers are drowning and true that the savings are documented.90 So let me be exact: I would let a teacher use these tools tomorrow for planning, differentiation, drafting, and administrative work, gladly and without a permission slip. What I would not do is let a district book those hours as a staffing reduction, or place a machine in the chain of any decision about a specific child without a named human being who can be appealed to. "You are romanticizing struggle. Every generation said this about calculators." They did, and the calculator argument is better than its critics admit — arithmetic fluency did decline, and we decided the trade was worth it, deliberately, and we restructured the assessment around it. We changed what we tested to protect the reasoning we cared about. The calculator generation's mistake was not adopting calculators; it was, where it happened, giving students calculators while still testing computation, which produced students who could neither compute nor reason. That is precisely the configuration we are in now, at scale, without having had the conversation.

"The equity argument cuts the other way. A free machine tutor is worth far more to a poor child than a rich one." This is the strongest objection in the set and I think it is half true, which makes it dangerous. Yes: an infinitely patient tutor at three in the morning is worth more to the kid whose parents work nights than to the kid with a $200-an-hour human. If we deployed a well-designed, independently evaluated tutor to the students

furthest behind, I would support it, and the recent gains at the tenth and twenty-fifth percentiles suggest the floor is where movement is possible.91 But that is not what is being sold to poor districts. What is being sold is a substitution — the software in place of the counselor, the aide, the small class — while the affluent version keeps every adult and adds the software.92 I want the tutor. I want it delivered on top of a staffing floor, in writing. "You cannot ban the future, and states that try will fall behind." I am not proposing a ban. Every requirement below is a condition on public procurement and public authority — what a school district may buy with tax money and what may be done to a minor without a human answerable for it. A vendor that cannot publish an error rate is not the future. It is a supplier that has not finished its work. "Assessment reform is a decade-long fight and the tools are here now." Also true, and it is the reason I put assessment first rather than second. If we buy the tools before we change the tests, the tools will be optimized to produce the thing the tests measure, the vendors will build to that specification, and the contracts will be renewed on the strength of scores that no longer mean anything. Then the reform becomes impossible, because the entire apparatus will have a financial interest in the old measure. The sequence is not a preference. It is the difference between a tool and a trap. VIII. Three questions I would ask about any system placed between a teacher and a child. They are not technical. Anyone can ask them out loud, in a public meeting, and the quality of the silence that follows is the answer. Who can witness it? Can a parent see the model, or the rule, or the score that sorted their own child? Today, in almost every district in the country, the answer is no. A risk score is generated by a private model on a minor and is not disclosed, not explained, and not appealable. A track recommendation arrives as a fact. A flagged essay arrives as an accusation. If no one outside the vendor can see the mechanism, nothing about it can be contested, and an uncontestable decision about a child is not an educational judgment. It is an administrative event that happens to them.

Who carries the burden if it is wrong? The student carries a misclassification for years — into a track, into a disciplinary record, into an IEP that was drafted by a system and signed by a person who did not write it. The vendor carries it for one contract cycle, and sometimes not even that: the Los Angeles vendor was gone in twelve weeks and its founder's exposure was to investors, not to students.93 When the party making the decision is structurally immune to the consequence, the decision will get worse over time. Not from malice. From the ordinary physics of incentives. Can it be undone? There is usually no mechanism to reverse a track assignment, expunge a behavioral flag, delete a monitoring record, or claw back exfiltrated data. The Social Security numbers and disability records taken in the PowerSchool breach are not recoverable.94 A seventh grader's risk score, once generated, has no expiry that anyone has published. We are making permanent records about people who are legally incapable of consenting to them, and we have built no exit. Any system that fails all three — invisible, unaccountable, irreversible — should not be within a hundred feet of a child, regardless of what it does for test scores. IX. What I would insist on. Not aspirations. Conditions of purchase and conditions of authority, each with a name attached and a consequence for failure. 1. No instructional or assessment system in a public school without published error rates, subgroup performance data, and independent evaluation. Published means on a public page, in plain language, before the contract is signed: what the system gets wrong, how often, and for which groups of students — English learners, students with disabilities, students by race and income. Independent means not the vendor and not a study the vendor commissioned. If a detector's false-positive rate on non-native English writers is unknown, it is not deployable; the one published study on that question should have ended the practice in 2023.95 Who: the school board that signs, the superintendent who recommends, the state education agency that maintains the approved-vendor list. If not: the contract is void and unenforceable, and no public money may be spent under it.

Full prohibition on the sale, licensing, or secondary use of minors' educational data, with statutory deletion timelines. No resale. No model training on identifiable student work without affirmative parental opt-in that is not a condition of enrollment. No transfer of student data as an asset in a bankruptcy or acquisition. Hard deletion deadlines — behavioral and monitoring records within one year, everything not part of the permanent academic transcript within two years of a student leaving the district — with certification of deletion filed with the state. Who: the vendor's named data officer, the district privacy officer, the state attorney general. If not: statutory damages per student, payable to the students, and disqualification from public contracts for a period of years. The Texas action shows the mechanism exists; what is missing is that it should not require a catastrophe to trigger it.96 3. Assessment reform first, tools second. Before a district buys generative tools at scale, it must show how it assesses the things a machine can now produce: oral defense of written work, in-process drafting visible to the teacher, portfolios across a year, problem sets done in the room, demonstrated reasoning rather than delivered answers. This is more expensive in teacher time than a scored essay. That is not an argument against it; it is the actual cost of knowing whether a child can think. Who: the state assessment authority, the district curriculum office, the accrediting body. If not: no procurement approval for instructional AI. The tests change first. 4. Funded, paid teacher preparation as a precondition of any district contract. Not a webinar. Paid time, during contract hours, before deployment, covering what the tool does, how it fails, when to override it, and how to teach students to interrogate its output. Eighteen percent formal guidance is a scandal that every party to it can see and no party is accountable for.97 Who: the district — training funded as a line item within the contract, not from professional-development funds that get cut in March. If not: deployment is prohibited until the training is complete and documented.

A universal verification-and-civics literacy requirement for graduation. Every student, before a diploma: how to trace a claim to a primary source; how to tell a study from a press release; how to read a percentage and a base rate; how automated decisions about people are made; how to find out whether one was made about you; and how to contest it. Taught as a practiced skill and assessed by demonstration, not by multiple choice. Tool operation is not literacy. Source evaluation is literacy. The confidence to say I don't believe that, out loud, with a reason, is literacy. Who: the state board of education, in graduation requirements. If not: it will be an elective that four percent of students take. 6. Human authority over every placement, discipline, and eligibility decision, with a named appealable person. A machine may inform. A machine may not decide, and may not be the reason no one will say who decided. Every automated recommendation touching a specific child carries a name, a title, and a working address; the family is told in writing that an automated system was involved, what it considered, and how to appeal; and the appeal goes to a human with the actual authority to reverse it. This is the provision that would let the special-education harms currently invisible in the record become visible at all.98 Who: the named decision-maker; the superintendent, jointly and personally. If not: the decision is vacated on request and the record expunged. 7. A floor on human staffing that AI spending cannot be used to reduce. In writing, in the board's budget resolution: teachers, aides, counselors, librarians, and special-education staff per student may not fall below a stated floor, and no efficiency gain from any AI system may be booked as a staffing reduction for the life of the contract. If the software genuinely creates capacity, spend it on smaller groups and on the students furthest behind. This is the single provision that separates the version of this technology that helps poor children from the version that abandons them, and it is the one no vendor will ever propose. Who: the school board, by recorded vote, in public. If not: the contract terminates, and the savings are clawed back into staffing. I have written enough governance documents to know the objection: this is a lot of process for a district with three people in its central office. Yes. So cut it down to the two that cannot be given up. Number

six, because a decision about a child with no human name on it is intolerable in a free country. And number seven, because everything else in this list is decoration if the adults leave the building. X.

I do not believe the machine is the enemy of learning. I have used it. It made this book faster and, in places, better, and it was useless to me precisely where I lacked the judgment to check it. That is the finding. The tool amplifies a capacity it cannot create. It rewards the person who already knows how to think and hollows out the person who does not yet, and school is the entire span of years in which a human being is in the second category. We are proposing to hand the amplifier to people before they have the thing to amplify, and we are doing it without changing the assessments that would tell us whether it worked, without training the adults, without publishing the error rates, without a name on the decisions, and having just withdrawn the one federal document that explained how these systems hurt children. What I keep returning to is the state an hour from my house. Louisiana went from fiftieth in the country to sixteenth in fourth-grade reading, and it did it with phonics, screening, tutoring, teacher training, hard conversations with parents, and a decade of adults refusing to look away.99 There was no algorithm in it. There was attention, applied by named human beings who could be held responsible, sustained past the point where it was politically comfortable. That is the standard. Any technology that can survive inside it is welcome. Any technology that requires us to abandon it — that asks us to stop looking, stop training, stop staffing, stop explaining ourselves to parents — should be refused. A machine can write the essay. A machine cannot want anything for a child. Every single thing that has ever worked in education ran on somebody wanting something for a specific child and being answerable for whether they got it. Do not automate the wanting, and do not, under any circumstances, automate the answerability. If we get this chapter wrong, the rest of the book has no audience capable of reading it. — D.M.D.

Chapter Three: Memory

and the Record

Who keeps history, and whether it stays correctable

If the record is now generated, curated, and stored by private systems, what happens to the past? I. On the twelfth of August, 2026, the Louisiana Public Service Commission voted three to one to let Meta keep secret the figures behind its Hyperion data center in Richland Parish — the investment amount, the job-creation claims, and the electricity-demand calculations that everything else in the proceeding depends on. An administrative law judge had ruled the month before that the company should answer the subpoenas. The Commission overruled her. One commissioner, Davante Lewis, dissented.100 I want to be precise about what happened, because it is easy to hear it as a fight about a company and miss that it was a fight about a record. Nobody voted that the numbers were false. Nobody voted that the numbers did not matter. They voted that the numbers would not exist in the place where the public keeps things — that when some researcher in 2041 tries to reconstruct how Louisiana decided to build out several gigawatts of generation for one customer, the file will contain the decision and not the basis for it. That is how the past gets edited now. Not by burning anything. By declining to make it. Here is the admission I owe, because I have been on the other side of this table. I ran a company. I signed a document-retention policy. I understood, in the way that everyone in that chair understands, that a retention policy is two things at once: a compliance instrument and a liability instrument. It says what we keep. It therefore says what is gone by the time anyone asks. I did not think of myself as editing history. I thought of myself as managing risk, which is what the lawyers called it, and they were right, and it was also editing history. I have sat in board meetings where a decision was recorded in the minutes as a resolution, and the reasoning — the part a historian would want, the part that would show whether we were serious or merely covered — was never written down anywhere by anyone, on purpose. So I am not going to pretend that private control of the record is a novelty the machines invented. It is an old habit. What is new is the scale, the speed, and the fact that the tools which now generate our account of the world are owned by parties with retention policies of their own, and no obligation to tell us what they are. II. Assemble the pressures. Each is defensible standing alone.

Storage is cheap, so agencies stopped printing. Digitization is expensive, so libraries partnered with vendors. Vendors need revenue, so access acquired a price. Litigation is ruinous, so retention windows shortened. Moderation is mandatory, so deletion became automated. Models need training corpora, so text was collected at a scale that made asking permission impractical, and impracticality became a legal argument. Security is real, so infrastructure records became exempt. Economic development is competitive, so investment figures became trade secrets. Classification is necessary, so the count of classifications became a burden to publish. And every one of those choices was made by someone reasonable, in a meeting, with a memo, for a reason. The result is not a conspiracy. It is a drift, and drifts are harder to stop than conspiracies, because there is no one to indict and nothing to repeal. III. Before blaming the machine, let us be honest about the ground it landed on. The record was already rotting. The Pew Research Center sampled roughly a million web pages that existed at some point between 2013 and 2023 and checked them in October 2023. Twenty-five percent were gone. Of pages that existed in 2013, thirty-eight percent were gone; of pages from 2023 alone, eight percent had already vanished within the year. Twenty-one percent of government webpages contained at least one broken link. Of about five million tweets collected in the spring of 2023, twenty-three percent were no longer visible.101 It is

worse in the places we assume are permanent. Jonathan Zittrain, Kendra Albert, and Lawrence Lessig found that roughly half of the URLs cited in United States Supreme Court opinions no longer point to the material the Court cited. More than seventy percent of URLs in three Harvard journals had rotted. And here is the finding I think about most: of 353 Supreme Court URLs that still returned a live page, only about three-quarters actually led to the cited content. A working link is not a correct link. The citation resolves, the reader is satisfied, and the authority is gone.102 A Harvard Library Innovation Lab study of the New York Times — 553,693 articles containing more than two million external links — found a quarter of deep links inaccessible, and drift in the content of thirteen percent of the still-live links that humans checked by hand.103 That is a working paper, not a peer-reviewed result, and I flag it as such. The direction is not seriously in dispute. Now the complication, and I want to give it real weight.

When the removals began in 2025, institutions fought back and sometimes won. In Doctors for America v. OPM, Judge John D. Bates ordered HHS, CDC, and FDA to restore identified webpages and datasets to their January 30, 2025 versions — and set the deadline for that same evening.104 When the administration moved to dismantle the Institute of Museum and Library Services by executive order in March 2025, twenty-one state attorneys general sued in April; a temporary restraining order issued May 1, a preliminary injunction May 14, and in November 2025 the court permanently blocked the order. The terminated grants were reinstated that December, and the appeal was withdrawn under a binding settlement in April 2026.105 Volunteers built the Data Rescue Project, which reports more than nine hundred people working through federal data holdings.106 So the record is not defenseless. Courts intervened. Librarians organized. Archivists worked nights. Any argument in this chapter that treats loss as inevitable is wrong, and is also insulting to the people who spent eighteen months preventing it. What I take from the counter-example is narrower: every one of those saves was a rescue. Emergency response, after the fact, dependent on someone noticing in time, on a plaintiff with standing, on a judge willing to move in days. That is not a preservation system. That is a fire brigade. And the drift I described in the previous section is a building code problem. IV. Here is what is actually happening, named. Custody migration. The record has moved onto infrastructure the public does not own. What that means in practice became visible in 2025, when the Prosecutor of the International Criminal Court, Karim Khan, lost access to his email following a United States executive order sanctioning the court; the Associated Press reported the loss of access, frozen bank accounts, the departure of six senior officials, and three resulting lawsuits.107 Microsoft has consistently denied cutting off the institution: its president, Brad Smith, said the company's actions "did not in any way involve the cessation of services to the ICC."108 In February 2026 Microsoft told a UK parliamentary committee that its own earlier testimony contained an inaccuracy, apologized, and asked for the record to be corrected — its corrected position being that it was "in touch with the ICC throughout the process that resulted in the disconnection of its sanctioned official from Microsoft services," while at no point ceasing service to the court itself.109 Read that formulation carefully. It is not a denial that a chief prosecutor was disconnected. It is a clarification of who was disconnected and by whose decision. Microsoft has since made five "European Digital Commitments," including a pledge to contest in court any order to suspend European cloud services.110 I take the pledge as sincerely meant. It is a corporate commitment, revocable by the same board that issued it, and it is currently the strongest protection the institutional memory of an international court enjoys. Retention set by the vendor and by whoever is suing the vendor. In May 2025, Magistrate Judge Ona Wang ordered OpenAI to preserve and segregate output log data that would otherwise have been deleted.111 In October 2025, a stipulated order terminated that obligation effective September 26, 2025, with continued retention narrowed to a specified domain list and excluding users

in the European Economic Area, Switzerland, and the United Kingdom.112 I am not arguing the order should have persisted — the privacy objections to indefinite retention of everyone's conversations are strong, and I share them. I am pointing at the mechanism. Whether hundreds of millions of records of how people actually used these systems in 2025 survive was decided by litigation posture between two private parties. No archivist was in the room. No statute governed it. Removal at the source. The Environmental Data & Governance Initiative, which monitors more than four thousand federal environmental webpages, recorded seventy percent more changes in the first hundred days of 2025 than in the same period of 2017 — and notes that its four thousand pages represent only about a fifth of the scope it was able to monitor during the first Trump administration.113 The Data Rescue Project's estimates of datasets removed from public access range from three to four thousand.114 These are advocacy counts, self-reported, measuring different things, and I will not blend them into a single number. They are also the only counts that exist, which is itself the finding.

Communications built to leave nothing. Federal law is not ambiguous: 44 U.S.C. § 3101 requires each agency head to make and preserve records documenting the agency's decisions and essential transactions.115 In March 2025, after senior national security officials discussed military operations on Signal, American Oversight sued; Judge James Boasberg ordered preservation from the bench within two days. In June 2025 the court granted a preliminary injunction in part — requiring officials to notify the acting Archivist of at-risk records and requiring the Archivist to refer the matter to the Attorney General — but declined to order recovery of the messages.116 On April 30, 2026, the parties filed a joint stipulation dismissing the case with prejudice.117 The statute held. The remedy did not. The messages were never recovered, and as far as the public record shows, no one was penalized for their absence. Laundering, in the technical sense: severing a work from its author. In Bartz v. Anthropic, Judge William Alsup held in June 2025 that training on lawfully acquired books could be fair use while retaining pirated copies was not.118 The case settled for $1.5 billion, given final approval on July 20, 2026 — roughly $3,000 per work, against a corpus of about seven million downloaded copies from which some five hundred thousand titles qualified for the class.119 In Kadrey v. Meta, Judge Vince Chhabria granted Meta partial summary judgment on fair use in June 2025, expressly limited to the thirteen named plaintiffs and that record, while the claim concerning Meta's torrenting of the corpus survived.120 These are compensation events and narrow holdings. Neither restores attribution. A book absorbed into a model's weights is not cited by it, cannot be traced from its output, and does not appear in any bibliography a reader could check. Payment settles the debt to the author. It does nothing for the reader downstream who needs to know where a sentence came from. Deletion as moderation, applied to evidence. Human Rights Watch checked 5,396 pieces of content it had cited in 4,739 of its own reports and found 619 — eleven percent — had been removed from the platforms where they lived. In the same report it cited YouTube's own transparency data: over six million videos removed in a single quarter of 2020, 93.4 percent first flagged by automated systems, and half taken down before receiving a single view.121 The Syrian Archive, run by Mnemonic, states it has reinstated more than 650,000 records to platforms — its own count, not independently audited.122 Automated systems, built to remove atrocity, removed the evidence of atrocity, at machine speed, before human review, and a small nonprofit has spent a decade putting some of it back. Enclosure of what libraries may lend. In September 2024 the Second Circuit affirmed against the Internet Archive in Hachette Book Group v. Internet Archive, holding that controlled digital lending of 127 in-copyright books was not fair use.123 The recording industry's separate "Great 78" suit settled on confidential terms in September 2025 — no damages figure exists, and anyone who cites one is guessing.124 Meanwhile the statute libraries actually rely on, 17 U.S.C. § 108, permits one copy for most purposes and up to three for preservation or replacement of obsolete formats — a numeric limit written

for a world of photocopiers.125 The Copyright Office has a reform track open, including a proposal to replace the fixed three-copy limit with a "reasonably necessary" standard. Nothing has been enacted.126 Toll-gating the public file. PACER, the federal courts' records system, charges ten cents a page, capped at three dollars per document, with an automatic waiver under thirty dollars a quarter. Effective January 1, 2027, and for five years, the fee rises to twelve cents and the waiver threshold to forty dollars — the first change in about fifteen years.127 In National Veterans Legal Services Program v. United States, plaintiffs argued the courts had charged more than the law allowed; the case settled for $125,000,000, and the Federal Circuit affirmed the settlement on March 20, 2026.128 A bill to change the underlying structure, the Open Courts Act of 2026, was introduced on June 2, 2026 and has not been enacted.129 The judiciary's own budget documents report PACER processing just over 619 million requests in fiscal 2024, with eighty-four percent of active users paying nothing and projected annual revenue around $142 million.130 The fee is small, most users are exempt, and the system still generated a nine-figure revenue stream and a nine-figure settlement. Small tolls at scale are not small.

The page that is closed, and the count that is no longer kept. This is the finding that genuinely surprised me. The Information Security Oversight Office's annual report to the President for fiscal 2024, transmitted in August 2025, does not publish the number of original or derivative classification decisions made across the federal government, and gives no explanation for the omission.131 It reports other things carefully — nineteen agencies with delegated original classification authority, 1,661 delegations, five on-site inspections in which 161 classified documents were reviewed and thirty-four percent were found to contain errors or discrepancies, and roughly 1,200 mandatory-declassification-review appeals pending before the interagency panel.132 But the headline number is simply absent. It used to be there. For fiscal 2012, the government reported 73,477 original and 95,180,243 derivative classification decisions.133 For 1998 it reported 137,005 original decisions.134 I am not able to tell you the current figure, because the government no longer says. Separately, and in the same period, the Justice Department reported that federal agencies received a record 1,501,432 Freedom of Information Act requests in fiscal 2024 and ended the year with 267,056 backlogged — up thirty-three percent in a single year, with the Department of Homeland Security alone accounting for fifty-three percent of the backlog.135 There is a closed page in every state. Some material genuinely must be preserved without being published — sources, victims, minors, weapons design. I am arguing that a state which will not tell you how many pages it closed last year has stopped being able to promise that the closures are lawful, and that a request queue of a quarter million is a denial system operating under the name of a disclosure system. The local record, going dark. Medill's State of Local News 2025 counts 213 counties with no local news source, 1,524 with only one, roughly fifty million Americans with limited or no access to local news, and 136 newspaper closures in a single year — more than two a week — with newspaper employment down seven percent in the year and more than three-quarters of newspaper jobs gone since 2005.136 The archives go with the papers. The Columbia Journalism Review documented the pattern: the Canadian Record in Texas ended its print edition and donated its archive to the University of North Texas; Gannett gave roughly ten million photographs from the Louisville Courier Journal to an institution that could hold them; a Chicago Tribune photo editor found that of her own archival work, two images remained.137 The Poughkeepsie Journal's photo morgue of more than 200,000 images became a rescue project.138 Where a local paper closes and no university is nearby, there is no story to document. A century of a town's photographic record leaves in a dumpster and generates no record of its own leaving.

Authentication offered in place of truth. The Coalition for Content Provenance and Authenticity has built the leading technical answer: cryptographically signed provenance metadata, marketed as Content Credentials. It is

real engineering and I support it. But note where it stands. As of August 2026 the specification site publishes version 2.4, while the downloadable specification PDF is version 2.2 dated May 2025 and the coalition's own announcement page celebrates the launch of 2.3 — three live version numbers at once.139 The corresponding international standard, ISO/CD 22144, remains under development with no published edition.140 On the regulatory side, Article 50 of the EU AI Act requires providers of systems generating synthetic audio, image, video, or text to ensure outputs are marked in a machine-readable format and detectable as artificially generated — qualified, in the text itself, by "as far as technically feasible" and by "the generally acknowledged state of the art."141 Those obligations are scheduled to apply from August 2, 2026, with systems placed on the market before that date reportedly receiving a four-month grace period for machine-readable marking under the Digital Omnibus agreement.142 My worry is not that provenance fails. It is that provenance succeeds and we mistake what it proves. A signature establishes that a file has not been altered since signing, and who signed it. It does not establish that the thing depicted occurred. An authenticated photograph of a staged event is authenticated. A signed dataset built on a false premise is signed. If we build a civic habit of checking the credential instead of checking the claim, we will have replaced verification with a receipt — and the people best equipped to issue receipts at scale are the largest institutions, which is exactly backwards from where scrutiny should point. Irrevocable inference. Nasr, Carlini, and colleagues demonstrated in 2023 that a simple divergence attack made a deployed commercial chatbot emit memorized training data at roughly 150 times its normal rate; the authors report extracting several megabytes for around two hundred dollars of compute and conclude that current alignment techniques do not eliminate memorization.143 That is a preprint, and I cite it as one. Going the other direction — reliably removing a specific fact from a trained model — has not been demonstrated. Researchers surveying seventy-two unlearning papers argue the field's benchmarks are weak measures of progress, and showed that substituting an incorrect answer in place of the target knowledge caused a twenty-eight percent accuracy drop for a leading method, indicating the "unlearned" information was still there.144

Now put that next to the right of correction. Article 16 of the GDPR gives a person the right to have inaccurate personal data rectified without undue delay.145 The advocacy group noyb has been testing whether that right reaches a model: first in Austria in 2024 over a fabricated date of birth, then in March 2025 in Norway on behalf of a man about whom a chatbot generated a detailed account of murdering two of his children and serving twenty-one years in prison.146 That complaint remains pending, with the Irish Data Protection Commission now the lead authority.147 Meanwhile the most-cited AI privacy penalty in the world no longer stands: the Italian Garante's €15 million fine against OpenAI, imposed in December 2024, was annulled by the Court of Rome in March 2026 on procedural grounds.148 Whether a person can compel the correction of a false fact about themselves inside a model is an open question — legally unresolved, technically unsolved, and being litigated one fabricated life at a time. And the commons itself thinning. In October 2025 the Wikimedia Foundation reported that after revising its bot-detection logic and reclassifying traffic from March through August 2025, human pageviews were down roughly eight percent year over year, which it attributed to generative search summaries answering questions without sending anyone to the source, and to younger users' preference for social video.149 That is self-reported and methodology-dependent, and the causal attribution is the Foundation's own. But the structure of the concern does not depend on the exact number: the systems that summarize the encyclopedia depend on the encyclopedia, and they are positioned between it and the volunteers who would have written the next article. V. The concerns, stated plainly enough to read aloud at a hearing. VI. Take it together and the cumulative issue is not that things are being lost. Things have always been lost. Libraries burned, negatives faded, minutes went unwritten, and every generation before ours knew less about its

grandparents than it wanted to. The issue is a change in who holds the pen — and more precisely, in the direction of the burden of proof.

For most of the modern era, the presumption ran one way: a public act generated a public record, and the state had to justify withholding it. That presumption is inverting. Increasingly the record exists only if some private party's retention schedule permits, only if no automated moderation system flagged it, only if a fee was paid, only if a trade-secret exemption was not asserted, only if a court moved within days of someone noticing. The citizen who wants the basis of a decision must now assemble it from fragments, at cost, against the clock, and frequently cannot. The party who declined to produce it faces, in the ordinary case, no consequence. That is not a loss of information. That is a transfer of power, and it is the specific transfer that makes every other protection in this book unenforceable. You cannot hold anyone to a standard if you cannot establish what they did. VII. Objections I take seriously. "Preservation is not free, and you are asking for enormous public expenditure." Correct, and I should be honest about the trade. Comprehensive preservation in open formats under public custody costs real money that competes with roads and clinics. My answer is that we already pay for it, badly, in three currencies: emergency rescue by volunteers and litigators, permanent loss we never quantify, and access fees charged to the people least able to pay them. Preservation is the cheapest form of accountability there is, precisely because it is paid once and drawn on forever. "Some things genuinely must be closed." Yes. Sources whose lives depend on it. Victims. Minors. Grand jury material. Weapons design. Location data on shelters. Sacred material a community has asked not be displayed. I am not making an argument for total transparency, which is a young person's argument. I am making an argument about who decides and whether the decision is countable — that a closure should be made by a named authority, on a stated ground, for a stated period, recorded in a public tally, and reviewable. The state that stopped publishing the number of things it classified did not become more secure. It became less answerable.

"Deletion is a right." It is, and one I support. The strongest form of this objection is that permanent, searchable, inference-reconstructible records are themselves an instrument of oppression, particularly for people whose worst moment is indexed forever. I hold both. The resolution I can defend is a distinction between the record of public power and the record of private life. An agency's decision, a utility's demand forecast, a court's docket, a regulator's vote: these are acts of power and should be permanent and public. A person's medical history, adolescence, immigration file, or mistaken arrest is not an act of power and should be governed by the individual. Where the two collide — a public official's private communications about public business — the public character of the business governs. That line is genuinely hard in the middle and clear at both ends, and most disputes are at the ends. "The platforms are private property and can moderate as they choose." As a matter of current law, largely true. But when a platform becomes the sole repository of the only footage of a war crime, the property framing stops describing the situation. I am not proposing to compel a company to host anything. I am proposing that automated removal of material with evidentiary character should route to an escrow rather than to oblivion — that the choice is not "host it" or "destroy it," and that treating those as the only options is a design decision, not a legal necessity. "Courts and volunteers handled it." They did, and I said so. My concern is the reliance profile. Every save in 2025 and 2026 depended on someone noticing quickly, a plaintiff with standing, a judge willing to act in days, and a foundation willing to fund the labor. Remove any one and the material is gone. A preservation regime that only works when all four hold is not a regime. "You are romanticizing paper." Somewhat, and I should own it. Paper burned, flooded, was pulped for the war effort, and was destroyed on purpose by people who understood exactly what they were doing. I do not think our ancestors were better custodians. I think the failure mode has changed in a way that matters: paper decays

visibly and slowly, and a missing file announces itself as a gap. A digital record can be silently altered, or can resolve to a live page containing different content than it did when cited, and the gap does not announce itself. Three-quarters of live Supreme Court citation links leading to the cited content is the whole problem in one statistic. Silent failure is worse than loud failure, because it does not summon a rescue. VIII. Three questions I would ask about any system that holds, generates, or filters the public record. They are not technical. Anyone can ask them at a hearing, and the shape of the answer tells you what you are dealing with. 61

Who can witness it? An archive that no one outside the holding institution can inspect is not a record. It is a claim. This is the test the Louisiana vote failed in the plainest possible way: the figures exist, the Commission has seen some of them, and the public has been told to accept the conclusion without the basis. It is the test the classification system now fails at the level of arithmetic, since we cannot count what was closed. And it is the test a model fails structurally — a system that cannot say which source produced a given sentence has not published a record at all, however accurate it happens to be. Who carries the burden if it is wrong? Almost always the party with the thinnest presence in the archive. A community whose history was never well documented absorbs every inference later made about it — every model that fills the gap by extrapolating from populations that were documented, every risk estimate, every allocation formula, every historical account assembled by a system that treated absence as evidence of nothing rather than as evidence of a prior failure to record. The person whose fabricated criminal history a model produced in Norway bears the cost, and is still waiting. When the party who controls the record is immune to the consequence of its errors, the record degrades — not from malice, but from the ordinary physics of incentives. Can it be undone? This is the one that decides everything. A corrigible history can be amended by evidence: someone finds the letter, files the correction, and the account changes. An authenticated history cannot. If the signature is the proof, then challenging the account means challenging the signer, and the signer is the party with the keys. Extracted training data can be reproduced; targeted facts cannot reliably be removed. A destroyed Signal thread cannot be reconstructed. A pulped photo morgue cannot be reprinted. A dead citation cannot be resolved. And the one legal right built precisely for this — rectification — has not yet been shown to reach inside a model at all. Any system that fails all three — unwitnessable, unaccountable, uncorrectable — is not keeping the record. It is issuing an official history, and the fact that it is machine-generated rather than state-decreed should be no comfort whatsoever. IX. What I would insist on. Conditions, with a name attached and a consequence for failure.

Public records stay in public custody, in open formats, with statutory obligations that survive the vendor. A public body may use a private platform, but the authoritative copy resides in an archive the public controls, in a documented non-proprietary format, with a full export deliverable within thirty days of any request or contract termination. Every contract for a system that holds public records must contain a preservation clause that binds successors, assignees, and acquirers, and survives bankruptcy. The federal government has already stated the principle for itself, requiring agencies to manage all permanent and temporary records electronically with metadata and closing agency-operated storage facilities.150 What is missing is the same requirement at the state, parish, and municipal level, where most public records actually live. Who: the records officer of the public body, personally named in the contract; the state archivist; the procurement officer who signs. If not: the contract is void, no public funds may be disbursed under it, and the public body may not renew with the vendor. 2. No infrastructure decision on a sealed factual basis. Where a public body decides on the basis of numbers submitted by a private party — investment, jobs, load, water withdrawal, emissions, subsidy — those numbers enter the public record, in full, before the vote. Aggregate or time-delay them if a genuine competitive harm can be

specifically demonstrated, on the record, item by item, by a named official. If the numbers are too sensitive to disclose, they are too sensitive to bind the public to a thirty-year obligation. Who: the commission or council itself, by recorded vote; the state's public records custodian; the attorney general on petition. If not: the decision is void and the proceeding reopens. This is the provision I would have wanted in Louisiana on August 12, 2026.151 3. Provenance and chain of custody for official material, published — and never confused with truth. Every official record carries who created it, when, under what authority, what was changed, when, and by whom. Cryptographic provenance is the right tool and should be required for official releases. It must also be labeled for what it is. No public body should describe a signed file as "verified" when what has been verified is the signature. Content Credentials should be cited by version and retrieval date, because the specification is a moving target, and no one should call it an ISO standard, because there is not yet a published one.152 Who: the issuing agency's records officer; the state chief information officer. If not: the record is inadmissible as an official version, and the agency must produce the underlying material on request.

Independent archives, funded on multi-year appropriations, with escrow for material removed from platforms. Public and university archives, digital preservation programs, and the nonprofits that have been doing this work unpaid need funding that cannot be cut in a single budget cycle or terminated by executive order — a lesson the 2025 library litigation taught expensively.153 Separately: when a platform's automated systems remove material with potential evidentiary value, the material goes to a designated independent escrow — accessible to courts, investigators, and accredited researchers under protocol — rather than to deletion. Human Rights Watch found eleven percent of what it had cited was gone.154 Escrow is the difference between a prosecution and a rumor. Who: the appropriating legislature; the platform's named compliance officer; a designated national archival authority. If not: removal of evidentiary material without escrow is itself a sanctionable act, and the appropriation is a floor that cannot be reprogrammed. 5. A real right of correction to the record about oneself, including inside a model. Not a support ticket. A named recipient, a statutory clock, a written decision with reasons, an appeal to a body with authority to compel, and — this is the part that has no answer yet — an obligation on the operator of a generative system either to correct a demonstrated false statement of fact about an identified person at the output layer with an auditable, testable, published mechanism, or to state plainly that it cannot, and accept the consequence of that admission. Rectification is already law in Europe.155 Whether it reaches a model is unresolved, the technical literature suggests reliable removal has not been demonstrated, and the largest fine to date has been annulled.156 I am not going to pretend the engineering is solved. I am going to insist that "we cannot correct it" is a statement with legal consequences, not an excuse that ends the conversation. Who: the operator's named data protection officer; the supervisory authority; the court on appeal. If not: the false output is enjoined for the affected person, with penalties that scale to the operator's revenue rather than to the individual's damages.

Community control of community data, with the labels and licenses to enforce it. Where material documents a people, that people governs its use. The framework exists: the CARE Principles for Indigenous Data Governance — collective benefit, authority to control, responsibility, ethics — published in 2019 by the Research Data Alliance's international Indigenous data sovereignty group, deliberately positioned as the people-and-purpose complement to the FAIR principles, which the authors note focus on facilitating data sharing "while ignoring power differentials and historical contexts."157 The operational tooling exists too: Local Contexts' Traditional Knowledge Labels — six provenance labels and nine protocol labels — let a community attach its own terms to material held in someone else's collection.158 Any public institution that holds such material should implement both, and no dataset drawn from such holdings should be licensed for model

training without the community's affirmative consent. Who: the museum, university, library, or agency director; the tribal or community authority as counterparty. If not: the license is void and the material is withdrawn from the training corpus and from public display pending consent. 7. Access to the public record at the marginal cost of delivery, and never as a revenue line. Court dockets, agency records, regulatory filings, and legislative material are the raw material of accountability journalism and civic litigation, and charging for them is a tax on scrutiny that falls hardest on the people with the least. A judiciary that processes 619 million requests a year and books $142 million in revenue from doing so has a business, not a public record.159 Fees should recover delivery cost, be waived for journalists, researchers, pro se litigants, and nonprofits by default rather than on application, and should be reported publicly against actual cost each year. A $125 million settlement over overcharging is a finding about the size of the toll.160 Who: the administrative office of the court system; the appropriating legislature; the agency FOIA officer. If not: fees above documented delivery cost are refundable, and the collecting body publishes the shortfall. I have written enough governance documents to hear the objection: this is a great deal of machinery for a parish clerk with two staff. So take the two that cannot be given up. Number one, because a public record held only on someone else's property is not a public record. And number two, because a decision made on a sealed basis cannot be reviewed by anyone, ever, including by the people who will pay for it for thirty years. X.

I have been arguing throughout this book that a society cannot consent to what it cannot see and cannot contest what it cannot verify. This chapter is the third leg: it cannot learn from what it cannot remember, and — more urgently — it cannot correct what it is not permitted to examine. History must remain correctable. That is the entire claim. Not accurate, which no history has ever been on the first pass. Correctable. Amendable by evidence. Open to the person who finds the letter in the attic, the researcher who notices the discrepancy in the numbers, the reporter who files the request, the descendant who says: that is not what happened to us, and here is what did. Every mechanism in this chapter narrows that opening. Some narrow it by accident — link rot, format obsolescence, a paper closing with its photographs still in the basement. Some narrow it by policy — a retention schedule, a trade-secret exemption, a fee, a vote to seal. Some narrow it by architecture — a model that cannot cite its sources and cannot forget a specific fact, a moderation system that deletes before anyone looks, a provenance signature that certifies the file and says nothing about the world. The accidents are fixable with money. The policies are fixable with law. The architecture is the one we are choosing right now, while the choice is still open, mostly without noticing that we are choosing. What I am most afraid of is not a lost archive. It is a complete one. A record that is beautifully preserved, cryptographically signed, instantly searchable, conversationally accessible, and authenticated rather than true — where the answer arrives fluent and sourceless, where the citation resolves to a page that no longer says what it said, where the thing that is missing generates no gap because a system has smoothly interpolated across it, and where challenging any of it requires access that costs money, standing you do not have, and keys held by the party you are challenging. That is not a dark age. Dark ages are obvious. That is something newer and harder to fight: a well-lit room with no windows. So keep the file. Keep it in public hands, in a format you can open in forty years without asking permission, with a name on every closure and a number on every seal. Keep the door open for the person who will prove us wrong. It is the only provision in this entire book that has to survive us, because it is the one that lets our grandchildren correct everything else we got wrong. — D.M.D.

XI. Requirements of Part One Three chapters end here, and between them they ask for twenty-two things: eight for the information commons, seven for the schools, seven for the record. That is too many for any citizen to

carry into a meeting, and I know it, because I have watched people walk into hearings holding a list that long and lose the room by item four.

So this section does the compression. Not a summary — a reduction to what has to be true before any of the rest becomes possible. Take these three, in this order, because they are sequential: the first makes the system visible, the second makes the visibility trustworthy, and the third is what the first two amount to in the one institution that touches every citizen before they are old enough to consent to anything. 1. Disclose before you deploy. No consequential system operates on the public until its purpose, its data sources, its measured error rates, and its known failure modes are published in plain language, in advance, in a place the affected public can find without a lawyer. Plain language is a real requirement and not a courtesy. A model card written for engineers, a privacy policy written for litigation, and a procurement notice written for vendors are all publications and none of them are disclosures. The test is whether a parent, a ratepayer, or a person whose claim was denied can read it and know what was done to them. In advance is the part that gets negotiated away. Disclosure after deployment is a press release. Disclosure before deployment is a decision point, which is precisely why it is resisted. Who: the public body that procures, by resolution; the vendor, as a condition of the contract; the state agency that maintains the approved-vendor list. If not: the deployment is unlawful, the contract is void and unenforceable, and no public funds may be disbursed under it. 2. No self-certification. Where a consequential system is evaluated, the evaluation is performed by a party with no financial interest in the finding, and the finding is public. The reason this needs its own number is that the first requirement, standing alone, is satisfiable by a vendor grading its own homework — and that is the ordinary practice, not the abuse case. The published study on essay-detection false positives among non-native English writers was independent, which is exactly why it was inconvenient; a vendor-commissioned version would not have found what it found.161 Independence has a practical definition I would put in the statute: the evaluator is not the developer, not the deployer, not a subsidiary or affiliate of either, not compensated on a basis contingent on the result, and not bound by a contract term restricting publication of an adverse finding. Add to that a right of vetted access for accredited researchers and working journalists, because a right to audit that only regulators may exercise is a right that lapses whenever a regulator is captured, defunded, or simply out of staff.

Who: the auditor, licensed and liable; the procuring body, which must accept no other form of assurance; the legislature that funds the capacity. If not: the system may not be relied upon for any determination about a specific person, and refusal of access is itself the violation. 3. Schools buy training and judgment, or they buy nothing. Published subgroup performance before the contract is signed. Funded, paid teacher preparation inside the procurement budget rather than out of the professional-development line that gets cut in March. Human authority, with a name and a working address, over every placement, discipline, and eligibility decision about a child. No sale, licensing, or secondary use of a minor's data, ever, with statutory deletion deadlines and certified compliance. I give the schools their own requirement rather than folding them into the first two for a reason that is not sentimental. A school is the only institution in this book that acts on a person who cannot consent, cannot exit, cannot hire counsel, and will carry the determination for decades. Every asymmetry the rest of this book describes is at its maximum there. And the deficit is measurable: Gallup found that a substantial majority of teachers had received no formal guidance at all on the use of these tools.162 A system deployed on children, operated by adults nobody trained, with no published error rate and no appealable human, is not an educational technology decision. It is an unconsented experiment with a purchase order. Who: the school board, by recorded public vote; the superintendent, personally; the state board of education, in graduation requirements and on the vendor list. If not: the contract terminates, the savings are clawed back into staffing, and the

decisions made under it are vacated on request. The two that cannot be given up Each chapter in this part ended by naming the pair it would keep if it could keep only two, and I want to set those six side by side, because they turn out to be the same claim in three costumes. Chapter One kept a funded watcher and a preserved authenticated record. Chapter Two kept a named human with authority to reverse a decision about a child, and a staffing floor that efficiency gains may not be booked against. Chapter Three kept public records in public custody, and no infrastructure decision on a sealed factual basis. Read them together and they reduce to one sentence, which is the whole of Part One and the only thing on this page I would defend under any circumstances: Someone outside the institution must be able to see the thing, and someone inside it must be answerable for it by name.

Everything above is machinery for producing those two people. The watcher, the auditor, the researcher, the journalist, the parent reading the error rate — all one role, seen from different doorways. The records officer, the superintendent, the commissioner casting a recorded vote, the named appealable human — all one role, seen from different chairs. Where both exist, a mistake is a problem. Where either is missing, a mistake is a condition. What Part One does not settle I will not pretend this part is sufficient, because it is not, and the gap is structural rather than a matter of adding more items. Every requirement above presumes a place to bring it. A board that meets. A commission with a docket. A court that will hear a case about a school district. A comment period that closes after the notice arrives rather than before. Perception protects your ability to know that something was done to you and to prove it later. It does not, by itself, give you anywhere to go. That is the subject of what follows. Part Two is about will — the capacity to decide and to act — and about the two ways it goes wrong: force that supplies its own justification, and emergency that never ends. The chapters there ask who decided, at what speed, under what claim of necessity, and whether there remains a room in which the decision can be argued with. A society that can see clearly and has nowhere to bring what it sees has not been protected. It has been made a very well-informed spectator.

What Perception Requires

1. The record must be obtainable. A citizen should not need wealth, specialized credentials, or a personal connection to discover what public power is doing.

2. The record must be independently checkable. A claim cannot be verified only by the institution that made it.

3. The record must be durable. Public memory cannot depend solely on a platform, vendor, political administration, or private archive.

4. Synthetic material must be identified. Disclosure of how a public artifact was made protects deliberation without authorizing government to decide what people may believe.

5. The watcher must be supported. Local journalism, libraries, archives, educators, researchers, and public-interest institutions are civic infrastructure, not decorative extras.

6. Education must preserve judgment. The point of learning is not the production of a correct-looking answer; it is the development of a person capable of evaluating one.

A Practice for the Reader

Find one public decision near you. Identify the body that will decide it. Locate the agenda, docket, contract, permit, or meeting record. Save a copy. Ask what source supports its central claim. Ask who is missing from the record. Then tell one other person where to find it.

This is small work. It is also how a public begins to recover the power to see.

Coda: The Right to See Is Not Yet the Power to Act

To see clearly is not the same as being free.

A person may recognize a deception and still have no way to correct the record. A family may understand why a decision was made and still be unable to reach anyone with authority to reconsider it. A community may discover the truth only after the contract is signed, the permit issued, the system deployed, or the harm already underway.

Perception restores reality to public life. But truth that cannot be acted upon becomes another form of powerlessness.

The right to see must become the power to act. That is the work of the next part.

NEXT: PART TWO — WILL

POLITICS AND DEMOCRATIC PRACTICE

Governing at Machine Speed

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."

