A CITIZEN'S MANIFESTO FOR THE AGE OF AI

PART FOUR — EXCHANGE



WORK, WAGES, AND FINANCIAL EQUALITY

The distribution question is the whole question

A CITIZEN'S MANIFESTO FOR THE AGE OF AI

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

PART FOUR — EXCHANGE

WORK, WAGES, AND FINANCIAL EQUALITY

The Distribution Question Is the Whole Question

DAVID M. DOYLE

The Series Charter

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

Part Four concerns Exchange: the fair distribution of gains, costs, risks, labor, and public benefit. Every system distributes something—wealth, risk, time, opportunity, security, labor, and power.

The movement of the six-part manifesto is deliberate:

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

The Three Tests

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

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

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

How to Read Part Four

If automation increases productivity, workers should not bear the instability alone. If public resources help build an industry, the public should share in its gains. If a community supplies land, water, power, labor, tax support, and infrastructure, it deserves enforceable return.

Exchange asks who benefits, who pays, and who is left carrying the consequences.

The central question is not whether a system creates value. It is for whom, at whose expense, and under whose terms?

PART FOUR — EXCHANGE

Reciprocity, provision, and the honest accounting of costs

Its corruption: dependency, extraction, and debt as a form of government.

Contents

Chapter Ten — Work, Wages, and Financial Equality: The distribution question is the whole question

Chapter Eleven — Poverty and the Automation of Scarcity: Systems that decide who is deserving

Chapter Twelve — Provision and Movement: Food, logistics, and the systems that get things to people

Part Opening

Between 1987 and 2015, labor productivity in the American computer and peripheral equipment industry rose by an average of 16.0 percent a year.

Real compensation per hour in that same industry fell by an average of 3.4 percent a year.

This is not an activist's statistic; it is a table published by the Bureau of Labor Statistics, in an article whose own summary notes that productivity outpaced compensation in 83 percent of the 183 industries it studied.568 Hold that beside two more facts.

On the fourth of July, 2025, the United States enacted a law that made permanent a 100 percent first-year depreciation deduction for qualifying equipment acquired after the nineteenth of January of that year — the full cost of a machine, deducted in the year it is bought.569 Treasury and the Internal Revenue Service confirmed the terms eighteen months later, describing it as a "permanent 100% additional first year depreciation deduction."570 There is no corresponding provision for hiring a person.

Payroll is taxed.

Equipment is expensed.

The tax code is a statement, in the only language a balance sheet reads, about which input a society would prefer you to buy.

And on the far side of the same ledger: in 2021 the Federal Court of Australia recorded that a government had used an automated system to assert debts "totalling at least $1.763 billion against approximately 433,000 Australians," that the Commonwealth "admitted that it did not have a proper legal basis to raise, demand or recover" them, and that the proceeding "exposed a shameful chapter in the administration of the Commonwealth social security system and a massive failure of public administration."571 Three facts, one subject.

The machine's gain is recognized immediately and in full.

The worker's share is a lagging indicator.

And when the accounting runs the other way — when the state believes a citizen owes it money — the demand can be automated, issued at scale, and enforced years before anyone is made to prove it was ever owed.

Let me say what I am not arguing, because Part Four is the part most easily mistaken for a complaint about commerce, and it is not one.

Exchange is not extraction.

Exchange is how a society feeds, moves, employs, insures, and houses itself, and I am not aware of any mechanism other than a market that does those things at scale without producing queues, favouritism, and rot.

Prices carry information.

Profit is the signal that somebody wanted the thing.

I have raised capital, met payroll, and priced products, and I have no interest in a book that pretends the productivity of the last two centuries was a trick.

The discipline of exchange is narrower and older than any of that: the party receiving the benefit also carries the cost, and the accounting is honest.

A firm that captures a gain and leaves the cost with a ratepayer, a watershed, or a laid-off workforce has not out-competed anyone; it has moved an entry from one column to another and called the difference profit.

A benefits system that automates denial while leaving provision manual has not become efficient; it has decided which of its two possible errors it can live with.

Part One asked whether you can see.

Part Two asked whether anyone can act, and whether the actor can be reached.

Part Three asked what is permanently altered and who pays for it.

Part Four asks the household question: if the machine makes the work cheaper, why is my life not cheaper?

The mechanism has two halves, and they are running at different speeds.

The first half is the asymmetry of the ledger.

Gains from automation land where the tool is owned, immediately and in cash.

Costs land where the work was done, slowly and in ways that never appear as a line item — a title that now pays less for the same duties, a schedule set by software, a pension exposed to the equity of the four firms that own the models.

Nobody has to conspire for this to happen.

Until 2025 the structure at least made a firm spread the deduction for a machine over the machine's useful life.

Now it does not.

The second half is coordination without agreement.

This is the part that is genuinely new, and the clearest record of it is in a federal courthouse in North Carolina.

The Justice Department, with eight state attorneys general, alleged that a single software vendor — holding roughly 80 percent of the market for commercial revenue-management software used to price apartments — "contracts with competing landlords who agree to share with RealPage nonpublic, competitively sensitive information about their apartment rental rates and other lease terms to train and run RealPage's algorithmic pricing software," and that the software "then generates recommendations … based on their and their rivals' competitively sensitive information."572 An amended complaint added six of the country's largest landlords, "more than 1.3 million units in 43 states and the District of Columbia."573 In November of 2025 the Department filed a proposed settlement requiring the vendor to "cease having its software use competitors' nonpublic, competitively sensitive information to determine rental prices in runtime operation" and to limit model training to nonpublic data "aged for at least 12 months."574 Read what that settlement concedes about the world before it.

Competitors did not need to meet, and no one needed to agree to anything, because the model was the meeting.

Price coordination became a subscription.

And the remedy — necessary, and I am glad of it — is a data-freshness rule, which tells you how thin the line now is between a market and a cartel that shares a vendor.

Here is the pattern I would ask you to recognize, because it will run through all three chapters of this part.

The price stops being a fact about the good and becomes a fact about you.

In January of 2025 the Federal Trade Commission published initial findings from a study of six intermediaries that had worked with at least 250 clients, reporting that "retailers frequently use people's personal information to set targeted, tailored prices for goods and services—from a person's location and demographics, down to their mouse movements on a webpage," and that a consumer "profiled as a new parent may intentionally be shown higher priced baby thermometers on the first page of their search results."575 That is the corruption of exchange stated precisely.

In a market, I pay what the thing costs plus what the seller can get.

In this arrangement I pay what my situation reveals about my capacity to refuse — which means the person with the least room to shop, compare, wait, or walk away is charged the most.

Dynamic pricing on a hotel room is an inconvenience.

The same logic applied to insulin, electricity, rent, or a fare is not pricing.

It is a tax on desperation, levied privately, and disclosed to no one.

A CITIZEN'S MANIFESTO FOR THE AGE OF AI 197 On the public side, the same logic wears the opposite face.

The Australian scheme averaged annual tax data across fortnights, assumed a steady income that casual and intermittent workers do not have, and generated a demand — and the burden of disproving it fell on the person least equipped to carry it.

Even the remedy shows the asymmetry: refunds were confined to debts raised using "averaged ATO income information," which meant that being wronged was not enough; you had to have been wronged by the specific

method the Commonwealth had conceded.576 The Royal Commission that examined all of this reported in July of 2023.577 Why does this belong in a book about artificial intelligence?

Because the largest productivity claim in a century is arriving at exactly the moment when every institution that once transmitted productivity into wages has been weakened — union density, sectoral bargaining, a tax code neutral between labour and equipment, the assumption that a job's title fixes its duties.

And because the same computational capacity that might make provision cheap is being applied first, and much more competently, to the machinery of exclusion: to deciding who is deserving, what you can be charged, and how quickly a debt can be asserted against you.

I want to be honest about the limit of my own claim.

The BLS series ends in 2015 and covers a period in which many things other than automation moved wages, and no table proves causation.

The direction, though, is not seriously contested by anyone who has read it, and the burden of proof properly sits with whoever wants to argue that this time the gain will reach the ground on its own.

The chapters that follow ask four questions of any economic arrangement built on these systems, and I would ask you to carry them into any negotiation where efficiency is being promised: Who receives the gain.

Who carries the cost.

Whether the accounting is public.

And whether the person on the wrong end of an automated determination can get it reversed before it costs them the rent.

Here is the standard I would hold to, stated as one test: an exchange is legitimate to the extent that the party capturing the benefit also carries the cost under a published accounting; that no price is set from a person's revealed inability to refuse; and that no adverse determination takes money or benefit from a household before a named human being has confirmed it and a real appeal has been available.

Each clause is falsifiable, which is why I state it that way.

Every arrangement in the three chapters that follow — the ratepayer carrying a private load, the fossil plant with no bonded retirement date, the rent set by a shared model, the debt raised by an average — fails at least one clause, and each of them is defended in public as efficiency. — D.M.D.

DAVID M.

DOYLE 198

Chapter Ten: Work, Wages, and Financial Equality

The Distribution Question Is the Whole Question

The Distribution Question Is the Whole Question

If the machine makes the work cheaper, why is my life not cheaper?

I.

On the third of September, 2026, the Bureau of Labor Statistics published a routine quarterly release. It carried a document number, USDL 26-1434, and an embargo notice, and it was reported for a day and forgotten.578 Here is the sentence I believe will be quoted for a long time. "The labor share, which is the percentage of output that accrues to workers in the form of compensation, was 52.8 percent in the second quarter of 2026, the lowest level in the series, which begins in the first quarter of 1947."579 The share of everything the American economy produces that reaches the people who produce it is the lowest it has been since the government started counting, and the government started counting in the year the Marshall Plan was drafted. The sentence immediately before it reads: "Real hourly compensation, which takes into account consumer prices, decreased 3.3 percent in the second quarter of 2026 and decreased 0.1 percent over the last four quarters."580 And four sections down, concerning the nonfinancial corporate sector: "Productivity increased 2.2 percent, as output increased 3.9 percent and hours worked increased 1.7 percent. Productivity increased 3.1 percent over the last four quarters. Unit profits of nonfinancial corporations increased at an annualized rate of 43.0 percent in the second quarter of 2026, the highest rate since the second quarter of 2021. Over the last four quarters, unit profits increased 17.8 percent, the highest rate since the fourth quarter of 2021."581 I will not claim that artificial intelligence caused this. The labor share has been falling for decades, and the causes are contested: trade, the decline of union density, the tax treatment of capital, the structure of corporate governance. What I will say is narrower. The largest claimed productivity revolution in a century is arriving into an economy whose distributive machinery is already at its weakest recorded setting. Whatever the machine produces will be distributed by that machinery. And on the current record, that machinery is not distributing.

This is the distribution question, and in my view it is the whole of the AI debate. A technology that doubles output and delivers none of the gain to the people whose work it replaces has not made the country richer. It has made a smaller number of people richer and called the result a country.

II.

The Bureau publishes both series on a common base, 2017 equals 100. In the second quarter of 2026, nonfarm business labor productivity stood at 120.0. Real hourly compensation stood at 109.2.582 Output per hour has risen twenty percent since 2017. The purchasing power of an hour of work has risen nine. The American worker produced a fifth more and took home a tenth more. The remainder went to unit nonlabor payments, which stood at 144.0 on the same base, and to the forty-three percent annualized rise in unit profits.583 Nor is this technological stagnation. "During the current business cycle, starting in the fourth quarter of 2019, labor productivity has grown at an annualized rate of 2.1 percent through the second quarter of 2026 … The 2.1-percent annualized rate of nonfarm business productivity growth in the current business cycle thus far is higher than the 1.5-percent rate of the previous business cycle."584 So I want to retire one argument before it is made:

that wages will rise when productivity rises, and that the answer to stagnant pay is more technology, more investment, more automation. We have had more technology. Productivity growth in this cycle is running above the previous cycle's. The labor share hit its lowest recorded level in the same quarter. The transmission belt between the two is broken, and no amount of additional productivity repairs a broken belt. It only spins it faster.

III.

Where is the machine closest to the work? The Bureau measures that too, industry by industry. In the release on service-providing industries published on the twenty-sixth of August, 2026, covering calendar 2025: "Productivity ranged from a 10.8-percent decline in wired telecommunications carriers to a 12.9-percent gain in software publishers."585 "Productivity exceeded 8.0-percent growth in three industries: software publishers (+12.9 percent), line-haul railroads (+11.3 percent), and wireless telecommunications carriers (+8.8 percent). Output increased in all three industries while hours worked decreased." Output up, hours down. That is the signature. The question a citizen should ask is what happens next to the hour that was removed, and to the person who was working it. DAVID M. DOYLE 200 The pattern repeats. In wholesale and retail trade: "In wholesale trade, output grew 3.1 percent and hours worked decreased 1.2 percent. In retail trade, output grew 2.5 percent and hours worked decreased 0.4 percent."586 And in one line the Bureau records something close to a parable: vending machine operators "posted its largest productivity gain on record (+18.2 percent). Output decreased (-0.7 percent) and hours fell more sharply (-16.0 percent)." Productivity at a record, the gain produced entirely by removing one hour in six. These numbers do not isolate artificial intelligence; the Bureau does not attribute industry productivity changes to particular technologies, and I will not put that attribution in its mouth. They show the shape of the thing. Where output rises while hours fall, someone has been made surplus, and the statistical system records the surplus as an improvement.

IV.

I want to turn from what the economy did to what the law rewards, because the two are not independent. There are three ways a company can spend a dollar. It can buy a machine. It can pay a person. It can hand the dollar back to its shareholders. American tax law prices those three choices very differently, and in 2025 the difference was widened and made permanent. Buy a machine. Under Internal Revenue Code section 168(k), as amended in July 2025: "the depreciation deduction … for the taxable year in which such property is placed in service shall include an allowance equal to 100 percent of the adjusted basis of the qualified property."587 The full cost, deducted in year one. This was not new in 2025, but it had been scheduled to phase out; the 2025 law struck the phase-down and the expiration. The Code's own amendment notes record the surgery: the statute "substituted '100 percent' for 'the applicable percentage'" and "struck out cl. (iii) which read as follows: 'which is placed in service by the taxpayer before January 1, 2027.'"588 The same law added section 168(n) allowing the same immediate write-off for the structure itself — the factory or plant — for property whose "construction … begins after January 19, 2025, and before January 1, 2029."589 And it added section 174A: "Notwithstanding section 263, there shall be allowed as a deduction any domestic research or experimental expenditures which are paid or incurred by the taxpayer during the taxable year."590 The statute adds: "any amount paid or incurred in connection with the development of any software shall be treated as a research or experimental expenditure."591

So: the server is deductible in full, in year one. The building housing it is deductible in full, in year one. The software running on it is deductible in full, in year one. Now pay a person. Wages are deductible too — section 162(a)(1) allows "a reasonable allowance for salaries or other compensation for personal services actually rendered."592 But the same dollar triggers a tax the machine does not. Section 3111 imposes "on every employer an excise tax, with respect to having individuals in his employ, equal to 6.2 percent of the wages," and a further 1.45 percent for hospital insurance.593 Seven and sixty-five hundredths percent, levied for the act of employing a human being. The worker pays the same again out of the other side of the same dollar.594 I understand the history: those taxes fund Social Security and Medicare, and I am not proposing to defund either. But name the structure honestly. The machine is subsidized at one hundred percent in year one and the worker is taxed at 7.65 percent at the moment of hire, and we have done it while asking why firms are substituting machines for workers. And the third choice, handing the dollar to shareholders? Section 4501: "There is hereby imposed on each covered corporation a tax equal to 1 percent of the fair market value of any stock of the corporation which is repurchased."595 One percent. One percent to return capital. Seven and a half percent to employ a person. A full and immediate deduction to buy a machine. A citizen does not need an economics degree to predict what gets bought.

V.

Does the money follow the incentive? Two filings answer in their own words. Alphabet's annual report for the year ended the thirty-first of December, 2025: "Capital expenditures, which primarily reflected investments in technical infrastructure, were $91.4 billion for the year ended December 31, 2025."596 The same filing records $45,709 million in stock repurchases and $10,049 million in dividends in the same year.597 Microsoft's annual report for the fiscal year ended the thirtieth of June, 2026: additions to property and equipment of $115,948 million, against $22,271 million of common stock repurchased and $26,445 million of dividends paid.598 The filing explains the jump: "Cash used in investing increased $66.9 billion to $139.5 billion for fiscal year 2026, primarily due to a $51.4 billion increase in additions to property and equipment."599 Two companies. Two fiscal years. Two hundred seven billion dollars into physical plant, and one hundred four billion dollars to shareholders. Neither of those numbers is a wage.

Alphabet's filing closes the circle. "Changes to U.S. tax law enacted on July 4, 2025, allow, among other things, for immediate expensing of domestic research and experimentation costs and accelerated depreciation on eligible capital expenditures, the effects of which are included in operating cash flows for the year ended December 31, 2025."600 The company is telling its investors, in the section where it must tell them the truth, that the tax change improved its cash flow. I take it at its word.

VI.

One more number, and then I will stop counting. The federal minimum wage is $7.25 an hour. The Department of Labor's own history chart records the last change: "Jul 24, 2009 | $7.25 for all covered, nonexempt workers."601 Seventeen years. The iPhone had been on sale for two years. There was no Instagram, no Uber, no Siri, no large language model that could hold a conversation. Every technology this book concerns was invented, commercialized, and deployed while the federal wage floor did not move by a single cent. A floor that has not moved in seventeen years is not a floor. It is a fossil, and its immobility tells you precisely how much political weight the wage side of this ledger carries.

VII.

If the gains are not going to wages, where are they accumulating? The Federal Reserve answers this quarterly. In the Distributional Financial Accounts for the first quarter of 2026, the top one percent of households by wealth held 31.6 percent of total household net worth. The bottom fifty percent held 2.5 percent.602 But net worth is not the number that matters here. The gains from an AI boom accrue, in the first instance and overwhelmingly, to the owners of corporate stock. On that measure: the top one percent of households held 50.2 percent of all corporate equities and mutual fund shares. The bottom fifty percent held 1.1 percent.603 Half of the corporate equity of the United States is held by one household in a hundred. The standard objection is that these figures ignore retirement accounts, that ordinary Americans own stock through their 401(k)s, and that the concentration is an artifact of measurement. It is not. The Federal Reserve defines this series explicitly as "Holdings of corporate equities and mutual fund shares excluding equities and mutual fund shares owned through DC pensions."604 Defined-contribution and defined-benefit entitlements are tracked as separate line items. The Survey of Consumer Finances fills in the participation picture. "As in previous years, participation in the stock market in 2022 rose with usual income groups: 34 percent of families in the bottom half of the distribution held stock, compared with 78 percent of families in the upper-middle group and 95 percent of families in the top decile."605 Within the owning population, the concentration repeats: for directly held stocks, the median holding was $15,000 and the mean was $403,900.606 When a mean is twenty-seven times a median, you are not looking at a distribution. You are looking at a few very large holdings and a long tail of token ones.

VIII.

Now the pension question. American households held about $53.1 trillion in retirement assets at the end of 2025 — the Federal Reserve's figure, covering "[h]ouseholds' retirement assets in tax-deferred accounts, including employer sponsored pension plans, IRAs, Roth IRAs, and annuities."607 Of the pension funds' own financial assets, roughly $10.0 trillion sat in directly held corporate equities and a further $6.0 trillion in mutual fund shares.608 A very large share of American retirement security is now a bet on the stock market. That has been true for forty years; it is the explicit design of the defined-contribution system. The new fact is what the stock market has become. From S&P Dow Jones Indices' own factsheet for the S&P 500: "NUMBER OF CONSTITUENTS: 503." "WEIGHT LARGEST CONSTITUENT [%]: 7." "WEIGHT TOP 10 CONSTITUENTS [%]: 33.6."609 Ten companies out of five hundred and three carry a third of the index. The median constituent is worth $35.6 billion; the largest is worth $3.34 trillion.610 The ten, in order of weight: Apple, Microsoft, Nvidia, Amazon, Meta, Berkshire Hathaway, Alphabet Class A, Broadcom, Alphabet Class C, Tesla.611 Nine are technology, AI-infrastructure, or AI-platform firms. The one that is not is a holding company. The retirement security of a hundred and thirty-five million American households rests, to a third of its equity exposure, on the continued earnings growth of the same handful of firms building the compute, selling the models, and buying each other's capacity. The ordinary worker does not own enough equity to capture much of the upside of the AI boom. But through a pension fund, that same worker is carrying a meaningful share of its downside. That is the worst of both positions. Not enough ownership to benefit. Enough exposure to be hurt.

IX.

I have been describing the machine as an economic actor. I want now to describe it as a supervisor, because for tens of millions of Americans the first experience of artificial intelligence at work was not a chatbot. It was a manager that never sleeps. In October of 2022, the General Counsel of the National Labor Relations Board issued a memorandum to every regional office in the country. Its subject line: "Electronic Monitoring and Algorithmic Management of Employees Interfering with the Exercise of Section 7 Rights."612 It described the practice concretely: "In warehouses, for example, some employers record workers' conversations and track their movements using wearable devices, security cameras, and radio-frequency identification badges. On the road, some employers keep tabs on drivers using GPS tracking devices and cameras. And some employers monitor employees who work on computers—whether in call centers, offices, or at home—using keyloggers and software that takes screenshots, webcam photos, or audio recordings throughout the day."613 It then stated what current law already forbids: "It is clear under extant law that employers violate Section 8(a)(1) if they discipline employees who concertedly protest workplace surveillance or the pace of work set by algorithmic management."614 And it added the point that bears directly on this book: "if employers rely on artificial intelligence to screen job applicants or issue discipline, the employer—as well as a third-party software provider—may violate Section 8(a)(3) if the underlying algorithm is making decisions based on employees' protected activity."615 Its central proposal was a presumption. The General Counsel would urge the Board "to find that an employer has presumptively violated Section 8(a)(1) where the employer's surveillance and management practices, viewed as a whole, would tend to interfere with or prevent a reasonable employee from engaging in activity protected by the Act," with a disclosure remedy if the employer's business need prevailed.616 On the fourteenth of February, 2025, that memorandum was rescinded. The Acting General Counsel's rescission notice lists it among others: "GC 23-02 on electronic monitoring."617 The stated reason was caseload. The federal government's only general statement of enforcement policy on algorithmic management was withdrawn. The same thing happened at the Department of Labor. The Wage and Hour Division had issued Field Assistance Bulletin 2024-1 on AI in the workplace under the Fair Labor Standards Act. The Division's own index now carries this footnote: the Bulletin "was issued pursuant to Executive Order (EO) 14110 … EO 14148 … subsequently revoked EO 14110 on January 20, 2025. Accordingly, FAB 2024-1 is withdrawn."618

Two agencies. Two documents. Both withdrawn within four weeks of each other, in the same season in which the compute-reporting threshold discussed in the previous chapter was revoked. I do not allege coordination. I observe a direction.

X.

Where federal guidance retreated, two states wrote statutes worth reading because they show what a concrete rule looks like. California's warehouse quota law defines the thing precisely: a "quota" is "a work standard under which an employee is assigned or required to perform at a specified productivity speed, or perform a quantified number of tasks, or to handle or produce a quantified amount of material, within a defined time period and under which the employee may suffer an adverse employment action if they fail to complete the performance standard."619 It requires written disclosure of every quota on hire.620 It forbids a quota "that prevents compliance with meal or rest periods, use of bathroom facilities, including reasonable travel time to and from bathroom facilities, or occupational health and safety laws."621 And it gives the worker a right to the data held

about them: on request, the employer must provide "a written description of each quota to which the employee is subject and a copy of the most recent 90 days of the employee's own personal work speed data," within twenty-one days.622 That last provision is, to my mind, the single most important sentence in American algorithmic-management law. It converts the worker from the object of the measurement into a party who can read it. The law is enforced by a rebuttable presumption of retaliation for adverse action within ninety days of a data request,623 an injury-rate trigger for investigation at 1.5 times the industry average,624 and a private right of action for injunctive relief.625 Oregon has since adopted its own version.626 That the quotas were producing injury is not a matter of inference. In February 2023 OSHA announced citations at six Amazon warehouses, finding that the company "exposed warehouse workers to a high risk of low back injuries and other musculoskeletal disorders," related to the "[h]igh frequency with which employees must lift packages," the "[h]eavy weight of the items," awkward twisting and bending, and "[l]ong hours required to complete assigned tasks."627 The agency added that it "reviewed on-site injury logs required by federal law and discovered that, in fact, Amazon warehouse workers experienced high rates of musculoskeletal disorders."628 The proposed penalty for the three facilities named was $46,875.629

In December 2024 the matter settled corporate-wide, ten cases on the eve of trial, with Amazon agreeing to "assess ergonomic risk across its facilities, including through annual updates, and investigate and implement controls to reduce ergonomic risk," and to pay $145,000.630 Whatever else the American system of workplace enforcement is, it is not priced to deter. The penalty is a rounding error on a rounding error.

XI.

Now to the moment before the job: the screen. In 2022 the EEOC sued a tutoring company over its hiring software. "In 2020, iTutorGroup programmed their tutor application software to automatically reject female applicants age 55 or older and male applicants age 60 or older. iTutorGroup rejected more than 200 qualified applicants based in the United States because of their age."631 It settled for $365,000 and five years of monitoring. That case was simple because the rule was written into a form. The harder question is what happens when the rule is learned. In Mobley v. Workday, the plaintiff alleges that a widely used applicant-screening system discriminates on the basis of race, age, and disability. In May 2025 the court granted preliminary certification of a nationwide collective: "The proposed collective is similarly situated because Mobley has substantially alleged the existence of a unified policy: the use of Workday's AI recommendation system to score, sort, rank, or screen applicants."632 The common question survives the variation between employers: "that issue is susceptible to common proof—it cuts across the proposed collective, regardless of the degree to which particular employers place weight on those recommendations."633 The court named the injury in a phrase I find exact: the members "were allegedly required to compete on unequal footing due to Workday's discriminatory AI recommendations."634 The collective was defined as "[a]ll individuals aged 40 and over who, from September 24, 2020, through the present, [] applied for job opportunities using Workday, Inc.'s job application platform and were denied employment recommendations."635 Then comes the number, offered by the defendant as a reason not to allow notice: "Workday states that '1.1 billion applications were rejected using Workday' during the time period at issue and 'any notice would still invite potentially hundreds of millions of potential plaintiffs to file their claims in this case.'"636

One point one billion rejections, through one vendor's platform, in under five years. It is the scale at which a single privately designed screening function now stands between Americans and employment — not a law

anyone voted on, but a product feature licensed by employers. In March 2026 the court let the age-discrimination claim proceed, rejecting the vendor's argument that applicants cannot bring disparate-impact claims under the age statute.637 Nothing has been proven. I am not asserting that the system discriminates; that is what the litigation will determine, and it may determine the opposite. What is already established is the architecture. A single vendor's model sits in the path of over a billion applications. The employers using it did not build it. The applicants rejected by it are told nothing about it. And whether it may be challenged at all under the civil rights statutes of the 1960s — written for personnel officers, not for models — has required three years of litigation and is still being worked out. One further detail deserves mention. When the plaintiffs sought the underlying data, a magistrate judge concluded that "plaintiffs have not met their burden of showing that Workday has actual possession of its customers' data where it resides on segregated, password-protected customer tenants that Workday is contractually prohibited from freely accessing."638 The data needed to prove discrimination by the model sits in thousands of customer accounts that the vendor says it cannot reach. The architecture that delivers the screening at scale also disperses the evidence of what the screening did.

XII.

Against all this, what have workers won when they had the standing to bargain? The Writers Guild's 2023 agreement added an entire article on generative systems. Its operative move is definitional: "because neither traditional AI nor GAI is a person, neither is a 'writer' or 'professional writer' … and, therefore, written material produced by traditional AI or GAI shall not be considered literary material."639 From that follows the economics — machine-produced material "shall not be considered assigned material for purposes of determining the writer's compensation" and "shall not be considered source material for purposes of determining writing credit."640 And consent runs in the worker's favor: "A Company may not require, as a condition of employment, that a writer use a GAI program which generates written material that would otherwise be 'literary material.'"641

The performers' agreement took a different route, governing the likeness itself. It defines an "Employment-Based Digital Replica" and an "Independently Created Digital Replica," and requires consent that "must be clear and conspicuous" and obtained in a signed writing "that includes a reasonably specific description of the intended use."642 Scanning time is work time; a day spent being digitized is paid as a day.643 The longshoremen went furthest: "There shall be no fully-automated terminals developed and no fully-automated container handling equipment used during the term of this Master Contract."644 Anticipating the lawyers, they defined the term: "The term 'fully-automated' in the sentence immediately above is defined as devoid of human interaction."645 The same agreement bars remotely operated ship-to-shore cranes,646 and provides that signatories "shall not use automation or artificial intelligence or quantum computing for the performance of clerical functions."647 The package-delivery agreement addressed surveillance directly, in the closest thing in American private-sector law to a due-process rule for algorithmic discipline: "No employee shall be disciplined based solely upon information received from GPS, telematics, or any successor system that similarly tracks or surveils an employee's movements unless they engage in dishonesty … A driver's failure to accurately recall what is reflected by the technology shall not by itself be considered dishonesty. The Company must confirm by direct observation or other corroborating evidence any other violations warranting discharge."648 Four agreements, four strategies: define the machine out of the credit, require consent for the likeness, prohibit the technology outright, require human corroboration before discipline. Every one is a rule no federal agency currently imposes. The counter-example is as instructive. The auto workers' 2023 national agreement contains

no provision I could locate using the terms "artificial intelligence" or "algorithm." Technological displacement is still handled through language of an older vintage: when "changes in methods, products or policies would otherwise require the permanent laying off of employees, the seniority of the displaced employees shall become plant-wide."649 A good clause, written for a world of retooled assembly lines, routing the question through seniority rather than any right to know what the system is or how it decides. The lesson is narrower than that unions are the whole answer: every enforceable right over workplace automation currently existing in the United States was bargained for, not legislated. Which means the roughly nine in ten American private-sector workers without a contract have none of them.

XIII.

If bargaining is how these rights are won, the question of who is legally permitted to bargain becomes decisive. That question is being answered in the direction of fewer. In January 2024 the Department of Labor finalized a rule on independent-contractor classification, replacing a 2021 rule with "an analysis for determining employee or independent contractor status that is more consistent with the FLSA as interpreted by longstanding judicial precedent." Its stated stake: "The misclassification of employees as independent contractors may deny workers minimum wage, overtime pay, and other protections."650 In May 2025 the Wage and Hour Division stopped applying it.651 The Bulletin notes that "the 2024 Rule remains in effect for purposes of private litigation"652 — meaning the rule is law for a worker who can afford a lawyer and is not enforcement policy for one who cannot. In February 2026 the Department proposed to rescind it formally, reverting to the 2021 framework in which two "core factors" — control, and opportunity for profit or loss — "typically carr[y] greater weight in the analysis than any other factor."653 The proposal's stated purpose: "The department's proposed rule seeks to protect these workers' entrepreneurial spirit and simplify compliance for American job creators."654 There are genuine independent contractors in America and their independence is worth protecting. But the "control" factor is where the algorithm does its work. When a platform sets the price, assigns the task, measures completion time, scores performance, and deactivates the account, it exercises precisely the control the classification test asks about — while presenting itself as a marketplace. Instead, the proposed rule specifies that requiring compliance with "quality control standards" and similar contractual terms "does not constitute control that makes the individual more or less likely to be an employee."655 Some state courts have been more direct. New Jersey's highest court, applying that state's three-part test, described the decisive question this way: "Prong C of the ABC test broadly asks whether a worker can maintain a business independent of and apart from the employer. If the worker 'would join the ranks of the unemployed' when the relationship ends, the worker cannot be considered independent."656 "Would join the ranks of the unemployed." That is the right question, and it is the one the federal test is moving away from.

XIV.

Now the other end of the labor market — the work the machine cannot do.

In the most recent occupational wage survey, the single largest occupation in the United States is home health and personal care aides: 4,305,810 people, at a median wage of $17.21 an hour.657 Nursing assistants: 1,448,910, at $20.32. Childcare workers: 518,910, at $16.82.658 The all-occupations median is $24.51. In the table of occupations with the most projected job growth through 2035, home health and personal care aides is first: an increase of 847,300 jobs, or 18.1 percent. The median annual wage printed in the same row: $35,800.659 The

Bureau explains the driver: "the aging population is expected to boost demand for a wide variety of healthcare and social services, including home health and personal care services."660 So the occupation that will add more jobs than any other over the next decade — work that is least automatable, most dependent on patience and human touch — pays about half the national average annual wage. This is not an accident of supply and demand. Demand is rising sharply and predictably. Under ordinary market logic the wage should be rising with it. It is not, because the payer is largely public, the rates are set administratively, and the workforce is heavily female, heavily immigrant, and almost entirely unorganized. The wage is not a price discovered in a market. It is a decision, made by public bodies and private operators, about what this work is worth. Which sets up the defining economic fact of the next twenty years. We are pouring capital into automating cognitive work that pays well, while the work that is growing fastest cannot be automated and pays badly. If nothing changes, the machine will displace the accountant and the paralegal and the junior programmer into a labor market whose principal growth opportunity is a job paying thirty-five thousand dollars a year. Every serious discussion of AI and employment eventually arrives at retraining, and a citizen is entitled to ask: retraining into what?

XV.

Is the displacement actually happening? Here I have to be careful, because this is where public argument is loosest and where I found the evidence most equivocal. The clearest statement from a chief executive is Amazon's, in a June 2025 message to employees: "It's hard to know exactly where this nets out over time, but in the next few years, we expect that this will reduce our total corporate workforce as we get efficiency gains from using AI extensively across the company."661

But watch what happened when the reductions were announced. In October 2025: "it will mean an overall reduction in our corporate workforce of approximately 14,000 roles."662 The stated reason was not artificial intelligence. It was structure: "We're convinced that we need to be organized more leanly, with fewer layers and more ownership."663 A January 2026 announcement of roughly 16,000 more roles again cites "reducing layers, increasing ownership, and removing bureaucracy," with no mention of AI. The CEO says AI will reduce the corporate workforce. The announcements reducing the corporate workforce cite bureaucracy. Either the cuts are AI-driven and the announcements avoid saying so, or they are organizational and the earlier message overstated the AI connection. I cannot tell you which from the documents. The delivery company's disclosure is more concrete, precisely because it is filed with the Securities and Exchange Commission where the standard is different. Its January 2026 earnings release describes a "Network of the Future" initiative "intended to enhance the efficiency of our network through automation and operational sort consolidation," and states the result: "We reduced our operational workforce by approximately 48,000 positions, including 15,000 fewer seasonal positions, and closed daily operations at 93 leased and owned buildings during 2025."664 The savings: "approximately $3.5 billion in 2025," with "approximately $3 billion in 2026" expected.665 That filing attributes the reduction to automation, sort consolidation, and a planned volume decline from its largest customer — not to AI, which it does not mention. It is automation in the older sense. But it is the best-documented instance I found of a company stating, under securities-law obligation, the headcount consequence of an efficiency program and the dollars it captured. And here is the structural problem a citizen should understand. The federal statute meant to give workers warning of mass displacement barely covers it. In the only comprehensive federal evaluation I could locate, the Government Accountability Office found that "[a]bout one quarter of the 8,350 plant closures and mass layoffs in 2001 appear subject to WARN's advance notice requirements," and that of those covered, "[e]mployers provided notice for an estimated 36 percent."666 That study is from 2003, and I

could locate no more recent comprehensive federal evaluation. But the statutory thresholds have not changed, and the displacement pattern this chapter describes — rolling, distributed, white-collar, a few hundred roles at a time across many sites — is precisely the pattern those thresholds were never designed to catch. X

VI.

Everything so far concerns income.

I want to close with what happens on the other side of the ledger, where the same technology sets what you pay. Start with rent, because it is the largest line in most household budgets and because here the government actually sued. The Justice Department's amended complaint alleges that a pricing vendor "sells software to landlords that collects nonpublic information from competing landlords and uses that combined information to make pricing recommendations,"667 and quotes the company's own materials describing the product as one that "helps curb [landlords'] instincts to respond to down-market conditions by either dramatically lowering price or by holding price when they are losing velocity."668 A company executive is quoted saying "there is greater good in everybody succeeding versus essentially trying to compete against one another in a way that actually keeps the entire industry down," and that with enough adoption landlords would "likely move in unison versus against each other."669 The government's characterization is the sentence to remember: "RealPage replaces competition with coordination. It substitutes unity for rivalry."670 The complaint alleges the software ingests "nonpublic rental rates, future apartment availability, and changes in competitors' rates and occupancy" daily,671 that the company "controls at least 80 percent of that market,"672 and quotes a landlord customer writing: "I always liked this product because your algorithm uses proprietary data from other subscribers to suggest rents and term. That's classic price fixing." The case settled in November 2025. The consent judgment's terms describe, by negation, how the machine worked. The company must "cease using current or historical Unaffiliated Property Data in the Runtime Operation of any Revenue Management Product."673 It must not offer a product that "includes any Governor Guardrail that does not have symmetrical upper and lower bounds for the recommended change in rental price."674 It must allow users "to set parameters that permit pricing recommendations to go below the Pricing Floor to the same extent such pricing recommendations exceed the Pricing Ceiling."675 And it may not offer a product that "prohibits, disincentivizes, or impedes in any way a user's or licensee's ability to reject or override any price recommendation."676 Read those backwards and you have the design that was alleged: a system fed by competitors' confidential data, with asymmetric guardrails that allowed prices to rise more easily than they fell, and features that discouraged landlords from overriding it. The remedy is symmetry. A CITIZEN'S MANIFESTO FOR THE AGE OF AI 213 New York then wrote the prohibition into statute. It defines an "algorithmic device" as "any machine, device, computer program or computer software that on its own or with human assistance performs a coordinating function," and defines the coordinating function as the three-step sequence: collecting price and occupancy data from two or more unaffiliated owners, "analyzing or processing" it "using a system, software, or process that uses computation, including by using that information to train an algorithm," and "recommending rental prices, lease renewal terms, ideal occupancy levels, or other lease terms."677 That statute regulates the function rather than the label. It does not matter what the product is called or whether the vendor describes its output as a "recommendation." If it collects rivals' data, computes on it, and returns a price, it is covered. The same pattern — data in, price out, individualized — is now documented well beyond housing. The Federal Trade Commission's staff study of surveillance pricing describes tools that let a seller "charge or set prices, or determine quantity and availability for Individual Consumers based on location, purchase history, return history,

customer service history, browsing behavior, demographic information, or any other variables."678 The examples are worth quoting exactly. On new parents: "if a consumer is profiled as a new parent, the consumer may intentionally be shown higher priced baby thermometers on the first page of their in-app search results, based on their residential zip code and time of purchase." The report continues that these tools "could also potentially be used to collect behavioral details that a retailer could use to forecast a customer's state of mind, like using a shopper's selection of 'fast-delivery' shipping on an order of infant formula to infer that a shopper could be a rushed parent who may be less price-sensitive."679 On people in a disaster: "A hypothetical health supplies company that has identified some of its customers as homeowners living in an area that is undergoing flooding could target them with promotional campaigns for stress-relief supplements."680 On a gambler trying to leave: if a customer "demonstrates hesitation by lingering on the homepage longer than expected or moves their cursor towards the button to close out their browser tab, the website may trigger a pop-up showing popular sporting events to incentivize the visitor to remain on the website and place a bet."681 DAVID M. DOYLE 214 What is collected to make this possible: "micro interactions shoppers might have with a virtual storefront, whether they've used their mouse to highlight the name of a particular product on the page, how far down that page they scroll while shopping, or whether they watched a given percentage of an on-page video before closing it."682 From which a seller may infer "a shopper's emotional state, purchase intent, or financial sensitivity," including "whether the consumer is likely to be interested in products eligible for food assistance benefits."683 And the business case, in the respondents' own framing as staff summarized it: the tools support "revenue growth at between 2–5%" and "increases in margins of 1-4%."684 That is the corruption documented by the government's own study: the price ceases to be a fact about the good and becomes a fact about the buyer. A thermometer does not cost more because thermometers are scarce. It costs more because the system believes you are tired, frightened, and will not shop around. Two states have responded, and both responses are instructive for what they exempt. New York now requires a seller using personalized algorithmic pricing to display: "THIS PRICE WAS SET BY AN ALGORITHM USING YOUR PERSONAL DATA."685 Maryland has a parallel disclosure and a substantive ban on dynamic pricing of tax-exempt food.686 But both statutes carve out insurance and financial institutions.687 The two sectors that price on personal data most aggressively, and where the consequences of mispricing are most severe, are exempt. In credit, the rule is at least clear in principle. The Consumer Financial Protection Bureau has stated that lenders "must be able to provide applicants against whom adverse action is taken with an accurate statement of reasons," regardless of the technology, and that "[a] creditor cannot justify noncompliance … based on the mere fact that the technology it employs to evaluate applications is too complicated or opaque to understand." Its formulation is one of the best sentences any American regulator has written about this technology: "A creditor's lack of understanding of its own methods is therefore not a cognizable defense against liability."688 A follow-up circular closed the obvious evasion, holding that creditors "that simply select the closest, but nevertheless inaccurate, identifiable factors from the checklist of sample reasons are not in compliance with the law."689 In insurance, Colorado has gone furthest, prohibiting insurers from using "external consumer data and information sources, as well as any algorithms or predictive models that use external consumer data and information sources, in a way that unfairly discriminates,"690 and placing the burden on the insurer to test and demonstrate that it does not, with an officer's attestation.691

One state. One line of insurance at first, extended in 2025 to auto and health plans.692 That is the entire American structure for algorithmic insurance pricing.

XVII.

I promised in the introduction that this book would judge power by a fixed standard rather than by whether I like the people holding it. The standard is that no authority, public or private, elected or capitalized, may simultaneously control the means of force, of food and material provision, of communication and public truth, of transit and movement, of historical memory, of defining an emergency, and of predicting the future. This chapter has been about provision, and the mechanism by which a person obtains it: work, and the wage that work commands. Let me apply the three questions this book has used throughout. Can it be witnessed? Only partially. The productivity and wage series are public, the tax provisions are public law, the capital expenditures are in audited filings. But the screening model that rejected 1.1 billion applications is not public, and its own vendor says the data needed to test it sits in customer accounts it cannot access. The rent-setting algorithm became visible only because the Justice Department sued and quoted internal documents. The pricing tools became visible only because the Federal Trade Commission compelled six companies to produce materials under a statutory study power. Where the government has compulsory process, we can see. Where it does not, we cannot. That is litigation-dependent visibility, and it arrives years late. Who carries the burden? Not the party receiving the benefit. The firm deducts the machine in full in year one and pays 7.65 percent to employ a person. The worker absorbs a 3.3 percent quarterly decline in real hourly compensation. The displaced worker falls largely outside a notice statute that covered a quarter of separation events and was honored in about a third of those. In every instance in this chapter, the optimization is performed by one party and the cost lands on another. Is there a way back? This is the test the record fails most clearly. An applicant screened out by a model is not told the model exists, let alone given reasons. A tenant priced by a coordinating system had no appeal until the government sued on their behalf. A consumer shown a higher price because the system inferred financial desperation does not know they were shown a higher price. The one bright exception is the warehouse quota statute that gives a worker the right to demand ninety days of their own work-speed data. One state, one industry, one category of data.

Measured against the standard, provision is not yet concentrated in the way communication and prediction already are. But the terms on which an ordinary person obtains provision — the wage, the schedule, the screening, the rent, the price — are increasingly set by systems that cannot be witnessed, whose burden falls on the person, and from which there is no return.

XVIII.

So here is what I would insist on. Not what I predict; what I would refuse to accept the absence of. A right to the reasons, at the moment of the decision. Any system used in hiring, scheduling, promotion, discipline, wage-setting, or termination must, on the request of the person affected, produce a specific statement of the principal factors that produced the outcome, and the data the system used. The credit regulator has already written the correct rule: opacity is not a defense, and a creditor's failure to understand its own model does not excuse it. There is no principled reason a person is entitled to the reasons for a declined credit card and not to the reasons for a declined job. A right to your own work-speed data. Ninety days of it, on request, within twenty-one days, with a presumption of retaliation for adverse action following the request. One state has already written this statute. It should be federal law, it should not be limited to warehouses, and it should cover every metric a system records about a person's pace, idle time, keystrokes, location, or attention. Tax neutrality between labor and equipment. If a firm may deduct the full cost of a machine in the year it is installed, then the employer-side

payroll tax on the first tranche of a new job's wages should be offset to a comparable degree. The present structure is an industrial policy in favor of substituting the first for the second, and it was made permanent in 2025 without ever being debated in those terms. Bargaining rights over the introduction of workplace automation. Every enforceable protection documented in this chapter was bargained, not legislated. The duty to provide information about, and bargain over, the implementation of tracking and automated management technologies should be statutory and should not depend on the enforcement priorities of whoever holds an agency's general counsel office in a given year. A restored and expanded notice standard for displacement. The existing warning statute reaches about a quarter of separation events and is honored in about a third of those. It should be amended to capture rolling and distributed reductions aggregated across a company over a twelve-month period, and the notice should state whether an automated system's deployment contributed to the reduction.

A funded transition, financed before the displacement rather than after it. If a company captures three and a half billion dollars of annual savings from a network automation program, a defined portion should fund retraining, wage insurance, and health coverage continuity for the workers whose positions produced it. Not a voluntary pledge. A condition, sized to the savings, posted in advance. Prohibition of price-setting from revealed desperation. Individualized pricing of essentials — food, medicine, fuel, housing, utilities, and credit — should be prohibited outright rather than disclosed. Where disclosure is the chosen remedy for non-essential goods, the insurance and financial-institution exemptions in both existing state statutes should be removed. And the coordinating-function definition written by New York should be adopted nationally and extended beyond rent to any market where a common vendor's model prices the output of competing sellers. A wage floor that moves. Seventeen years is not a policy. Whatever the right number is, a statutory minimum that cannot move without an act of Congress will not move. It should be indexed, and the indexation should be automatic.

XIX.

I want to tell you how to check every claim in this chapter yourself, because the habit matters more than my conclusions. Read one government release. Go to the Bureau of Labor Statistics' productivity page and open the most recent quarterly release. It is four pages. Find the paragraph containing the phrase "labor share." That paragraph will tell you, for the most recent quarter, what output did, what real pay did, and what share of the result reached workers. You should not accept an intermediary — including me. Read two numbers together. In the same release, compare the labor productivity index to the real hourly compensation index on the same base year. The distance between those two numbers is the distribution question expressed arithmetically. Read one company's cash flow statement. Not the press release — the filing. Find three lines: purchases of property and equipment, repurchases of stock, dividends paid. You will learn more about a company's actual priorities in ninety seconds than from any quantity of corporate communication about its commitment to its people. Look for the absence. When a chief executive says a technology will reduce the workforce, and the announcement actually reducing the workforce does not mention that technology, the gap between those two documents is itself a finding. When a statute requiring price disclosure exempts insurance and banking, the exemption is the story. When an agency withdraws a guidance document and issues no replacement, the empty space on the index page is the policy.

And ask the question at a public meeting. When a project, a contract, or a procurement comes before a body you can attend: who carries the cost if this is wrong, is there a named human who can reverse an adverse

decision, and can an ordinary person obtain the data used against them? If the answer to the third question is that the data belongs to a vendor, you have learned the most important thing about the system in that room.

XX.

The strongest objection to this chapter is not that its numbers are wrong. They are the government's numbers and the companies' own filings. The strongest objection is about causation: the labor share has been declining since the 1970s, the causes are contested, artificial intelligence is too new to have produced a seventy-nine-year low, and to place these facts in a book about AI is to imply a causal story the evidence does not support. I accept much of that. My claim is not that artificial intelligence caused the low labor share. My claim is that the low labor share is the condition into which artificial intelligence is arriving, and that this condition determines what the technology will mean for ordinary people. A distributive system already at its weakest recorded setting, a tax code just made permanently more favorable to equipment than to employment, a wage floor frozen for seventeen years, enforcement guidance withdrawn at two agencies within a month, and a screening architecture standing between a billion applications and employment — that is the machinery into which the largest productivity claim in a century is being fed. A second objection: that this is an argument for slowing the technology, and that slowing it would forfeit real gains in medicine and science. It is not. Nothing I asked for in Section XVIII slows a model's training or restricts a capability. A right to the reasons for a decision, a right to your own performance data, tax neutrality between a person and a machine, a notice standard that captures the layoffs actually happening, and a ban on pricing a thermometer by a parent's exhaustion — none of these makes any system less capable. If a business model cannot survive telling a person why they were rejected, the problem is the business model. I began this chapter with a single sentence from a routine government release: the labor share was 52.8 percent, the lowest since 1947. I want to end with what that sentence describes, because a percentage is an abstraction and the thing it measures is not.

It describes a country in which more is being produced, by fewer hours, for a smaller share. It describes the home health aide earning $17.21 an hour in the fastest-growing occupation in America. It describes the applicant among a billion rejections who was told nothing and could appeal to no one. It describes the tenant whose rent was set by a system that its own vendor's customers recognized, in writing, as classic price fixing. It describes the parent shown the more expensive thermometer because an inference had been drawn about her state of mind. None of those people lack productivity. The country has never had more of it. They lack a share, and a reason, and a way back. Efficiency is not a public benefit until somebody's rent goes down. — D.M.D.

Chapter Eleven

Chapter Eleven: Poverty and the Automation of Scarcity

Whose Error the System Prefers

Why is it so much harder to get help than to be denied it?

I.

Every year the United States Department of Agriculture publishes two numbers about the food assistance program. One is famous. The other is not. The famous one is the payment error rate. For fiscal year 2024 it was 10.93 percent nationally, of which 9.26 points were overpayments and 1.67 points were underpayments.693 This is the number that triggers financial liability for states. It gets discussed as waste. The other is the case and procedural error rate. It measures a different sample: not households receiving benefits, but households that were denied, suspended, or terminated. For fiscal year 2024 it was 43.81 percent.694 In Georgia it was 82.30 percent. In Texas, 66.47. In Arkansas, 61.40.695 Close to forty-four percent of the negative actions reviewed were invalid or procedurally deficient. The government knows this. The government measures it. The government publishes it on its own website, in a PDF, with a state-by-state table. And nothing happens. The reason is written into the regulation. The stated goals of the quality control system include "a basis for establishing State agency liability for payment errors that exceed the National performance measure."696 The liability machinery attaches to money that went out the door wrongly. It does not attach to help that never went out the door at all. We have built capable machinery for detecting one kind of mistake and left the other unpunished. Then we call the resulting savings efficiency.

II.

In March of 1970 the Supreme Court decided Goldberg v. Kelly. The question was whether a state could terminate public assistance benefits without a prior hearing. The Court said no. The syllabus states the balance: "The interest of the eligible recipient in the uninterrupted receipt of public assistance, which provides him with essential food, clothing, housing, and medical care, coupled with the State's interest that his payments not be erroneously terminated, clearly outweighs the State's competing concern to prevent any increase in its fiscal and administrative burdens."697 The holding: "A pre-termination evidentiary hearing is necessary to provide the welfare recipient with procedural due process."698 And its content: "Such hearing need not take the form of a judicial or quasi-judicial trial, but the recipient must be provided with timely and adequate notice detailing the reasons for termination, and an effective opportunity to defend by confronting adverse witnesses and by presenting his own arguments and evidence orally before the decision maker."699 The Court quoted the district court below: "By hypothesis, a welfare recipient is destitute, without funds or assets. . . . Suffice it to say that to cut off a welfare recipient in the face of . . . 'brutal need' without a prior hearing of some sort is unconscionable, unless overwhelming considerations justify it."700 Brutal need. That is the standard the American constitutional order set for itself in 1970, before anyone had automated anything. And this: "In almost every setting where important decisions turn on questions of fact, due process requires an opportunity to confront and cross-examine adverse witnesses."701 Nothing in this chapter asks for a new right. The law has demanded it for fifty-

six years. The question is what happens to that demand when the adverse witness is a scoring model, the decision maker is a batch process, and the notice is a status code in a portal the claimant has no reason to open.

III.

In October of 2013 the State of Michigan brought online the benefits component of a system called the Michigan Integrated Data Automated System. The state auditor's own description is worth quoting for its blandness: "In October 2013, UIA implemented the benefit section of its new computer system, Michigan Integrated Data Automated System (MiDAS). MiDAS provides for more efficient processing of claims than UIA's predecessor system."702 Here is what the efficiency produced. The Michigan Supreme Court summarized the allegations in Bauserman v. Unemployment Insurance Agency: claimants "specifically alleged that defendant had used an automated fraud-detection system — the Michigan Integrated Data Automated System (MiDAS) — to determine that plaintiffs had received unemployment benefits for which they were not eligible and then garnished plaintiffs' wages and tax refunds to recover the amount of the alleged overpayments, interest, and penalties that defendant had assessed without providing meaningful notice or an opportunity to be heard."703

Justice Cavanagh, writing for the Court: "Plaintiffs allege that defendant, Michigan's Unemployment Insurance Agency (the Agency), adjudicated allegations of fraud, seized plaintiffs' tax returns, and imposed penalties on plaintiffs without providing meaningful notice or an opportunity to be heard in violation of Michigan's constitutional right to due process, Const 1963, art 1, § 17."704 The State eventually settled a related civil rights class action for twenty million dollars. Its own Attorney General's office described the instrument as an "auto-adjudication system" and said it had falsely accused "thousands of Michiganders" of unemployment fraud.705 A figure circulates in connection with this episode: that roughly ninety-three percent of the automated system's fraud determinations were wrong. The Michigan Supreme Court's opinion recites it, but its citation is a newspaper article, not a state document.706 I could not locate the underlying agency study, and I am not going to build an argument on it. What I can put weight on is the state's own audit. Between October 1, 2013 and March 31, 2015, the agency "issued 60,324 (re)determinations finding intentional misrepresentation on 47,350 claims."707 The auditors pulled a sample of sixty. In forty-six, the agency "could have improved its efforts to contact" claimants who had not responded to requests for information. And the agency "did not inform claimants that, absent new information provided by another source, failure to respond to the requests for information would result in a finding of intentional misrepresentation. Instead, the requests stated that failure to respond would result in a (re)determination based on available information."708 Silence was the input. Fraud was the output. The transformation happened inside a system that was described, accurately, as more efficient.

IV.

If you want to know what a notice about an algorithmic decision ought to say, the best answer I have found in American law is two pages long and was signed by a federal district judge in Jonesboro, Arkansas, on the first of November, 2016.

The background: Arkansas had moved its home-care allocation for disabled and elderly Medicaid beneficiaries onto an assessment instrument and a scoring methodology. By the state's account: "Before 2013, DHS used a paper form to determine eligibility for the department's home- and community-based Medicaid waiver services for adults with physical disabilities and the frail elderly. In 2013, DHS began to use the ArPath Assessment tool

for determining eligibility only." Then: "The Centers for Medicare and Medicaid Services (CMS) approved the ARChoices waiver with the understanding that DHS would begin to use an objective tool to allocate attendant care hours for beneficiaries. … DHS then began using this new tool, known as the Resource Utilization Groups (RUGs) methodology, in conjunction with the ArPath Assessment tool to determine the number of attendant care hours ARChoices beneficiaries receive."709 Ethel Jacobs's attendant care hours were cut. She sued. Judge D.P. Marshall Jr. entered a permanent injunction, and I quote the operative passage nearly in full: "Jacobs's benefits may not be reduced unless and until DHS completes a new assessment and provides notice that satisfies the Due Process Clause and the Medicaid Act. The notice must be in plain and clear terms. The notice must explain why benefits are being reduced. The notice must be accompanied by a copy of the beneficiary's ARPath assessment. The notice must explain the benefit reduction with specific references (as applicable) to the beneficiary's ARPath assessment, the beneficiary's Resource Utilization Group, and the ARChoices system, including the algorithm. The notice must be as specific as reasonably practicable about the beneficiary's health conditions and reduced benefits. Finally, the notice must alert the beneficiary about any applicable limitations on the presentation of evidence at any later hearing about the benefits decision."710 Including the algorithm. A federal judge, in a two-page order, wrote the disclosure standard that much subsequent policy debate has been groping toward. Not a summary. The assessment itself, the specific group assignment, the system, the algorithm, in plain and clear terms, plus advance warning about what evidence the beneficiary will and will not be permitted to present on appeal. Separately, in state court, Arkansas beneficiaries challenged how the scoring methodology had been adopted at all. The court ruled that DHS had not substantially complied with notice requirements in adopting the RUGs methodology, and barred its use until the rule was properly promulgated.711

A limit on my sourcing: the Arkansas Supreme Court's opinions in that appeal are available online, so far as I could find, only on commercial legal databases; I have therefore relied on the Arkansas agency's own public description of the rulings rather than quoting the appellate opinions.712

V.

Idaho ran the same experiment with a different instrument, and the resulting case is the deepest judicial treatment of a proprietary scoring tool I know of. The class: "Plaintiffs are developmentally disabled adults who qualify for benefits under Medicaid. They are eligible for long-term institutional care but choose to live instead in their own homes or in community settings. Over a decade ago, when their Medicaid payments were reduced, plaintiffs sued the Idaho Department of Health & Welfare."713 In March 2016 the court granted partial summary judgment, finding "that: (1) the budget tool's unreliability and arbitrary nature violated plaintiffs' right to due process of law; (2) the notices utilized by IDHW were inadequate to satisfy due process; and (3) to ensure due process, IDHW must 'receive a commitment from a suitable representative to assist the participant before proceeding to informal review and taking any action to confirm a budget reduction produced by the budget tool.'"714 Years later the state wanted to keep the successor instrument's user's manual from participants. Magistrate Judge Candy W. Dale refused: "Without access to the User's Manual, the SIS-A assessment remains a black box for participants, where they see only certain inputs and outputs, but not the internal workings (e.g., the algorithm applied in the SIS-A software). Therefore, as was the situation without disclosure of the SIB-R booklets, precluding participant access to the User's Manual prevents the participant from challenging errors or effectively cross-examining IAPs who claim their assessments are accurate. In essence, it would be impossible for the participant to prove error in the administration of the SIS-A."715 And the principle underneath it: "the

very foundation of the adversary process assumes that use of undisclosed information will violate due process because of the risk of error."716 On March 31, 2025 — more than nine years after the court held the original budget tool arbitrary — the court found the department in civil contempt and ordered it to pay class counsel's monitoring fees, stating that "a partial award is necessary to incentivize IDHW to comply with the Settlement Agreement."717

Nine years, a class action, a constitutional ruling, a settlement, and a contempt finding, and the replacement tool still was not in place. If that is what winning looks like, we should be honest that the remedy is not working at the speed that human lives run.

VI.

The cases above are about individual determinations. Now consider the same logic at population scale. When the pandemic-era continuous enrollment condition for Medicaid ended, states resumed eligibility redeterminations. Across renewals due from March 2023 through June 2024: "Total Renewals Due 94.3 million," "Individuals Disenrolled from Medicaid or CHIP Coverage 20.8 million (22% of total due)."718 Of those 20.8 million: "Determined ineligible for Medicaid or CHIP based on the return of a renewal form 6.5 million (7% of total due) (31% of total terminated)" "Disenrolled for procedural reasons 14.3 million (15% of total due) (69% of total terminated)"719 More than two out of every three people who lost coverage lost it for paperwork reasons, not because anyone determined they were ineligible. The agency itself is careful here. Its caveats note that some states were holding procedural terminations as a mitigation strategy, and that "some states have restored eligibility for people initially found ineligible due to systems or other errors," which "is frequently not identifiable in the data."720 A procedural termination is not by itself proof that the person was eligible. But we do not need to guess, because the agency identified a specific automated defect throwing eligible people off. In a letter to all State Medicaid Directors dated August 30, 2023, it wrote: "More recently, however, CMS has learned of additional systems and operational issues affecting multiple states, which may be resulting in eligible individuals being improperly disenrolled. Specifically, we understand that some states are conducting ex parte renewals at the household level, without regard to differing eligibility statuses and income thresholds for individuals within the household. As a result, while a state may have sufficient information during the ex parte process to renew Medicaid or CHIP coverage for some individuals in a multi-member household, states are sending renewal forms requesting information for all household members, and, if the renewal form is not returned, states are disenrolling all individuals in the household, including those who should have been determined to be eligible through the ex parte process."721

Read that as an engineering description. The renewal was evaluated at the wrong unit of analysis: the system asked whether the household had been renewed, when the law requires an answer for each person in it. Adult eligibility in most states sits "at or below 133 percent of the federal poverty level," while "the median Medicaid/CHIP eligibility level for children across states is 255 percent FPL."722 The defect fell hardest on children, eligible on data the state already held, cut off because an adult in the same home did not mail back a form. The agency ordered states to pause procedural terminations, reinstate coverage for affected individuals, and fix their systems.723 Where a state cannot identify the affected individuals, it "must reinstate coverage for the full household."724 By the agency's count as of September 21, 2023, thirty states and territories were not conducting automatic renewals at the individual level or were still working to reinstate affected individuals.725

A clear majority of American jurisdictions were running an automated renewal process that disenrolled eligible people through a unit-of-analysis error no claimant could have detected, contested, or even named.

VII.

The Payment Integrity Information Act defines an improper payment as "any payment that should not have been made or that was made in an incorrect amount (including overpayments and underpayments)."726 The oversight framework sorts them into four buckets: overpayments, which "are payments that, in theory, should or could be recovered"; underpayments, "those in which recipients received less than what was due"; unknown payments; and technically improper payments, "those in which recipients received funds they were entitled to, but the payment failed to follow all applicable statutes or regulations."727 For fiscal year 2023, executive agency improper payment estimates totaled $236 billion, and "most of the $236 billion … consisted of overpayments."728 Every one of those four buckets describes a payment that was made. An eligible person who was wrongly denied, wrongly terminated, or who gave up halfway through an identity check received nothing. Nothing is not a payment. It cannot be an improper payment. It does not appear in the $236 billion. A government that measures only erroneous outflow will optimize only against erroneous outflow. Every threshold, every risk score, every automatic hold pending verification will be evaluated against a metric that counts the cost of a wrongful payment and is structurally blind to the cost of a wrongful denial. We did not decide to be cruel to the poor. We decided what to count, and the counting decided the rest.

To its credit, the food assistance program does measure denials. The negative case review exists to "verify, through documentation contained in the case record, whether the State agency's action to deny, terminate, or suspend the household was valid."729 The rate is defined as "an estimate of the proportion of denied, suspended, or terminated cases where the household was incorrectly denied, suspended, or terminated or where procedural deficiencies exist."730 But the review is a desk exercise that checks whether the file supports the denial, not whether the family was actually hungry and actually eligible.731 And the liability attaches elsewhere. The stated purpose of the system includes establishing "State agency liability for payment errors."732 Not denial errors. So a state with a 10.93 percent payment error rate faces fiscal consequences. A state with an 82.30 percent case and procedural error rate faces a line in a table.

VIII.

There is a newer gate, and it sits in front of everything else: proving that you are you. The federal tax agency contracted identity proofing to a single private credential service provider. The oversight review found that the agency "determined that it needed identity assurance level (IAL) 2," that "a private credential service provider, ID.me, is IRS's sole provider of level 2 identity-proofing products," and that these activities "include having individuals provide evidence, such as a driver's license, and biometric evidence, such as a selfie."733 The scale: "users accessed IAL 2 applications more than 150 million times between 2021 and 2024."734 The reviewers stated the trade-off in a single clean sentence: "If IRS makes the identity-proofing process too stringent, legitimate taxpayers may not be able to successfully prove their identities and access services or account information. Conversely, if the process is too easy for users to pass, then IRS's identity proofing may be more susceptible to a host of security failures including unauthorized access and fraud."735 There is a dial. Someone chooses where it sits. So: where does it sit, and how many people does it exclude? I could not find out. Neither could the auditors. The vendor's performance dashboard includes a "true pass rate," and here is what that

metric does: "ID.me's true pass rate excludes both users who abandon the process and users identified as highly probable fraudulent."736 The headline success metric is computed by removing from the denominator precisely the population the system failed.

The oversight findings compound it. "IRS officials were unable to show us that they had independently documented measurable goals or objectives to manage the outcomes of its identity-proofing program." "IRS officials were unable to show us that they evaluated or documented the outcomes of IRS's digital identity-proofing program" — "instead, they are relying on ID.me's own assessments of its solutions' performance." And: "IRS has not included ID.me on its inventory of programs using AI, limiting IRS's assurance that procurements using AI are used in alignment with IRS policy so that taxpayers' rights are protected and the technology is accurate, reliable, effective, and transparent."737 I searched for a published figure — any published figure — for how many Americans fail or abandon federal identity proofing. There is none in the public record that I could locate. That absence is the finding. The unemployment insurance side shows the same shape. Auditors found that "State efforts to address fraud can also lead to delays for legitimate claimants; officials in three states told us that they suspend affected accounts during an investigation, and claimants lack access until their accounts are restored."738 They recorded a claimant in Massachusetts saying: "My challenge is, because of the whole fraud thing, I had to verify every time. Even now when I log in, I have to do an extra step just to verify it's me by phone text— every time I log in, I have to verify who I am."739 The federal alternative is not obviously better. Its own oversight office "confirmed that Login.gov did not meet the IAL2 requirements and that GSA misled their customer agencies when GSA failed to communicate Login.gov's known noncompliance."740

IX.

I want to pause on something that complicates the easy version of this argument, because it complicates mine. The rules already exist. They are good rules. They are specific, they are binding, and they were written by the federal agencies themselves. On identity verification in unemployment insurance, the labor department issued a program letter in April 2021 that reads as if someone had studied the Michigan disaster and written the antidote. On document alternatives: "A state's list of acceptable documents to verify ID must include a sufficient variety to provide for equal access to UI benefits for all claimants. For example, if the state requires a state driver's license as proof for ID verification, it must offer a reasonable alternative for those who may not have that type of identification or provide sufficient time for other means of identification to be obtained."741

On non-digital access: "States must provide alternative mechanisms for individuals with access barriers, such as a disability or limited English proficiency."742 On reasoned decisions: "The determination must explain the reason for the decision, including the facts: what the individual was asked to provide; what was received; and, if information was insufficient, why."743 On portal notices: "if the state posts such notice(s) on a claimant portal, the message must remain in the claimant's login portal until the issue has been resolved."744 The Medicaid regulations are equally explicit about electronic notice. A beneficiary must be given "a choice to receive notices and information … in electronic format or by regular mail" and "must be permitted to change such election." If they choose electronic, the agency must "[e]nsure that the individual's election to receive notices electronically is confirmed by regular mail," must "[p]ost notices to the individual's electronic account within 1 business day," and must "[s]end a notice by regular mail within three business days of the date of a failed electronic communication if an electronic communication is undeliverable."745 The food assistance regulations guarantee

a right to file that no portal outage can extinguish: "All households have the right to apply or to re-apply for SNAP in writing. The State agency shall neither deny nor interfere with a household's right to apply or to re-apply in writing." And the minimum viable application is defined as narrowly as it can be: the state "must provide a means for all applicants applying through any mechanism to immediately begin the application process by filing an application with only the name, address and signature."746 So the paper regime is strong. Which raises the only interesting question: why does none of it show up in the experience of the people I have described? Part of the answer is that some of these duties are written to be elastic. The language-access obligation under civil rights law is governed by a four-factor balancing test, and the translation safe harbor is explicitly one-way: "The failure to provide written translations under the circumstances outlined in paragraphs (a) and (b) does not mean there is non-compliance."747

Part of the answer is that a right you must know about, and invoke, and insist upon, against a call center with a three-hour hold time, is not the same kind of thing as a right the system implements by default. The food agency's own civil rights guidance names the problem when it advises "[a]dequately funding and staffing call centers to ensure they are accessible without prolonged delays," noting that people with limited English proficiency and people with disabilities "may have limited access to internet access, or limited transportation to apply for services in person, and therefore may rely more on call centers."748 But the largest part of the answer is this: the rules bind the agency's formal conduct, and the software determines the citizen's actual experience. Nobody violated the right to file in writing. They built systems in which the digital path is fast, staffed, and measured, and the paper path is slow, thin, and unmeasured — and then reported the resulting digital adoption as a preference.

X.

Now the clawback, which is where automation becomes something closer to a collections operation. The retirement and disability agency's own report to Congress sets out the scale. In fiscal year 2024 it "recovered approximately $4.9 billion in OASDI and SSI program overpayments," had "approximately $10.3 billion scheduled for repayment," and "waived approximately $302 million."749 Where do overpayments come from? The inspector general's office reports that from fiscal 2015 through 2022 the agency "made nearly $72 billion in improper payments, most of which were overpayments," with "an uncollected overpayment balance of $23 billion" at the end of fiscal 2023.750 Among the leading causes on the retirement and disability side are the agency's own computations. The agency makes an arithmetic error. The beneficiary repays it, with interest of a kind: not financial interest, but the compound interest of dread. Until March 2024 the default recovery rate was one hundred percent of the monthly benefit. The Commissioner who changed it said so bluntly: "It's unconscionable that someone would find themselves facing homelessness or unable to pay bills, because Social Security withheld their entire payment for recovery of an overpayment."751 That is a real reform and I credit it plainly. But look at the mechanics in the implementing instruction: "If the overpaid individual does not repay, request a waiver, or request a reconsideration prior to the end of the 60-day due process period, we will, in most cases, automatically recover the overpayment by withholding 10 percent of their Title II monthly benefit credited (MBC) amount, or $10 per month, whichever is more."752

Silence for sixty days initiates automatic collection against a person's subsistence income, with no human required to touch the file. The vehicle for stopping it — a waiver — requires first clearing a fault threshold: the agency grants a waiver when "[t]he individual was without fault in causing the overpayment" and recovery

would "defeat the purpose" of the program, which it defines in part as the individual needing "substantially all of their income … to meet ordinary and necessary living expenses."753 So the person must prove they did not cause an error the agency may itself have made, and then separately prove they are poor enough to be excused. At a forgiveness rate of six cents on the collected dollar.

XI.

Everything so far concerns whether you get money. This section concerns whether you keep your children. The Allegheny County family screening tool is the best-documented predictive system in American child welfare, and I credit the county for publishing its methodology, evaluation, and limitations. What follows is drawn from the county's own documents. The tool "is a predictive risk model that supports child welfare call screeners and supervisors in deciding how to respond to reports of suspected child abuse or neglect," in use since August 2016. It "analyzes integrated data from the Allegheny County Data Warehouse and generates a risk score that estimates the likelihood of a child experiencing harm within two years."754 The tool uses over 100 predictive factors per child, weighted through a regression model to produce scores "ranging from 1–20."755 "The score ranges from 1 to 20 (where 20 is the highest 'risk' and 1 is the lowest), indicating the ventile into which the AFST falls. A recommendation for 'auto screen-in' occurs when the AFST falls above 18 for the placement score."756 Every adult and every child in the home is scored on the strength of a single hotline call. The county lists the data sources: "Child welfare records · Jail records · Juvenile probation records · Behavioral health records," with birth records added in the second version — and, critically, the first version also used "Public benefit records (e.g., Temporary Aid to Needy Families [TANF], Supplemental Nutrition Assistance Program [SNAP])."757 In the tool's first version, the fact of having received food assistance was an input into a score estimating whether your child would be removed from your home. Not neglect. Not injury. Participation in an anti-poverty program. The county removed those records in the second version.

The county is also candid about a related defect. It stopped modeling re-referral because "initial incoming referral rates also represent the most racially disproportionate step of the referral pathway, and so a model predicting future referrals tends to overrepresent black children relative to white."758 That is a real methodological improvement. It is also a confession about what a predictive model trained on administrative contact actually predicts: not harm, but the system's own past attention.

XII.

What did the tool accomplish? The county commissioned an outside evaluation comparing about 31,000 children referred in the eighteen months before deployment with about 34,000 after.759 The findings, in the evaluation's own words: "Overall, the AFST did not lead to increases in the rate of referred children screened-in for investigation." "Use of the AFST did not lead to decreases in re-referral rates for children screened-out without investigation. Re-referral rates among children screened-out stayed the same for children overall, with the exception of children who were 4-6 years of age. … Unfortunately, for this age group there was a slight but statistically significant increase in the likelihood of the Post-AFST group of children being re-referred." "The AFST led to reductions in disparities of case opening rates between black and white children." "There was no evidence that the AFST resulted in greater screening consistency within individual call-screeners."760 That record does not make Allegheny County a villain. It makes Allegheny County the only honest data point we have. Which is why what happened next should concern us.

XIII.

In October 2022 the federal health department's own policy office published an analysis of racial bias in child welfare predictive risk models. Its warning was explicit: such models "have the potential to exacerbate existing racial and ethnic disparities if not developed with careful attention to the methods used."761 Its recommendations were specific: agencies "should assess the risk and potential benefits of using the model, consider how best to balance PRM misclassification rates," should "continuously examine its predictive performance and evaluate its impact on equity and other outcomes," and should "engage community members and caseworkers in making key decisions and be transparent about their approach."762

"Overall, about 37% of all children in the United States experience a CPS investigation by the age of 18," with rates "of approximately 53% and 32%" for African American and Hispanic children respectively, "compared to 28% of White children."763 More than half of Black children in America will be the subject of a child protective services investigation before they turn eighteen. That is the base rate into which a scoring system is being introduced. On August 27, 2026, the federal children and families administration announced that it "awarded $6 million to 10 child welfare jurisdictions to test predictive risk modeling (PRM) and related predictive analytics approaches across the child welfare continuum—from prevention and hotline screening to placement stability and permanency."764 The recipients: "Indiana, Missouri, Nebraska, New Jersey, North Carolina, Ohio, Oklahoma, Texas, Muscogee Creek Nation, and the District of Columbia."765 The technologies: "machine learning, natural language processing, automated case summaries, dashboards, cross-agency data integration, text analytics, and predictive models embedded within Comprehensive Child Welfare Information Systems."766 I read the announcement looking for conditions: a bias-testing requirement, a disability-discrimination safeguard, a pre-deployment audit, a due-process provision. There are none.767 An announcement bulletin is not a grant agreement, and the underlying terms may contain safeguards the announcement does not mention. I could not locate the underlying issue brief on the agency's own site to check.768 But the public-facing document is what the public gets, and it contains no safeguard at all.

XIV.

The law that already governs this is disability law, and it is stronger than most people realize. In 2015 the civil rights division of the justice department, together with the health department's civil rights office and its children and families administration, issued joint technical assistance for child welfare agencies. The operative standard: "Individuals with disabilities must be treated on a case-by-case basis consistent with facts and objective evidence. Persons with disabilities may not be treated on the basis of generalizations or stereotypes."769 Agencies "should take steps to ensure, for example, that investigators, social workers, supervisors, and others base their assessments of and decisions regarding individuals with disabilities on actual facts that pertain to the individual person, and not on assumptions, generalizations, fears, or stereotypes about disabilities and how they might manifest."770 DAVID M. DOYLE 234 A predictive risk score is, definitionally, a generalization. It says: people whose administrative records resemble yours have historically had this outcome. And when behavioral health service records are among its inputs — as the county documents confirm they are, in the form of "aggregated indicators for whether each individual on the referral received any prior behavioral health service, as well as the number of days since the last behavioral health service"771 — the score is partly a function of having sought treatment. I want to be careful not to overstate the legal posture. The 2015 guidance does not mention algorithms; the terms "predictive analytics," "predictive risk modeling," and "algorithm" do not appear

in it.772 It was written about caseworkers, not models. The argument that it reaches automated scoring is mine, not the government's. But the government has come close, twice, recently. First, the four-agency joint statement on automated systems warned that "Many automated systems are 'black boxes' whose internal workings are not clear to most people and, in some cases, even the developer of the tool."773 Second, a voluntary resolution agreement signed in September 2025 between federal civil rights authorities and Washington State's child welfare department. Investigators "concluded that DCYF's policies, practices, and procedures were systemically deficient in preventing discrimination against parents with disabilities in Washington's child welfare system."774 The legal characterization maps onto a screening score with uncomfortable precision: "use of methods of administration that have the effect of discriminating on the basis of disability, and reliance on eligibility criteria that screen out or tend to screen out individuals with disabilities."775 That agreement does not mention algorithms either; its full text contains no reference to automated systems or scoring models.776 But "criteria or methods of administration that have the effect of discriminating" is exactly what an unaudited risk model built on behavioral health records is, and the first agency to be told so will have no excuse, because the standard has been on the books since 1973. Candor requires recording what I could not establish. Reporting has asserted that federal civil rights authorities made inquiries about the Allegheny tool specifically. I searched the justice department's own site and found no document naming the county or the tool, and I am not going to assert a federal investigation on the strength of secondary accounts.777

XV.

Housing is where the scoring economy is most fully built out, and where the enforcement record is longest.

The clearest description of a false-match algorithm in the American record comes from the trade commission's 2018 action against a tenant screening firm. The matching criteria required only "an exact match of the applicant's last name" plus "a non-exact match of a first name, middle name, or date of birth." So for an applicant named Anthony Jones born 10/15/1967, the records of "Antony Jones, 10/15/67," "Antonio Jones, 10/15/67," and "Antoinette Jones, 10/15/67" would all be treated as matches.778 Antoinette Jones. Under that rule, a woman's criminal record could attach to a man's rental application. The company "used broad criteria to match applicants to criminal records," "only applied limited filters to the results," and "did not have policies or procedures in place to assess the accuracy of those results."779 The consumer bureau later stated the rule directly: "the practice of matching consumer records solely through the matching of names is illegal under the Fair Credit Reporting Act," because name-only matching "does not assure maximum possible accuracy of consumer information."780 It warned that the opinion "does not create a safe harbor to use insufficient matching procedures involving multiple identifiers," since "name combined with a date of birth, could also lead to cases of mistaken identity."781 There is a documented racial gradient, and the mechanism is arithmetic rather than animus: "The risk of mismatching from name-only matching is likely to be greater for Hispanic, Asian, and Black individuals because there is less last-name diversity in those populations than among the non-Hispanic white population." Specifically, "[t]wenty-six surnames cover a quarter of the Hispanic population."782 Enforcement continued. In October 2023 a credit reporting agency and its subsidiary paid $15 million over tenant screening reports containing "inaccurate and incomplete eviction records," including failures to "prevent the inclusion of multiple entries for the same eviction case" and to "prevent the inclusion of sealed eviction records."783 In July 2026 another screening firm paid $2.25 million for, among other things, not "preventing the inclusion of duplicate case records and multiple entries for the same criminal or eviction action," which

"gave the false impression that applicants had more criminal convictions or had been sued for eviction more times than they actually had."784

On the underlying data, the market report is devastating. "A study of 3.6 million eviction court records from 12 states found that 'on average, 22% of state eviction cases are ambiguous or false records.'" Some "34% of eviction filings over the past decade ended in default judgments where the tenant did not attend the court hearing, often because they were not aware of the hearing." And "fewer than 3% of tenants have access to counsel, compared to 81% of property owners."785 Assemble the pipeline. A tenant is sued. They are not effectively served, or cannot take a day off work. A default judgment enters. The record is ambiguous or false in roughly a fifth of cases. It is sold to a screening company, matched by surname to possibly the wrong person, compressed into "a score, a recommendation, or any other sort of explicit or implicit opinion," and returned to a landlord who sees a number.786 The scores "are not easily understood due to low transparency of the underlying models and data inputs used by each company, nor is there independently validated, publicly available evidence that they reliably predict rental outcomes."787 There is no evidence they work. There is extensive evidence that their inputs are wrong. They are used anyway, on a market of "over 100 million people across 45 million rental units."788

XVI.

The housing regulator has written the clearest algorithmic-accountability language in American administrative practice, and almost nobody has read it. From its April 2024 guidance on tenant screening: "A housing provider or a tenant screening company can violate the Fair Housing Act by using a protected characteristic – or a proxy for a protected characteristic – as a screening criterion. This is true even if the decision for how to screen applicants is made in whole or in part by an automated system, including a system using machine learning or another form of AI."789 On the futility of naive de-biasing: "Even if specific protected characteristics (and their more obvious proxies) are removed from a model's inputs, machine learning and other forms of AI — without corrective adjustments — may compensate for their removal by altering the weights of the remaining factors, thereby recreating the same discriminatory outcomes."790 A CITIZEN'S MANIFESTO FOR THE AGE OF AI 237 On necessity: "Only screening criteria and standards – including complex models – that are well-tailored to predict future behavior relevant to tenancy can be considered necessary to achieving a legitimate interest."791 On interpretability: "Even complex models should be interpretable enough that adjustments can be made to fix any fair housing problems that are found."792 And: "If a highly complex model has a discriminatory effect, the model's lack of transparency could make it hard to prove that a legally sufficient justification exists."793 Opacity is not a defense. It is an aggravating factor. The guidance also tells screening companies to stop hiding the reasons: "Denial recommendations or low 'grades' should not be provided to housing providers in a conclusory fashion. … Rather, screening reports should include all available relevant details about the basis for the determination (e.g., dates, locations, case numbers, and dispositions). … These details should be in plain language, so a lay person can easily understand the reasons for the denial."794 The justice department has taken the same position in court, filing a statement of interest in litigation over a tenant scoring product used against voucher holders: "Tenant screening policies are not exempt from the Fair Housing Act's protections just because decisions are made by algorithm."795 A model that ignores the single most predictive fact about whether the rent will arrive, and weights instead the debts that poverty generates, is not a prediction engine. It is a filter with a prediction engine's vocabulary.

XVII.

Then there is the direct surcharge on being poor. Start with applying for housing. "Application fees paid by prospective tenants typically range between $40 and $59," while screening companies charge landlords "approximately $25-$35 for each report."796 The fee exceeds the cost. And the burden is uneven: "The typical white renter reported paying $35 in application fees on their rental, while the typical Black, Latinx, and Asian renters all reported spending $50 on application fees." Worse, they apply more often: "Black and Latinx renters were also almost twice as likely to report submitting 5 applications or more (38 percent of Black and Latinx renters report submitting 5 or more, compared to 21 percent of white renters)."797 The people most likely to be falsely matched pay the most to be falsely matched, repeatedly.

Then consider release from incarceration, where the fee is charged on money that is already yours. The consumer bureau found that a release-card company "charged consumers unavoidable fees for prepaid cards used to return money owed to consumers at the time of their release from incarceration," having "leveraged its relationships with state and local departments of correction to impose fees on consumers exiting the prison or jail system."798 The card "replaced cash or check options previously offered by state departments of correction," and the company "did not provide a reasonable way for consumers to close their card accounts to obtain their card balances without paying fees."799 The bureau's phrase was exact: "By assessing fees on these captive consumers, JPay took advantage of them and caused harm."800 "Between May 2011 and June 2017 alone, over 500,000 consumers were forced to receive the funds owed to them at the time of their release on a feebearing Debit Release Card. Those cardholders incurred over $3 million in fees during that period."801 And the sentence that explains everything: the company "believed that entering into contracts with prisons for release cards would help it compete for other contracts for additional services; contract take-up was not driven by consumer demand."802 Even cash assistance carries a toll. Federal law requires states to ensure recipients "have access to using or withdrawing assistance with minimal fees or charges, including an opportunity to access assistance with no fees or charges."803 But the federal guidance concedes: "Even if the state TANF program provides fee-free transactions, recipients may still be subject to surcharges imposed by the ATM network owner."804 Eighty cents is not a scandal. Eighty cents charged to someone receiving cash assistance, to obtain the cash assistance, by a vendor under contract to the state that is providing the assistance, is a description of how the poverty premium actually works: never large, never singular, always there.

XVIII.

Finally, the utilities, where automation has quietly changed the physics of disconnection. For the first time, the federal energy agency has counted. In 2024 there were 13,447,935 residential electricity disconnections in the United States, against 11,370,551 reconnections, preceded by 94,870,709 final notices.805 Texas alone accounted for 3,028,150 disconnections.806

The federal form does not ask how many were executed remotely. But we can price the incentive. In a 2020 proceeding, a state utilities commission recorded a utility's weighted costs of "remote and manual reconnection … which are $4.08 and $39.75 respectively," and observed that "with smart meters the utilities can reconnect electric customers remotely, and this significantly reduces the costs."807 When disconnection required dispatching a human being to a house, cost imposed a natural limit. A truck roll is expensive, and the person in the truck sees the house — sees the child's bicycle, the oxygen concentrator, the window air conditioner in August. Remote disconnection removes the cost and the witness at once. One state has written the most explicit

rule I found. Utilities must "[p]erform all remote disconnections for nonpayment between the hours of 8:00 a.m. and 12:00 p.m." and only where the customer has "a reasonable opportunity to submit payment and have the utility reestablish service on the same day."808 They must "[d]isable remote disconnection functionality for medical facilities and critical infrastructure" and "limit the number of remote disconnections in a twenty-four hour period or take other reasonable measures to prevent unauthorized disconnections."809 And twice, the rule requires a human to go there. Before disconnecting a customer with a medical certificate, the utility must "visit the customer's premises and provide the customer with an opportunity to pay."810 And: "[p]rior to disconnecting a customer for nonpayment who the utility is aware has received low-income assistance in the prior two years," the utility must likewise "visit the customer's premises and provide the customer with an opportunity to pay via appropriate methods including providing payment to the dispatched utility representative."811 That is a regulator deciding that poverty is a reason to send a person rather than a signal — a refusal to let a decision get cheap enough to make carelessly. Most states have no equivalent. Disconnection procedure is overwhelmingly a matter of state and local regulation, and "policies vary among the nearly 3,000 electric utilities in the country."812

XIX.

Let me set out the strongest objections, because they are not weak.

The first objection: fraud is real, and somebody has to stop it. True. Improper payments in the tens of billions are a genuine stewardship failure. But a control tuned against a metric that counts only one kind of error will keep tightening until it hits a constraint, and right now the only constraint is litigation, which takes nine years and ends in contempt findings. Publish the false-denial rate next to the improper-payment rate and the control has something to push back against.

The second objection: procedural terminations are not wrongful terminations. Also true. Some of those 14.3 million people had moved, or had employer coverage, or were no longer eligible. But thirty jurisdictions were running a renewal process that the federal government found was improperly disenrolling eligible people, and the defect was invisible from the outside. The honest position is not that all 14.3 million were wronged. It is that we built a system in which no one can tell which ones were. The third objection: the screening tool is only a recommendation; a human decides. The county says exactly this, in good faith: "It was never intended or suggested that the algorithm would replace human decision-making."813 I accept the intent. But a score above 18 triggers auto screen-in,814 and the county's own evaluation found no evidence that the tool produced greater screening consistency within individual call-screeners.815 A recommendation that automatically becomes a decision at the top of the range, and does not measurably change decisions elsewhere, is not decision support. The fourth objection: these tools reduce disparity. Sometimes they do: the evaluation found "reductions in disparities of case opening rates between black and white children."816 A well-built, published, independently evaluated model, with its most poverty-loaded inputs removed, produced a real equity gain in one jurisdiction. That is the strongest argument for the position I am about to take — which is not prohibition, but audited deployment. The fifth objection: you are asking for costly process. Yes. Due process is costly; that is what the Supreme Court weighed in 1970 when it held that the recipient's interest and the state's interest in not erroneously terminating people together outweigh "the State's competing concern to prevent any increase in its fiscal and administrative burdens."817 The cost was considered and the answer was given.

What I will concede is the limit of my own evidence. I could not verify the widely circulated ninety-three percent figure from Michigan. I could not find any published figure for how many people fail federal identity proofing, because the metric is built to exclude them. I could not find the terms and conditions attached to the new child welfare modeling grants. I could not determine how many of the thirteen million shutoffs were executed remotely, because the form does not ask. In a chapter about measurement, I have had to note four places where the measurement does not exist. That is the same finding, appearing again.

XX.

Before the requirements, let me set the concerns out plainly, because they are not one problem but seven, and they compound.

Every row is the same shape. A decision that used to require a person now requires a threshold, and the threshold was set by someone optimizing a number that does not count you. Here is what I would insist on. Publish the false-denial rate, beside the improper-payment rate. Every program that reports a payment error rate must report a denial error rate, disaggregated by race, disability, primary language, and geography. The food program already computes one; it reported 43.81 percent nationally and above eighty percent in one state.818 Attach the same fiscal consequence to one that attaches to the other. Presumptive continuation of benefits during appeal. Not restoration after you win. Continuation while you contest. No algorithmic output may reduce, suspend, or terminate a subsistence benefit until the beneficiary has had timely notice and a real opportunity to be heard. The Arkansas notice standard, everywhere. Any adverse determination produced with the help of an automated system must be accompanied by the assessment itself, the specific classification assigned, plain-language reasons, and an explanation referencing the system "including the algorithm" — plus advance notice of any limits on what evidence may be presented on appeal.819 No black boxes in adverse determinations. If a scoring instrument is used against a person, that person gets the manual. The Idaho court's reasoning is sufficient: without it, "it would be impossible for the participant to prove error."820 Independent audit before deployment, for any predictive system in child welfare, housing, or justice. Pre-deployment, published, by someone who did not build it, with error rates reported separately for false positives and false negatives, by race and disability. No federal dollar for a predictive system without it.821 A named human decision-maker for every adverse determination. A name, a title, and a way to reach them. If no human being is willing to attach their name to cutting off a family's food, that is diagnostic. Automatic enrollment instead of burden-based means-testing. Where the government already possesses the data that establishes eligibility, it should enroll the person. The Medicaid rule already says as much for renewals — the agency "must make a redetermination of eligibility without requiring information from the individual if able to do so based on reliable information contained in the individual's account."822 Extend that principle from renewal to enrollment. Statutory non-digital access, with equal service levels. Not merely the right to file on paper, but a requirement that the non-digital path be resourced to comparable speed. The labor department already requires "alternative mechanisms for individuals with access barriers."823 A right to a slow channel is a rationed right.

Publish identity-proofing outcomes, including abandonment. No federal identity system should report a pass rate that "excludes both users who abandon the process."824 Count the people who gave up. Report the figure by demographic group.825 A right to the data used against you, in plain language, before the decision is final. Not after. Not on request.826 Before. Send a person before you cut the power to a poor household. One state already requires a site visit before disconnecting a customer known to have received low-income assistance

within two years.827 Remote disconnection dropped the marginal cost of cutting someone off from $39.75 to $4.08.828 When you make a grave act cheap, you should expect it to become common. An error rate is never a single number. It is a pair of numbers in tension, and every threshold ever set is a decision about how many wrongly denied people we will accept to avoid one wrongly paid one. What is remarkable is how completely this has been hidden from public debate, where "accuracy" is discussed as though it were one thing and improvements in it were free.

A system that produces one wrongful denial to prevent one wrongful payment has not been made accurate. It has chosen whose error it prefers, and it has chosen the poor. We have automated the machinery of exclusion faster than the machinery of provision. We can match your surname against an eviction record in milliseconds, score every adult in your household on more than a hundred variables, cross-reference your bank balances before ruling on your hardship waiver, and cut your electricity from a desk. We have not built, with anything like the same energy, the machinery that finds the eligible person who never applied, or restores the child wrongly dropped from coverage. Make it someone's job. Put their name on the letter. — D.M.D.

Chapter Twelve

Chapter Twelve: Provision and Movement

What reaches the town Who decides what reaches my town, and at what price?

I.

There is a sentence in the United States Code that I think every town council in the country should have framed on the wall. It was enacted in 1970, in the law that created the Postal Service. "No small post office shall be closed solely for operating at a deficit, it being the specific intent of the Congress that effective postal services be insured to residents of both urban and rural communities."829 Read what that sentence does. It does not say small post offices must never close. It does not deny that some lose money. It says something narrower and more radical: that losing money is not, by itself, a sufficient reason to cut a place off. Fifty-six years ago, Congress looked at the question this chapter asks — who decides what reaches my town? — and answered it in statute. Not the balance sheet alone. A place does not stop existing, for purposes of national service, because it is small and far away and expensive to reach. On September 8, 2026, the Postal Service's own Inspector General reported on what the current network plan will do when fully implemented. One initiative, the regional transportation optimization program, "will affect 72 percent of ZIP Codes in the contiguous U.S., covering 48 percent of the population and 29 percent of outgoing package volume. Rural areas will be impacted the most, with 86 percent of rural populations affected compared to 45 percent of urban populations."830 Eighty-six percent of rural people against forty-five percent of urban ones. Nobody broke a law that a court has yet enforced. No one set out to harm rural America. It is the predictable output of an optimization that counts volume, applied to a country where volume is unevenly spread. This chapter is about that kind of decision — what reaches a place and at what price — across the physical system that keeps people alive: food, water, medicine, fuel, freight, transit, and the emergency supply that is supposed to arrive when everything else has failed.

I will argue that a quiet change has happened across all of these systems at once. Each has been tuned, over forty years, to remove what an efficiency analyst would call slack and what an ordinary person would call a spare: the second supplier, the backup plant, the extra day of inventory, the local mechanic, the rural route that loses money. Each tuning was locally rational. Together they have produced a physical economy that is cheaper in normal times, more fragile in abnormal ones, and increasingly governed by models whose logic the people on the thin end of the route cannot see.

II.

Start with the Postal Service, because it is the cleanest case. It is public. It has a statutory mandate to serve everyone. It reports to a regulator. And even so, the logic of aggregate optimization has pulled it toward the same place a private carrier would go. The plan is called Delivering for America. Its stated aims are fair ones: the Service says it can "avoid $160 billion in projected losses by 2030," that it will "invest $40 billion in our people and infrastructure over 10 years," and that the plan "preserves our ability to fulfill our universal service mission."831 An insolvent postal system serves nobody. The Service asked the Postal Regulatory Commission for an advisory opinion on the service-standard changes the plan required, in a proceeding docketed as N2024-1. The Commission issued its opinion on January 31, 2025, and it is one of the most useful public documents on

routing that I have read, because it does something the Service had not done: it measured the effect place by place instead of letter by letter. The Service had analyzed impact by volume. The Commission found that approach misleading: "The Postal Service's volume-based analysis obscures and minimizes negative impacts on rural areas by giving greater weight to higher volume areas."832 When the Commission instead counted pairs of origin and destination ZIP Codes: "For Single-Piece FCM, 64.0 percent of rural 5-Digit to 3-Digit ZIP Code pairs will be downgraded, compared to 38.3 percent of non-rural pairs." Upgrades ran the other way: "only 1.7 percent of rural pairs will experience upgrades compared to 6.7 percent of non-rural pairs."833 Here is the mechanism in one sentence. An optimization weighted by volume discounts the places where volume is low, automatically, because that is what weighting by volume means.

The Commission weighed what the country would get in exchange. "Total projected cost savings, even if fully realized, represent approximately 4.4 percent of the Postal Service's FY 2024 operating expenses of $81.8 billion."834 It urged the Service "to reconsider whether the speculative, meager gains from this proposal outweigh the certain downgrade in service for a significant portion of the nation," and found the Service "irreversibly changing its network without laying a foundation for success."835 On the legal question, the Commission was careful. The statute forbids the Service to "make any undue or unreasonable discrimination among users of the mails."836 The Commission said that "although the Commission is very concerned about the disproportionate impact the proposed changes will have on rural communities, there is insufficient information on the record to render a definite determination of whether the proposal violates 39 U.S.C. § 403(c)."837 Concern, not a finding. The opinion is also advisory; the Service was not bound by it. The Inspector General's audits show what happened when the ideas were tried. A 2024 audit of the local transportation pilot reported that it "negatively impacted service to customers. The service performance scores for First-Class Mail declined after the implementation of LTO, more significantly impacting the rural population."838 And: "overall transportation expenditures increased by $7.13 million for the 15 LTO implemented regions when compared to the same period last year."839 Service got worse for rural customers, and the transportation line it was designed to cut went up. The 2026 assessment covers the whole network over the whole plan, and its verdict is sober: transportation expenses reduced "by about $755 million," workhour costs "by about $2 billion," and yet "the Postal Service's overall financial position has not materially improved compared to when these initiatives began."840 The network changes "coincided with service performance declines, with rural communities experiencing the most noticeable impacts."841 One more detail belongs here, because it shows how the law itself has drifted. In 2022, Congress amended the section governing how the Service chooses modes of transportation. The old text required "prompt and economical delivery of all mail." The new text requires "prompt, economical, consistent, and reliable delivery of all mail in a manner that increases operational efficiency and reduces complexity."842 Congress wrote operational efficiency into the transportation mandate, in the same statute that still says no small post office may be closed solely for losing money. Both sentences are now law. The question of which one governs a particular rural route is not being settled in public. It is being settled in network models.

III.

The postal story generalizes, and the federal government has said so in its own words.

In June 2021, the White House published its hundred-day review of American supply chains. Its central admission deserves to be quoted in full: "Our private sector and public policy approach to domestic production,

which for years, prioritized efficiency and low costs over security, sustainability and resilience, has resulted in the supply chain risks identified in this report."843 The Council of Economic Advisers, in its 2022 report to Congress, made the argument more precisely. "Low inventories by themselves do not necessarily lead to fragility; problems arise when low inventories are combined with low agility."844 Firms can buy resilience through "investing in backup capacity (redundancy)," the Council noted, "however, these strategies, especially redundancy, increase costs," and so "it is not cost-effective for firms to invest in completely avoiding all disasters."845 Then the Council named the market failure. "One consequence of a firm underinvesting in resiliency is that it increases the exposure of other firms in the networked system to the negative spillover effects of a disruption. This type of market failure is likely when the firm's investment decisions consider only its private costs and benefits, and it is unable to monetize the spillover benefits of its investment decisions for the rest of the system."846 That is the economic core of this chapter. Redundancy is a public good that private balance sheets are built to underprovide. A spare plant, a second supplier, an extra week of stock — each costs the firm that carries it, and each benefits everyone downstream who would otherwise have gone without. No firm is paid for the shortage that did not happen. By December 2024, the government's quadrennial review had reached the same conclusion. "Just-in-time" production, it said, "can come at the expense of investing in multiple suppliers, geographic supplier diversity, workers, and buffer inventory stocks that would improve the ability of supply chains to withstand external shocks." And: "Recent experience has starkly demonstrated that modern supply chains that are highly concentrated may be insufficiently resilient."847 Three administrations' worth of analysts have now agreed on the diagnosis. What has not happened is a corresponding set of obligations. Redundancy is still, for nearly every essential good, a voluntary choice made by a private firm against its own quarterly interest.

IV.

Now to food, beginning where most people meet it: the grocery store.

In March 2024, the Federal Trade Commission published a staff study of what happened inside the grocery supply chain during the pandemic. It used the Commission's compulsory study power to obtain internal materials from nine companies — "three grocery retailers, three grocery wholesalers, and three grocery producers."848 What it found was not primarily about price. It was about allocation: who got the product when there was not enough to go around. The mechanism was contractual. "To gain access to scarce products and therefore a competitive advantage, some companies—most often larger ones—used policies that imposed strict delivery requirements on their upstream suppliers and threatened fines for noncompliance. Walmart even tightened the delivery requirements its suppliers had to meet to avoid fines as the pandemic went on. In some cases, suppliers preferentially allocated product to the purchasers threatening to fine them."849 Specifically: "Walmart tightened its OTIF requirements in September 2020, requiring suppliers to achieve 98 percent OTIF compliance to avoid fines of 3 percent of the cost of goods."850 The consequence fell on smaller stores, which "did not have such contractual commitments or protections received less than their purchasing history would have allotted them, meaning they bore more than their share of the impact from shortages."851852 Think about who shops at the smaller store. In many rural counties and many poor neighborhoods it is the only store. So a contract term written at a distant headquarters, enforced through a supplier's allocation decisions, decided in 2020 which towns had the scarce goods on the shelf. Nobody in those towns was party to the contract. The Commission drew the structural lesson. "The potential for powerful retailers to distort product allocations during a shortage suggests that crises may create an opportunity for some firms to entrench their market

power."853 And: "Without policy intervention, markets that are fragile because of concentration may only further consolidate."854 The concentration is real. "The retail grocery sector has undergone significant consolidation over time, with the top four firms accounting for over 30 percent of the sales in 2019, as compared to around 15 percent in 1990."855 Nationally, that is doubled but not dominant. Locally it can be a different story, and the next largest consolidation was stopped in court.

In December 2024, a federal judge in Oregon enjoined the proposed merger of Kroger and Albertsons. The plaintiffs' expert found that "1,922 supermarket markets would have presumptively anticompetitive market concentration under the 2023 Merger Guidelines."856 The court found "that the merger would lead to undue market concentration in multiple geographic markets"857 and rejected the proposed divestiture of 579 stores.858 The companies called off the merger the next day. The same report also looked at margins. "Food and beverage retailer revenues increased to more than 6 percent over total costs in 2021, higher than their most recent peak, in 2015, of 5.6 percent. In the first three-quarters of 2023, retailer profits rose even more, with revenue reaching 7 percent over total costs."859 These are staff findings, not a Commission adjudication, and retailers dispute the interpretation. I report them as the government's own measurement.

V.

What does all this mean at the level of a household? The Agriculture Department's research service maintains the federal map of what is usually called food deserts, and is careful to call them something more precise: low-income, low-access census tracts. By its estimate, "in 2019, 11 to 27 percent of the U.S. population lived in low-income (LI) and low-access (LA) census tracts," the range depending on how distance and vehicle access are defined.860 The rural trend runs in one direction. The department's researchers found "that in rural counties, grocery stores have declined, while dollar stores and supercenters have increased steadily." In 2015, "there were 23 counties without food retailers of any type, and all of those were in rural nonmetro counties. Of the 44 counties with no grocery stores, 40 were rural nonmetro."861 I call food deserts a routing outcome, and I want to justify the phrase. A grocery store is the end of a distribution route. It exists where a chain's site-selection model and its distribution network say a store can meet a margin target, and it closes where they say it cannot. No one decided that a particular county should have no grocery store. A great many separate optimizations decided it together. The burden of price falls where you would expect. Between January 2020 and January 2023, "U.S. food-at-home prices increased 24.0 percent."862 And in 2023, "food costs ranged from 32.6 percent of after-tax income for households in the lowest income quintile… to 8.1 percent for the highest income quintile."863 The poorest fifth of households spends roughly a third of what it has on food. A ten percent rise in the price of groceries is, for them, a three percent cut in everything else.

Meanwhile the farmer's share of the food dollar is small. Under the department's revised methodology, "the farm share was 11.8 cents per dollar spent by U.S. consumers on domestically produced food in 2024."864 I want to flag a trap I nearly fell into myself: the department revised its method in March 2026, and the older series put the 2023 figure at 15.9 cents while the new one puts it at 12.1.865 The two series cannot be compared. Under either method, most of what a family pays for food goes to the stages between the farm and the table — processing, distribution, retail — and that is where the concentration and the routing decisions sit.

VI.

The narrowest point in the American food system is not the grocery store. It is the slaughterhouse. In January 2022, the White House stated the concentration plainly: "Four large meat-packing companies control 85 percent of the beef market. In poultry, the top four processing firms control 54 percent of the market. And in pork, the top four processing firms control about 70 percent of the market."866 "In 1977, the largest four beef-packing firms controlled just 25% of the market, compared to 82% today," the White House wrote in 2021.867 Here I must amend what I said in Chapter Ten. At the level of the whole food supply, four firms do not hold it. At the level of beef processing, four firms very nearly do. A concentration need not cover a whole sector to function as a gate; it only needs to sit at the point everything passes through. The action plan described the fragility in terms any rancher would recognize: "When COVID-19 or other disasters such as fires or cyberattacks shutter a plant, many ranchers have no other place to take their animals. Our overreliance on just a handful of giant processors leaves us all vulnerable, with any disruptions at these bottlenecks rippling throughout our food system." It added: "50 beef slaughter plants owned by just a handful of companies currently process nearly all the cattle in the United States."868 The response was "$1 billion in American Rescue Plan funds for expansion of independent processing capacity."869 The government concluded that the market would not build spare capacity on its own and paid to build some.

The counterweight belongs in the record. The Agriculture Department's research service summarized the economic literature this way: packers with fewer rivals "were also able to exercise some market power and pay lower prices for cattle than they would have had they faced more competitors. Even so, the impact of higher concentration—reducing prices paid for cattle—was more than offset by the impact of lower production costs."870 Scale made meat cheaper. I accept that. My argument is not that the large plants are inefficient. It is that their efficiency is bought partly with the spare capacity that a country needs when a plant goes down, and that the country is not told the price of that trade until it is paid.

VII.

The meat industry supplies the clearest example I have found of how data coordinates a market without anyone agreeing to anything. In September 2023, the Justice Department sued a company called Agri Stats, which sells benchmarking reports to meat processors. The complaint describes the business: "Each week, competing processors send competitively sensitive information from their internal accounting systems to Agri Stats. After auditing and standardizing these troves of data, Agri Stats creates and distributes comprehensive reports detailing competing processors' pricing, margins, inventories, and operations."871 The crucial allegation is about design. "Agri Stats designs its reports so that a processor does not need to communicate directly with other processors to determine their intentions, but instead can look at the reports to forecast what competitors will do."872 Processors used weekly reports ranking their prices against competitors' to target low-priced products for increases — "a practice some processors refer to as 'chasing price' or 'pricing with courage.'"873 An executive at one pork processor summarized the service's advice "in four words: 'Just raise your price.'"874 A former chief executive of a poultry company, on a 2009 call: "I see a lot of information from Agri Stats that tells me that nobody's going to ramp up."875 Notice the asymmetry. The service, the government alleged, provided "processors with unique insights about their competitors' production, costs, and pricing—and refusing to sell the same information to processors' customers, farmers, workers, or consumers."876 What makes this case instructive is the remedy. In May 2026, the government filed a proposed settlement under which the company is

"prohibited from sharing any sales reports or non-public pricing information between and among competing protein processors," "prohibited from providing any information to competing protein processors that is on average less than 45 days old," and "required to make the vast majority of its reports available to all interested purchasers on reasonable and non-discriminatory terms." A monitor is appointed.877 I could not confirm whether the court has entered the final judgment.

The principle is symmetry. If an information service is going to exist, the people on the other side of the transaction get to see it too. That is a rule that could be written generally, and I will propose it.

VIII.

Now to the farm itself, and to the machine the farmer drives. A modern tractor or combine is a computer on wheels, and its repair increasingly runs through software. In January 2025, the Federal Trade Commission, with Illinois and Minnesota, sued Deere & Company. The complaint's opening paragraph states the case: "For decades, Defendant Deere & Company… has throttled the ability of farmers and independent repair providers ('IRPs') to repair Deere equipment, leaving farmers wholly reliant upon Deere's network of authorized dealers ('Deere dealers') for many key repairs. Deere's increasingly sophisticated agricultural equipment requires a software tool to diagnose and repair problems that relate to electronic functions, and only Deere has the information and knowledge to create this essential tool."878 The full-function tool went only to dealers; farmers could buy "an inferior repair tool" that "lacks many of the key functions necessary to enable comprehensive repair."879880 The alleged result: "Deere's dealers are able to maintain a 100% market share and charge supracompetitive prices for restricted repairs."881 The harm is specifically agricultural. The restrictions "prevent them from more reliably planting, spraying, or harvesting crops on a schedule that would allow them to maximize yield."882 A combine down for a week in harvest is not an inconvenience. The company's defenses were environmental, safety, and intellectual-property concerns. The complaint called those "pretextual."883 The Commission's 2021 report had reached a similar judgment across industries: "there is scant evidence to support manufacturers' justifications for repair restrictions."884 I want to state this carefully. Emissions tampering is a real concern, and safety systems are real. The finding is not that such concerns never exist; it is that the record did not show they justify the restrictions actually imposed. In July 2026 the Commission announced a settlement requiring Deere, "for the next 10 years and under the supervision of the FTC and plaintiff states—to provide farmers and independent repair providers with the same equipment repair resources, including applicable software capabilities, that it currently provides to authorized Deere dealers."885 Stipulated orders take effect when the court signs them; I could not confirm that it has.

A state had already acted. Colorado's 2023 agricultural right-to-repair law requires a manufacturer to make available, "with fair and reasonable terms and costs," to "an independent repair provider or owner of the manufacturer's equipment any documentation, parts, embedded software, embedded software for agricultural equipment, firmware, or tools, or, with owner authorization, data that are intended for use with the equipment."886 It forbids parts-pairing887 and refusal on trade-secret grounds.888 It preserves the emissions laws.889 And it contains a provision that should worry anyone thinking about federal action: the article "will be repealed if the United States Congress enacts federal legislation establishing a right to repair agricultural equipment."890 A weak federal law would erase a strong state one.

IX.

The repair fight is about the machine. The quieter fight is about the data the machine collects.

Precision agriculture is not a niche. "Auto-steer guidance systems were used on only 5.3 percent of planted corn acres in 2001, growing to 58 percent in 2016," and by 2019 "GPS applications were used on 40 percent of all U.S. farm and ranchland acreage."891 The Government Accountability Office's 2024 assessment named the concern directly. "Increased precision agriculture technology use has raised concerns about ownership and control of farm data. More farmers are adopting precision agriculture technologies, and new devices are collecting increasingly larger amounts of real-time data on more farms. However, some farmers are concerned about losing ownership or control of their farm data if they decide to share that data with others."892 It noted the lock-in problem: "the industry lacks standards to drive compatibility between precision agriculture technologies," while "farmers often do not have a homogeneous set of equipment from a single manufacturer."893 The manufacturer's own filings show why the data matters to the business. Deere measures customer engagement by "engaged acres," a count of "unique acres with at least one operation pass recorded in the Operations Center in the past 12 months." It plans "to enhance how we deliver value by investing in a Solutions as a Service business model."894 I do not criticize a company for wanting recurring revenue. I note that a farm whose field history, yield maps, and application records live in one vendor's platform, and whose equipment reports to that platform, has a switching cost that no one wrote into the purchase price of the tractor.

What does federal law say about who owns that data? Almost nothing. There is a statute protecting information farmers give to the Agriculture Department: officials "shall not disclose" information "provided by an agricultural producer… in order to participate in programs of the Department," and participation "may not be conditioned on the consent" of the producer to disclosure.895 Those are good rules. They apply to the government. I found no comparable federal rule for the data held by private equipment makers and platforms. The farm bill passed by the House in 2026 would direct the department to develop "voluntary, consensus-based, private sector-led interconnectivity standards."896 Voluntary and private-sector-led. At the time of writing the House and Senate versions had not been reconciled.897 So the public protects the farmer's data from the public, and leaves it to contract when it is held by the vendor. That is backwards.

X.

Water is the input no farm can route around, and in the American West its allocation is now a matter of measurement and model as much as law. The Colorado River Compact divided the river on the assumption of a flow it no longer has. The Bureau of Reclamation put the gap plainly in July 2026: the Compact "apportioned 7.5 million acre-feet to both the Upper and Lower Basins based on the estimated pre-compact average inflow (1906-1921) of approximately 18 million acre-feet. However, inflows from 2000 to 2024 have only averaged 12.9 million acre-feet."898 "Since 2000 water use has exceeded total system inflow in most years."899 The first shortage was declared in August 2021, cutting Arizona by "512,000 acre-feet, which is approximately 18% of the state's annual apportionment."900 The rules governing the river were "all set to expire at the end of 2026."901 The states did not agree on what should replace them. Reclamation's decision page says it without ornament: "Consensus on long-term operations did not materialize. Instead, operating guidelines for a 2-year interval pursuant to a 10-year Decision Framework will be implemented by October 1, 2026."902 The August 2026 guidelines reduce Lower Basin deliveries by 1.25 million acre-feet in each of the next two years.903 When there

is not enough water, the question of who used how much becomes the question of who gets cut. Here the federal system is a partial model of what this book asks for. Reclamation publishes an annual water accounting report that "tabulates measured diversions, measured and unmeasured returns, and consumptive use of each user."904 It explains its method: "Unmeasured returns are calculated as a coefficient multiplied by the water user's measured diversion amount."905 It uses satellite-based estimates of evapotranspiration to refine those coefficients and verify conservation claims.906

This is exactly the sort of machinery that decides survival in a dry year: a coefficient, a satellite model, a verification rule. The difference from the grocery allocation software is that this machinery is published. Reclamation has even acknowledged its own changes of method: over time its revisions have "resulted in inconsistencies across the reports, when viewed continuously."907 I count that as a virtue. A method that admits its revisions can be argued with. The risk is in what happens next. Irrigation scheduling, soil-moisture sensing, and yield prediction are moving onto the same proprietary platforms described in the previous section. As conserved water becomes a creditable, tradable asset — the record of decision requires Reclamation to verify claimed conservation and publish the results908 — the model that estimates a farm's use becomes a model that determines its allocation. When that happens, the public-accounting standard Reclamation already follows should be the floor for any model whose output decides who gets water.

XI.

From food and water to medicine. In 2019 an interagency task force led by the Food and Drug Administration reported on why drug shortages persist. It studied "163 drugs that went into shortage in the 5-year period between 2013 and 2017," of which 63 percent were sterile injectables and 67 percent had a generic version.909 "In the year prior to going into shortage, the median per unit price was $8.73 for all the shortage drugs, $11.05 for injectables, and $2.27 for orally administered drugs."910 The first root cause the task force identified is the thesis of this chapter in pharmaceutical form. "Lack of Incentives to Produce Less Profitable Drugs. When market conditions limit manufacturers' profitability, they reduce a firm's motivation to maintain a presence in, or enter the market for older prescription drugs, and to invest in manufacturing quality and redundant capacity… Current contracting practices contribute to a 'race to the bottom' in pricing."911 The second: the market does not "reward drug manufacturers with price premiums for mature quality management, back-up manufacturing capabilities, or risk-management plans."912 The Department of Health and Human Services put it more bluntly in 2024. The key issues include "prices for generic drugs that are driven to levels so low that they create insufficient incentives for redundancy or resilience-oriented manufacturing, distribution, and purchasing. These market failures lead to pharmaceutical supply chains that are brittle, disruption-prone, and too slow to recover from shortages."913

Its analysts counted the result. "Between 2018 and 2023, a total of 258 unique active ingredients… went into national shortage." Injectable shortages last "roughly twice as long" as oral ones — a median of 4.6 years. "Shortages of essential products last much longer (median = 4.0 years)."914 A median shortage of four years for an essential medicine. That is not a supply shock. That is a steady state.

XII.

If I had to choose one episode to explain the whole chapter, I would choose infant formula in 2022. The product is, for many infants, the only food they eat. The FDA's long-term strategy states: "In 2022, four

companies controlled 99 percent of the infant formula market; in 2024, the number of major manufacturers was down to three."915 The firm with the largest share ran three plants. In February 2022, one of them, in Sturgis, Michigan, recalled product and stopped production. The Commissioner told Congress: "FDA knew that restarting the Sturgis, Michigan, facility was critical, because it was one of three plants run by a company with the largest market share, and many of its specialty formula products cannot be quickly manufactured at other facilities. We also became aware that Abbott Nutrition lacked a contingency plan to produce its lines of specialty metabolic and amino acid formulas that serve as a sole source of nutrition for thousands of infants with metabolic disorders."916 Sit with that. A regulator responsible for food safety found a plant with "significant operational deficiencies"917 and had to weigh enforcement against the fact that closing it would leave infants without food. The agency's own after-action review described the dilemma: the incident required officials "to assess and weigh risks associated with potential product contamination against risks of essential products being unavailable due to a shortage."918 That is what concentration does to the rule of law. When a single facility is too important to close, the regulator's leverage over it shrinks to what the facility is willing to concede. The review also named the missing authority. "The FDA does not have the authority to require manufacturers of infant formulas or certain medical foods to notify it when they become aware of a circumstance that could lead to a shortage of these products for U.S. consumers, or to require manufacturers to put in place risk management plans."919 Since then the agency has built a platform to track production "using data provided to the Agency voluntarily by infant formula manufacturers."920 Voluntarily, again.

Part of the concentration is a creature of public purchasing. Federal law has since 1989 "generally required WIC state agencies to use a single-supplier competitive system for infant formula," and as of August 2024, one manufacturer supplied milk-based formula to program participants "in 36 states and the District of Columbia."921 Here the public sector is not the rescuer of a fragile market. It is one of its architects, and it has the power to redesign it.

XIII.

The last mile for medicine is the pharmacy, and in rural America it is disappearing. The Federal Trade Commission's 2024 interim report on pharmacy benefit managers — the intermediaries who decide which pharmacies are in a health plan's network and what they are paid — found that "between 2013 and 2022, about ten percent of independent retail pharmacies in rural America closed." And "pharmacy closures are associated with significant declines in patient adherence to medication."922 The concentration is severe: "the six largest PBMs… manage nearly 95 percent of all prescriptions filled in the United States."923 Medicare's geographic access rules give some rural pharmacies a little leverage, but "where substitutes or customized network conditions exist, PBMs' internal documents show that those rural pharmacies may need to operate under worse, sometimes take-it-or-leave-it rates."924 The Commission's second interim report, in January 2025, found something that deserves to be described as routing. "Dispensing patterns suggest that the Big 3 PBMs may be steering highly profitable prescriptions to their own affiliated pharmacies (and away from unaffiliated pharmacies)," and they "reimbursed their affiliated pharmacies at a higher rate than they paid unaffiliated pharmacies on nearly every specialty generic drug examined."925 The intermediary decides where the prescription goes. When a chain closes stores, the logic is written down. Rite Aid's lender materials, filed with the SEC shortly after its 2023 bankruptcy petition, explain that when a store closes, the company will "decide whether to sell the scripts or to pour those scripts to a nearby location to improve that store's performance."926 The listed factors are "Financial Performance," "Market / Competitive Landscape," "Store Profile / Format,"

"Rent Relative to Market," "Regional Presence and Market Strategy," "Supply Chain and Regional Operations," and "Script and Inventory Liquidation Value."927

Every factor concerns the company. None concerns the town — the distance to the next pharmacy, the number of patients on maintenance medication, whether the nearest alternative is thirty miles away. I do not say that is wrong for a company in bankruptcy; its duty runs to its creditors. I say that a decision that ends a community's access to medicine is being made on a scorecard that does not contain the community, and that the public has no standing to ask that it should.

XIV.

Fuel and freight are the arteries, and both have shown how quickly a lean system stops. On May 7, 2021, a ransomware attack led the Colonial Pipeline Company to shut down its system. Its chief executive told the Senate: "The product we transport accounts for nearly half of the fuel consumed on the East Coast, providing energy for more than 50 million Americans," including fuel for "airports, ambulances and first responders." The shutdown was fast: "Overall, it took us approximately fifteen minutes to close down the conduit, which has about 260 delivery points across 13 states and Washington, D.C."928 The Energy Information Administration recorded the consequence: the national average price for regular gasoline reached "$3.03 per gallon (gal) on May 17, the first time retail gasoline prices averaged more than $3.00/gal since October 27, 2014," with "outages at many gas stations" in the Lower Atlantic region.929 Federal regulators issued an hours-of-service exemption for fuel truckers, allowed overweight fuel loads on interstates, granted emergency fuel waivers across more than a dozen states, and approved a Jones Act waiver.930 The redundancy did not exist in the system; it had to be improvised by suspending the rules. At the time, the Transportation Security Administration had not issued pipeline cybersecurity regulations, "relying instead upon industry compliance with voluntary guidelines."931 Freight rail tells the same story more slowly. In April 2022, the chairman of the Surface Transportation Board opened an emergency hearing: "I have raised concerns about the primacy Class I railroads have placed on lowering their operating ratios and satisfying their shareholders even at the cost of their customers. Part of that strategy has involved cutting their work force to the bare bones in order to reduce costs. Over the last 6 years, the Class Is collectively have reduced their work force by 29% – that is about 45,000 employees cut from the payrolls."932 On-time performance at the four largest carriers "had fallen from a pre-pandemic average of 85% in May 2019 to just 67% in the last week of May 2022."933 A labor witness drew the redundancy conclusion in the plainest terms: "It is no mystery why the railroads can't provide flexibility or surge capacity right now."934

The Board responded in 2024 with a rule allowing a captive shipper whose service fell below defined standards to petition for access to a competing railroad — a remedy that had essentially fallen into disuse: "no reciprocal switching orders have been issued since before 1985, and none have even been sought since 1989."935 In July 2025 the Seventh Circuit vacated the rule, holding that the statute requires a finding that the existing carrier's service is inadequate: "we GRANT the petition, VACATE the Final Rule, and REMAND to the Board."936 The court's reading of the statute may well be right. What it means is that the remedy has to come from Congress.

XV.

Movement is not only freight. It is people getting to the goods. The Bureau of Transportation Statistics reports that "in 2021, 85% of the nation's 88.8 million rural residents had reasonable access to intercity transportation,"

meaning a large airport within 75 miles or a smaller airport, bus stop, or rail station within 25.937 Excluding residents on the fringe of cities "increases the percentage of rural without reasonable intercity transportation access from 15% to 25%." In South Dakota, "62% of rural residents lacked reasonable access." Low-income households, 17 percent of rural households, "accounted for a disproportionately larger share (24%) of the rural households lacking reasonable access."938 This measures intercity travel, not local transit to a clinic or a store. The Government Accountability Office's work on older rural Americans fills in what the national figure hides. Officials described "long distances with little public transportation"; where transit existed, services "were often limited to certain days of the week or month, areas, or functions." In a Mississippi town that had lost its transportation service, "residents now must travel 25 miles to the nearest grocery store."939 A separate report found that transportation barriers "are a significant impediment to health care, leading to missed appointments and delayed care," and that the one federal transit program specifically aimed at medical trips had no evaluation plan to tell whether it worked.940 The Transportation Department defines transportation insecurity as occurring "when people are unable to get to where they need to go to meet the needs of their daily life regularly, reliably, affordably, and safely."941

The grocery store closed, the pharmacy closed, the mail slowed, and the bus runs on Tuesdays.

XVI.

When all of that fails at once, we call it a disaster, and we expect a public system to bring food and water. The record of what happened when it tried is sobering.

After Hurricane Maria, the Department of Homeland Security's Inspector General found that the Federal Emergency Management Agency "lost visibility of about 38 percent of its commodity shipments to Puerto Rico, worth an estimated $257 million. Commodities successfully delivered to Puerto Rico took an average of 69 days to reach their final destinations."942 Meals ready to eat were unavailable because "advance contract vendors were at capacity in providing nutritional meals following Hurricane Harvey and could not produce any more to support the effort in Puerto Rico." Puerto Rico officials reported "one meal for approximately every 12 snack boxes distributed to a survivor."943 The agency's own after-action report drew two conclusions that pull in opposite directions, and both are right. First: "closer partnerships with the private sector are crucial… No jurisdiction or federal agency has all the staff and resources it will need to respond to a catastrophic incident." Second: "governments need to be better prepared with their own supplies, to have pre-positioned contracts with enforcement mechanisms."944 After Irma, the agency "distributed more than 80 percent of its inventory for selected commodities from the Caribbean Distribution Center warehouse. Hurricane Maria struck before supplies were replenished."945 After Hurricane Helene in 2024, the geography itself failed. The administrator told Congress that floodwaters washed out "critical transportation routes, such as portions of I-40," and that "in western North Carolina alone, a total 655 roads were closed. As a result, many communities were isolated."946 I could not find an independent audit of the Helene supply response comparable to the one after Maria.947 The lesson is not that government logistics is better than private logistics. On the Maria evidence, it plainly was not. The lesson is that the public cannot contract out a capacity it does not itself understand, and that emergency supply assembled at the moment of need from vendors already at capacity is not supply at all. The capacity has to exist, and be exercised, before the storm.

XVII.

That leaves the routing decision in its purest form: the decision not to deliver to a place at all.

There is a modest federal precedent. In 2000, the Justice Department resolved a complaint that a national pizza chain's "policy of providing only limited pizza delivery in certain geographical areas had a discriminatory effect on African Americans." The company agreed that "decisions to limit pizza delivery are based upon a legitimate safety concern for their delivery drivers," that it could limit delivery only "where there is evidence that the safety of delivery drivers is threatened," that "the scope of any delivery limitations… must be narrowly confined to the area in which safety is a concern," and that "store managers also must conduct an annual review of any decision to limit delivery."948 Evidence, narrow scope, annual review. A complete framework for a routing exclusion in three phrases. The more ambitious model was federal and is now gone. In 2023 the FCC adopted rules against "digital discrimination of access," defined as "policies or practices, not justified by genuine issues of technical or economic feasibility, that differentially impact consumers' access to broadband internet access service based on their income level, race, ethnicity, color, religion, or national origin or are intended to have such differential impact."949 In May 2026 the Eighth Circuit vacated the rule, holding that "Congress did not authorize disparate impact liability."950 The definition itself is a good template for any essential service delivered by route: parcel, grocery, pharmacy, transit. The court did not say such a standard is unwise. It said the agency lacked authority to impose it without Congress saying so. So Congress should say so, in terms, for the services this chapter covers.

XVIII.

This book judges power by a fixed standard rather than by whether I like the people holding it.

The standard names concentrations that no authority, public or private, may hold together; two of them are the subject of this chapter — food and material provision, and transit and movement. Let me put the record to the book's three questions. Can it be witnessed? Where the public's stake is largest, the logic is least visible. The Postal Service's network models became legible only because a regulator ran a place-by-place analysis the Service had not published. The grocery allocation contracts became visible only because the Federal Trade Commission compelled nine companies to produce them. The benchmarking reports that let meat processors forecast one another were sold to processors and withheld from farmers and consumers until the Justice Department sued. The farm data sits in vendor platforms under contract terms the Accountability Office says farmers do not trust. The one clear exception is water: Reclamation publishes its accounting method, its coefficients, and its revisions.

Who carries the burden? The household on the thinnest margin, every time. The smaller store got proportionally less product during the shortage. The poorest fifth of households spends almost a third of its income on food. The rural ZIP Code pair was downgraded at nearly twice the non-rural rate. The rancher had no other place to take his animals when a plant closed. The infant with a metabolic disorder depended on one plant's specialty line. In every case the saving is spread across the network and the cost lands on the end of the route. Is there a way back? For a town cut from a route, no. The statute says no small post office shall be closed solely for losing money, but the regulator's opinion is advisory. The captive rail shipper's one remedy was vacated. The broadband discrimination rule was vacated. A community whose pharmacy's prescriptions were "poured" to another store has no forum at all. Return is available where a court or a regulator has intervened, and nowhere else.

XIX.

The first objection: efficiency is why food is cheap. Largely true.

The Agriculture Department's own summary found that the cost savings of large packing plants more than offset the market power they gained over cattle prices.951 I am not proposing to undo it. I am proposing to stop pretending it is free. The Council of Economic Advisers identified the precise failure: a firm that underinvests in resilience imposes costs on everyone downstream that it never pays.952 Where the downstream cost is an empty shelf of paper towels, let the market sort it out. Where it is an empty shelf of infant formula, it cannot.

The second objection: the Postal Service is losing money and must cut costs. Yes, and its losses are real. But the regulator found that the savings from the service-standard changes would be about 4.4 percent of operating expenses, and the Inspector General found that five years of network changes had not materially improved the Service's financial position.953 If the cuts do not fix the finances and do degrade rural service, the argument for them is weaker than it looks.

The third objection: repair restrictions protect safety and emissions. Some do, and a law can say so. Colorado's does: it expressly denies any right to "evade emissions, copyright, trademark, or patent laws."954 What the record does not support is the claim that every restriction actually imposed was necessary for safety. The Federal Trade Commission, after a public workshop and review of comments, found "scant evidence" for the justifications offered.955 A narrow safety exception is easy to write. A general veto disguised as one is what we have.

The fourth objection: mandated redundancy is a hidden tax. It is a cost, and I will not pretend otherwise. The Council was right that it is "not cost-effective for firms to invest in completely avoiding all disasters."956 That is why my proposal is limited to essentials — food, water, medicine, fuel, the means of reaching them — and asks first for measurement and planning rather than spare plants everywhere. The Food and Drug Administration already speaks of redundancy risk management plans for formula;957 the health department has asked Congress for authority to require resilience programs for drugs.958 Those are the right scale of intervention. The fifth objection: government is worse at logistics than business. Sometimes it is. After Maria, the federal government lost track of more than a third of its shipments. I have quoted that record against myself on purpose. But notice what failed: the agency depended on vendors already at capacity, used up its stock on the first storm, and did not track what it had contracted for. The answer is not to rely on private logistics more completely. It is for the public to hold, understand, and exercise a real capacity of its own, so that when the private system is at capacity or has withdrawn, there is something there. X

X.

Every row is the same shape. A decision that once allowed for slack — a second supplier, a backup plant, a spare day, a route that loses money — is now made by a model that treats slack as waste. And the model was built by a party that does not live at the end of the route. Here is what I would insist on. No individualized pricing of essentials. For food, infant formula, prescription medicine, home fuel, and water, I would prohibit it outright: one posted price for everyone in a market at a given moment, with lawful discounts available to all who qualify on published terms. A seller may run a sale. It may not charge a family more for bread because a model inferred the family could not go elsewhere. Route-impact statements, measured by place. Before any

network change by a carrier of essential goods — the Postal Service, a pharmacy chain above a size threshold, a regional grocery distributor, a Class I railroad — the operator must publish its effect by ZIP Code, not by volume. The Postal Regulatory Commission has shown the method. Make it mandatory and prior, not advisory. Disclosure of pricing and routing logic that affects essentials. Where an algorithm sets the price of an essential, allocates it in shortage, or decides which places a delivery network will serve, its operator must file the decision rules with a public regulator and explain them in plain language to anyone affected. Trade secrecy may protect the code. It may not hide the criteria.

Symmetry of market information. Any service that sells competitors' pricing, cost, or inventory data to sellers must sell the same reports, on the same terms, to buyers, farmers, and workers — the principle already in the proposed benchmarking settlement, written into law. Mandatory redundancy plans for essential supply. For every product the government designates essential, the manufacturer must maintain and file a redundancy risk management plan, notify the regulator of any event likely to cause a shortage, and identify its single points of failure. And end single-supplier public purchasing for essentials where a second supplier is available. Right to repair and data portability, with a federal floor that does not preempt. Owners and independent repairers get the same tools, software, parts, and documentation as authorized dealers, on fair terms, with parts-pairing prohibited and a narrow emissions and safety exception written in. Farm data belongs to the farm, is exportable in open formats, and may not be used by the vendor for purposes the farmer did not authorize. Any federal law must be a floor, expressly preserving stronger state laws, so that Colorado's statute is not repealed by something weaker. Public accounting for any model that allocates water. When a model's output determines an allocation, cut, or conservation credit, its method, coefficients, and revisions must be published to the standard Reclamation already meets on the lower Colorado. Universal service for delivery and transit. Extend the principle of the 1970 postal statute to the carriers of essential goods. No essential delivery zone may be withdrawn solely because it loses money; any exclusion must rest on published evidence, be narrowly drawn, and be reviewed every year — the three conditions a pizza company accepted in 2000. Public capacity to provide when a private system withdraws. Every state should have a standing, funded, exercised capability to deliver food, water, medicine, and fuel to its own communities: stock that is replenished after use, contracts with enforcement terms, vehicles, and people who have done it before the storm rather than learning how during it. Where a private provider exits a community's last pharmacy or last grocery, a public or cooperative alternative should have the right and the financing to step in. A market may allocate luxuries by algorithm. It may not quietly reprice survival, and it may not decide by routing model which places continue to exist.

The people who wrote the 1970 postal statute understood something= the language of optimization has made hard to say. A country is not= only a market. It is also a promise, made to every place inside its= borders, that being small, poor, and far away does not put you= outside the country. The promise costs money. It always has. The= question is whether we will keep paying for it on purpose, in public,= or let it be priced away one network change at a time, by models we= are not allowed to see, until one morning the truck does not come. Keep the promise. Write it down again, for the routes that matter= now. — D.M.D.

What Exchange Requires

1. Workers govern the tools that govern their hiring, scheduling, pay, evaluation, discipline, and termination.

2. Productivity gains must not be converted automatically into unemployment, wage suppression, or intensified surveillance.

3. Public subsidies carry enforceable public-return requirements.

4. Ratepayers and taxpayers are not the backstop for private infrastructure costs.

5. Contracts, incentives, and cost allocations are public where the public bears the risk.

6. Community benefit is measured by durable capacity, not sentiment, publicity, or unenforceable promises.

A Practice for the Reader

Choose a local subsidy, incentive agreement, major employer, or infrastructure project. Find the public terms. Ask: What did the public provide? What returns are promised? Who verifies delivery? What happens if the promise is broken?

Coda: Not Everything Can Be Priced

A fair distribution of benefits and burdens is necessary. It is not sufficient. Some harms cannot be justified by compensation. Some powers cannot be redeemed by convenience.

There are forms of surveillance, prediction, classification, and automated denial that remain unacceptable even when a market rewards them or an institution calls them efficient.

The next part draws those lines. It asks what must never be traded away in the name of progress.

NEXT: PART FIVE — THRESHOLD

HUMAN RIGHTS AND HUMAN DIGNITY

The Line That Does Not Move

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."

