What Is There

WHAT IS THERE

An Inventory of the Systems We Live In

J. Konstapel — Leiden, 2026

Contents

What Is There

An Inventory of the Systems We Live In

J. Konstapel, Leiden, 2026

Introduction: How to Read This Book

This book is an inventory. It records the current state of seven systems: science, engineering and energy, computing, money, institutions, the human being, and the earth. For each system it states what is there — how large it is, what it consists of, what it costs, what it loses, and how it has moved over time.

That is all Part I does. It contains no argument, no accusation and no theory. Every figure in it is documented, mainstream, and in almost every case produced by the system it describes: the physics community's own parameter tables, the energy laboratories' own flow charts, the publishers' own margins, the central banks' own money mechanics, the ministries' own evaluations, the statistical offices' own life measurements, the satellites the chain built to watch itself. A reader who distrusts any number in this book can walk to the source; the sources are the institutions.

One boundary of the inventory should be stated at the start. This book counts the cost side: the losses, the apparatus, and the ratio between the apparatus and the work. The delivery side is real, is larger, and is not in dispute. The same chains doubled life expectancy within a century, fed eight billion people, eradicated diseases, and put the recorded knowledge of the species in every pocket. That record is why every layer in this book was accepted at the moment it was added. The inventory does not weigh the two sides against each other; it measures how the cost side has been moving. Both statements can be true at once — the chains delivered, and their apparatus now grows faster than their work — and the second is the one this book documents.

Three rules governed the writing, and knowing them will help the reading.

First: no interpretation before the record is complete. Each chapter ends with a bare summary and one repeated observation — the system grows more complex, loses more, and nobody chose it — and then stops. No chapter explains another. No chapter names a culprit, because the record shows none: every layer, every rule, every link was added by someone for a defensible reason. If a pattern forms across the chapters, it forms in the reader. That is by design. A pattern the reader assembles is worth more than one the author asserts, and it cannot be taken away by disputing the author.

Second: the answers come as facts too. Appendix A records some twenty working answers — a care organization, a currency, a computer, a school, a check dam, an animal — each of which exists, runs or ran for decades, and has been measured outperforming the standard system. They are inventoried on the same terms as the problems: numbers, sources, no advocacy. The appendix ends by stating the one property all entries share. It states it as an observation about the list, because that is what it is.

Third: the description comes last and adds nothing. Appendix B gives, in a few pages, one description in which the seven chapters and the twenty answers are the same picture. It has a name and a body of work behind it, referenced there. But it is deliberately placed where it can be skipped. Part I stands without it. A reader who closes the book after Appendix A leaves with a complete, sourced inventory and a pattern of their own. A reader who continues finds that the pattern already has a language, arithmetic, and running implementations.

This order — facts, then working answers, then description — is the reverse of how books about systems are usually built. The usual book opens with a thesis and marshals facts toward it, which invites the reader to judge the thesis and discount the facts. This book cannot be read that way, because for seven chapters and one appendix there is no thesis to judge. There is only the record, and the record belongs to everyone.

One more thing should be said about tone. An inventory of losses can read as pessimism. It is not. Half of this book — the appendix of working answers, and everything it implies — is the opposite. The systems described in Part I are not failing because people are foolish or malicious; the record shows neither. They are failing in a particular, consistent, measurable way, the same way in every domain — and the answers that work, work in a particular, consistent, measurable way too, the same way in every domain. A problem with that much regularity is not a tragedy. It is a structure. And structures, once seen, can be worked with.

That is what the book offers: the seeing. What is there, first. What already works, second. What it all is, last — for those who want the name.

Begin with Chapter 1. It is about the system that produced almost every number in this book, which makes it the right place to test the method: read it, check a source or two, and decide whether an inventory can be trusted. The other six chapters ask nothing more than that.

The chapters may be read in any order after the first; each is self-contained, with its own annotated references. Readers pressed for time can read the final section of each chapter — "The record" — and Appendix A in full.

Chapter 1 — Science: An Inventory

Part I of What Is There — *J. Konstapel, Leiden, 2026*

This chapter records the current state of science as a system. It contains no theory and no proposal. It states what is there, with numbers. Every figure can be checked against the sources listed at the end.

1. The volume

Science now produces more than five million journal articles per year. Indexed article totals in Scopus and Web of Science grew by about 47% between 2016 and 2022. The growth is exponential, at roughly 5.6% per year. The number of practising scientists did not grow at that rate. The workload per scientist — writing, reviewing, editing — therefore rose.

In 1980 the world published about one million articles per year. Around 2020 it published more than five million in the major indexes, and over seven million when all sources are counted. Nobody reads this. A single specialist cannot follow a single subfield any more. The literature has become larger than the profession that produces it.

The response to this volume is a second literature about the first. Review articles now make up about 3% of all papers and their share is growing in most fields. Science summarizes itself because it can no longer read itself.

2. The state of fundamental physics

The Standard Model of particle physics is the most precisely tested theory in history. It is also incomplete by its own account.

The model contains 19 free parameters in its minimal form. These include the masses of the quarks and leptons, the mixing angles, and the coupling strengths. With neutrino masses included, the count rises to 26. None of these numbers is derived. All of them are measured and inserted by hand. The theory does not say why the electron has its mass. It records that it does.

Gravity is not part of the Standard Model. The two pillars of physics — quantum field theory and general relativity — have resisted unification for a century. Each works in its own domain. Where the domains meet, at black holes and at the beginning of the universe, the mathematics breaks down.

The main candidate for unification, string theory, has produced an estimated 10^500 possible vacuum states. This is not a prediction. It is a landscape in which almost any observation can be accommodated. After roughly fifty years of work, string theory has produced no experimentally confirmed result.

3. The dark sector

The standard cosmological model states that ordinary matter — everything ever observed in a laboratory — makes up about 5% of the universe. The remaining 95% consists of two postulated components: dark matter, about 27%, and dark energy, about 68%.

Neither has been directly detected. Dark matter was introduced to explain galaxy rotation curves that the equations did not reproduce. Dark energy was introduced to explain an accelerating expansion that the equations did not predict. Both are correction terms with names.

Decades of direct-detection experiments for dark matter particles have produced null results. The parameter space of candidate particles shrinks with each experiment, and the searches continue at increasing cost.

Meanwhile the model shows internal strain. The expansion rate of the universe, measured from the early universe, disagrees with the same rate measured from the nearby universe. This discrepancy, the Hubble tension, has grown more significant as measurements improved. Better data made the problem worse, not better.

The summary is plain. The reigning cosmological model describes 5% of its subject directly and 95%

through postulated components introduced to save the equations.

4. Replication

In 2005 John Ioannidis published "Why Most Published Research Findings Are False." The paper argued, from the statistics of research practice itself, that a majority of published claims in many fields are likely wrong. It became one of the most cited papers in science. Its argument has not been refuted. It has been confirmed.

In 2015 the Reproducibility Project repeated 100 published psychology experiments under close cooperation with the original authors. About 36% of the replications produced significant results, against 97% of the originals. Effect sizes in the replications were roughly half of those originally reported.

Cancer biology fared no better. A multi-year effort to replicate key preclinical cancer studies found that fewer than half could be reproduced, and that replication effect sizes were dramatically smaller than the originals.

In 2016 Nature surveyed 1,576 researchers across disciplines. More than 70% had tried and failed to reproduce another scientist's experiments. More than half had failed to reproduce their own. A majority agreed there is a significant reproducibility crisis.

These are not marginal fields. Psychology and preclinical medicine feed directly into therapy, policy and pharmaceutical development. The pharmaceutical industry reported the problem first: internal replication attempts at Bayer and Amgen in the early 2010s found that only 20–25% of published preclinical findings could be confirmed.

5. Retraction

In 2023 more than 10,000 research papers were retracted — a record. In 2022 the figure was about 5,500. The retraction rate has risen above 0.2% of all published articles.

More than 8,000 of the 2023 retractions came from journals of a single publisher subsidiary, Hindawi, after its peer-review process was found to be systematically compromised. Most of the retracted papers had appeared in special issues run by guest editors.

Behind the retractions stands an industry. Paper mills are commercial firms that sell finished fake papers and authorship slots. Estimates of their share of submitted manuscripts range from 2% to 20%, depending on field. Integrity investigators regard the roughly 50,000 retractions recorded worldwide as the visible part of a much larger body of fabricated literature.

Fake papers do damage even when nobody reads them. Their results are aggregated into review articles and meta-analyses. The fabricated enters the mainstream through summary.

6. The industry

Scientific publishing is one of the most profitable industries in the world. The largest publisher, Elsevier, reports operating margins around 38%. This exceeds the margins of Apple, Google and the major oil companies.

The business model is unique. The product — research — is delivered free by scientists. Quality control — peer review — is performed free by other scientists. The result is then sold back to the universities that employed both groups, through subscriptions and article-processing charges that can exceed 10,000 dollars for a single paper in a prestige journal.

The unpaid labour is enormous. Peer reviewers worldwide contribute an estimated 100 million hours per year. One study valued the reviewing time donated by researchers in the United States, China and the United Kingdom alone at over 2.5 billion dollars per year. None of it is compensated. Meanwhile editors report growing difficulty finding qualified reviewers at all, as the article volume outgrows the reviewer pool.

Careers, hiring, promotion and funding are all indexed to publication counts and journal prestige. The phrase for this is publish or perish. It is not a metaphor. It is the incentive structure that produced the volume of section 1, the replication failure of section 4, and the paper mills of section 5.

7. The administrative layer

Research funding runs through competitive grant applications. Success rates at the major funders hover around 20%, and lower in many programmes. Four out of five applications — each representing weeks to months of senior scientists' time — produce nothing but the application itself.

Surveys of researchers consistently find that around 40% of research time is spent not on research but on administration: applications, reporting, compliance, evaluation. Principal investigators at research universities routinely describe themselves as managers and fundraisers who no longer do science with their own hands.

Every control layer was added for a reason. Fraud justified oversight. Oversight required reporting. Reporting required administration. Administration consumed the time that was supposed to be protected. Each layer is individually defensible. The stack is not.

8. The record

The inventory of this chapter, in one place.

Five million-plus articles per year, growing at 5.6% annually, unreadable by the profession that writes them. A fundamental theory with 26 hand-inserted parameters and no account of gravity. A cosmology that directly describes 5% of its subject. A replication rate between 20% and 40% in the fields that were tested. Ten thousand retractions in a single year, with a commercial fabrication industry behind them. A publishing sector with 38% margins built on unpaid labour. And a working scientist who spends 40% of their time on administration.

Each of these facts is documented, mainstream, and undisputed in its own literature. No single person designed this system. No committee approved it. It grew, one defensible step at a time, into a structure that produces more, verifies less, and costs more per verified result every year.

That is what is there.

Annotated References

Hanson, M. A., Gómez Barreiro, P., Crosetto, P. & Brockington, D. (2024), "The strain on scientific publishing," Quantitative Science Studies. Why read? The definitive quantitative record of publication growth: 47% more indexed articles in 2022 than in 2016, exponential growth at 5.6% per year, and the widening gap between article volume and the reviewer pool. The factual backbone of section 1.

Particle Data Group, "Review of Particle Physics" (current edition). Why read? The authoritative listing of every measured Standard Model parameter. Reading the tables makes the point directly: the numbers are recorded, not derived. Read the summary tables first; the full review is thousands of pages.

Planck Collaboration (2020), "Planck 2018 results. VI. Cosmological parameters," Astronomy & Astrophysics. Why read? The source of the 5%–27%–68% division of the universe into ordinary matter, dark matter and dark energy, from the most precise cosmological dataset available.

Ioannidis, J. P. A. (2005), "Why Most Published Research Findings Are False," PLoS Medicine. Why read? The founding document of the replication debate. Short, statistical, and written before the confirming evidence arrived. Two decades later, its argument stands.

Open Science Collaboration (2015), "Estimating the reproducibility of psychological science," Science. Why read? The largest coordinated replication effort ever conducted: 100 studies, 36% replicated. Notable for its cooperative method — original authors helped design the replications, removing the usual excuse.

Errington, T. M. et al. (2021), "Investigating the replicability of preclinical cancer biology," eLife. Why read? The medical counterpart to the psychology project. Fewer than half of key cancer-biology findings replicated, with much smaller effects. Directly relevant to anyone who takes medicine.

Baker, M. (2016), "1,500 scientists lift the lid on reproducibility," Nature. Why read? The survey that quantified what researchers already knew: over 70% had failed to reproduce a colleague's experiment, over half had failed to reproduce their own.

Van Noorden, R. (2023), "More than 10,000 research papers were retracted in 2023 — a new record," Nature. Why read? The record year in retractions, the Hindawi collapse, and the country-level retraction rates. Includes the crucial observation that recorded retractions are the visible fraction of the fabricated literature.

Aczel, B., Szaszi, B. & Holcombe, A. O. (2021), "A billion-dollar donation: estimating the cost of researchers' time spent on peer review," Research Integrity and Peer Review. Why read? Puts a monetary figure on the unpaid labour that carries the publishing industry: billions of dollars of donated reviewing time per year, in three countries alone.

Buranyi, S. (2017), "Is the staggeringly profitable business of scientific publishing bad for science?," The Guardian. Why read? The most readable account of how the publishing business model came to be, including the origin of the 38% margins. Journalism, but carefully sourced and widely confirmed.

Chapter 2 — Engineering and Energy: The Long Chain

Part I of What Is There *— J. Konstapel, Leiden, 2026*

This chapter records how the industrial world handles energy. It follows the chain from source to use and states the measured loss at every link. It contains no theory and no proposal. Every figure can be checked against the sources listed at the end.

1. Two thirds

The United States is the best-documented energy system in the world. Lawrence Livermore National Laboratory publishes its full energy flow every year, in one diagram.

The 2024 diagram reads as follows. Total energy input: 94.2 quadrillion BTU. Energy that reached an end-use sector: 74.9. Energy delivered as actual services — motion, light, heat where heat was wanted, computation: about 33. Rejected energy — released to the environment without performing a service: 61.5.

Two thirds of everything extracted, refined, transported and burned performs no service. It leaves as waste heat. This ratio has been roughly stable for decades. Efficiency programmes, better engines and better buildings have not moved it materially, because the losses are not accidents. They are properties of the chain itself. A second mechanism holds the ratio in place and has been documented since 1865: efficiency gains are absorbed by growth in use. Jevons observed it with coal; the modern record repeats it with lighting, transport and computing — the cost per unit falls, the units multiply, and total consumption rises. The literature calls it the rebound effect.

The global picture is the same in kind. World primary energy consumption stands at roughly 620 exajoules per year. The majority is converted through combustion and heat engines, which carry the same structural loss everywhere they run.

2. The chain, link by link

The dominant energy path in the industrial world runs: extract, refine, transport, burn, convert, transmit, convert again, use. Each link takes its share.

Extraction and refining. Producing a fuel costs fuel. For conventional oil, several percent of the energy content is spent lifting, processing and refining it. For deeper and heavier resources the share rises steeply. Ore grades tell the same story on the material side: average copper ore grades have fallen to around 0.5%, which means that roughly 200 tonnes of rock are moved, crushed and processed for one tonne of metal.

Combustion for electricity. A conventional coal plant converts 33–40% of the fuel's energy into electricity. Modern combined-cycle gas plants reach about 60%. Everything else leaves through the cooling tower. This is not poor engineering. It is the Carnot limit: a heat engine can only convert the temperature difference, and the rest must be dumped. The cooling tower is not an accessory. It is the design.

Transmission. Moving electricity through the grid loses about 5% on average, more over long distances and in older networks.

The car. A petrol engine delivers roughly 20–25% of the tank's energy to the wheels. About 75–80% becomes heat in the engine block, the radiator and the exhaust. The arithmetic continues past the tank. The car weighs 1,500–2,000 kilograms and the passenger weighs 75. Of the energy that reaches the wheels, about 95% moves the vehicle and 5% moves the person. Combining the two ratios: roughly 1% of the fuel's energy performs the actual service — moving a human from one place to another.

Lighting, as the exception that proves the pattern. The incandescent bulb converted about 5% of its electricity into light. The LED converts 30–50%. Lighting is the one link where the loss was actually engineered away — and it was done not by improving the flame but by abandoning heat as the intermediate step. Everywhere the chain still passes through heat, the two-thirds loss remains.

3. Heating what was just cooled

The chain does not only lose energy. It routinely spends energy undoing its own previous step.

The power plant burns fuel to make heat, converts a third of it to electricity, and dumps the rest. The electricity travels to a building, where a large share of it drives air conditioning — machinery that removes heat. In warm climates, cooling is among the largest single electricity uses. The system heats the environment at the plant in order to cool the environment at the building.

Data centres repeat this inside one facility. Servers convert nearly all their electricity into heat. Cooling systems then consume additional electricity to remove that heat: about 7% extra in the most efficient hyperscale centres, over 30% in ordinary enterprise centres. The industry measure for this, PUE, is a ratio of total energy to computing energy. A PUE of 1.5 means: for every two units of computing, one unit is spent carrying away the heat the computing just made.

District heating projects invert the same loop at city scale. Residual heat from industry — the rejected two thirds of section 1 — is piped tens of kilometres to homes, through insulated pipelines that lose heat along the way, at capital costs of billions, to deliver back a fraction of what the source plant threw away. The waste of one link becomes the raw material of the next, with a new infrastructure and a new loss in between.

4. The digital layer

Computation was expected to dematerialize the economy. It has instead added a new consumption layer on top of the old ones.

Data centres consumed about 415 terawatt-hours of electricity in 2024 — around 1.5% of global electricity, growing at 12% per year, four times faster than total electricity demand. The International Energy Agency projects roughly 945 TWh by 2030, more than the entire present electricity consumption of Japan. In the United States, data centres are expected to account for nearly half of all electricity demand growth to 2030, and to consume more power than the production of aluminium, steel, cement and chemicals combined.

The driver is artificial intelligence. AI-focused server electricity use is projected to grow around 30% per year. A single generative-AI query consumes several times the energy of a conventional search. Inference — the routine running of models — already dominates, at 80–90% of AI computing.

Underneath the growth sits the register. Digital hardware is binary. A binary circuit has no rest state: memory must be refreshed, clocks must run, signals must be actively held at their levels. Every stored bit and every idle server draws power to remain what it is. The heat this produces is then removed by the cooling of section 3. The digital layer thus inherits the full structure of the chain: continuous input, continuous loss, and machinery to dispose of the loss.

5. Food

The food system runs the same chain with calories.

Industrial agriculture spends roughly 8–10 calories of fossil energy to deliver 1 calorie of food to a plate in a developed country. The inputs are fertilizer, diesel, processing, packaging, refrigeration and transport. Nitrogen fertilizer alone, produced through the Haber-Bosch process, consumes about 1–2% of the world's total energy supply.

At the end of this chain, roughly one third of all food produced is lost or wasted — about 1.3 billion tonnes per year. The energy embedded in that wasted food, having passed through every link of the chain, is discarded with it.

A century ago the ratio ran the other way: farming returned more calories than it consumed, because the sun did the work. The industrial food chain reversed the sign and called it productivity.

6. The response: a second chain

The declared answer to the losses of the fossil chain is the energy transition. As currently built, it is a second chain constructed beside the first.

One counter-series belongs in this record before the losses are counted: the cost of solar electricity fell about 90% between 2010 and 2020, making new solar the cheapest electricity in history in favourable locations. The panels are not the problem. The inventory below concerns the shape built around them — conversion, storage, transmission, backup.

Solar panels and wind turbines produce electricity intermittently. To match supply to demand, the system adds storage: batteries, at a round-trip loss of 5–15%; or hydrogen, produced by electrolysis at about 70% efficiency, stored, and converted back at about 50% — an overall round trip that discards half to two thirds of the input. To connect remote generation sites, the system adds thousands of kilometres of new transmission, with its own losses and its own decade-long construction times. To buffer the remaining gaps, the fossil plants are kept standing by.

On the material side, the second chain reopens the first link. Batteries, magnets and cables require lithium, cobalt, copper and rare earths, mined at the falling ore grades of section 2. The machinery of the transition is manufactured, today, mostly with the energy of the chain it is meant to replace.

None of this is hidden. Every figure in this section comes from the transition's own planning documents. The plan, stated plainly, is to rebuild the entire chain — extraction, conversion, transmission, storage, backup — in a new material, while keeping its structure intact.

7. The record

The inventory of this chapter, in one place.

Two thirds of all extracted energy performs no service and leaves as heat. A car delivers about 1% of its fuel to its actual task. Power plants dump the majority of their fuel through cooling towers by design, while buildings downstream spend electricity on cooling. Data centres, at 415 TWh and doubling by 2030, spend up to a third of their power removing the heat their own computation makes, because their binary hardware cannot rest. The food system spends 8–10 calories to deliver 1, then discards a third of the result. And the announced remedy is a parallel chain with the same shape, new materials, and its own conversion losses at every link.

Each figure is documented, mainstream, and undisputed in its own literature. No engineer chose the two-thirds loss. It follows from the shape of the chain: everything passes through heat, nothing may rest, and every loss is answered with a new link that has losses of its own.

That is what is there.

Annotated References

Lawrence Livermore National Laboratory, "Energy Flow Charts" (annual; 2024 edition). Why read? The single most informative diagram in energy analysis. One page shows the entire US energy system: 94.2 quads in, 61.5 quads rejected. Study the diagram before reading anything else about energy; most debates dissolve in front of it.

International Energy Agency (2025), "Energy and AI." Why read? The authoritative source for data-centre electricity: 415 TWh in 2024, 945 TWh projected for 2030, growth four times faster than all other demand, and the US data-centre sector overtaking heavy industry. Also documents the cooling shares per data-centre type.

Energy Institute, "Statistical Review of World Energy" (current edition). Why read? The standard reference for global primary energy (~620 EJ) and its composition. The long time series show how stable the structure of the system is beneath the yearly headlines.

US Department of Energy, "Where the Energy Goes: Gasoline Vehicles" (fueleconomy.gov). Why read? The official breakdown of what happens to a tank of petrol: about 75–80% lost before the wheels. Combined with vehicle-to-passenger mass, it yields the ~1% service figure. Short and unambiguous.

Smil, V. (2017), Energy and Civilization: A History. Why read? The wide-angle view of energy chains across history, including the reversal of the food-energy ratio under industrial agriculture and the energy cost of Haber-Bosch nitrogen. Smil is the standard author for energy realism; read the chapters on modern agriculture and fossil society.

FAO (2011), "Global Food Losses and Food Waste." Why read? The source of the one-third figure: roughly 1.3 billion tonnes of food lost or wasted per year, documented per region and per link of the food chain.

Pimentel, D. & Pimentel, M. (2008), Food, Energy, and Society. Why read? The classic accounting of calories in versus calories out in industrial food systems — the origin of the 8–10:1 input ratio. Older, but its method is transparent and its arithmetic has not been overturned.

Mudd, G. M. (2010), "The Environmental Sustainability of Mining in Australia," Resources Policy. Why read? Documents the long-term decline of ore grades, including copper toward 0.5%, and the rising energy and rock movement per tonne of metal. The material mirror of the energy chain.

Uptime Institute, "Global Data Center Survey" (annual). Why read? The industry's own measurement of PUE — the ratio that quantifies how much energy computing spends removing its own heat. Useful because it is self-reported by the sector, not by critics.

IRENA, "Renewable Power Generation Costs" (annual series). Why read? The documented 90% cost decline of solar electricity — the strongest counter-series in the energy record. Read it to see that the panel is cheap; the chain around it is where the inventory of section 6 lives.

IEA (2019), "The Future of Hydrogen." Why read? The transition's own numbers on the hydrogen round trip: electrolysis at ~70%, reconversion at ~50%. Read it to see that the second chain's losses are stated openly in its planning documents, not discovered by opponents.

Chapter 3 — Computing: The Stack That Cannot Rest

Part I of What Is There *— J. Konstapel, Leiden, 2026*

This chapter records the current state of computing as a system: the hardware beneath it, the software on top of it, and what it costs to keep both running. It contains no theory and no proposal. Every figure can be checked against the sources listed at the end.

1. The register

All mainstream computing runs on one register: the bit, with two states, on and off. Every design decision downstream follows from it.

A binary circuit holds no state without effort. The most common working memory, DRAM, stores each bit as charge in a capacitor. The charge leaks. Every cell must therefore be rewritten — refreshed — roughly every 64 milliseconds, whether the data is used or not. A memory module refreshes itself about fifteen times per second, continuously, for its entire operating life, purely to remain what it is.

The processor is paced by a clock: a signal that switches billions of times per second and is distributed to every corner of the chip. Clock distribution and the circuitry it drives consume a substantial fraction of a processor's power — historically up to 30–40% — before any useful computation is counted.

Idle is not free either. A server doing nothing draws a large share of its full-load power. Industry measurements have long shown idle servers consuming 30–60% of their peak draw. Whole fleets of "comatose" servers — powered, cooled, and doing no identifiable work — have been documented in enterprise data centres at rates around 25–30% of installed machines.

The consequence was recorded in Chapter 2: computation converts nearly all of its electricity into heat, and additional electricity is then spent removing that heat. The root sits here, in the register. A system whose baseline state must be actively maintained has a floor under its consumption that no efficiency programme can remove.

2. The distance from the physical limit

Physics sets a floor on the energy cost of computation. The Landauer limit states that erasing one bit at room temperature costs at minimum about 3×10^{-21} joules.

Actual hardware operates nowhere near it. A logic operation in current CMOS silicon costs on the order of 10^{-15} to 10^{-14} joules once the surrounding machinery — wires, clocks, memory movement — is included. The gap between what computation must cost and what it does cost spans roughly five to seven orders of magnitude: a factor of a hundred thousand to ten million.

Most of that gap is not in the transistor. It is in movement. Moving data between memory and processor costs far more energy than operating on it; in modern systems, data movement dominates the energy budget of computation. The architecture that separates memory from processing — the von Neumann design, standard since 1945 — makes this movement mandatory for every operation. The industry name for the resulting congestion is the von Neumann bottleneck. It has been named, studied and lamented for half a century, and it is still the design of essentially every computer sold.

3. The hardware chain

The machinery that produces chips has become one of the most complex and expensive industrial chains ever built.

A leading-edge semiconductor fabrication plant costs about 20 billion dollars. A finished advanced chip passes through well over a thousand process steps across several months. The lithography machines at the centre of the process, built by a single company, ASML, cost around 200 million euros each; the newest High-NA generation approaches 400 million euros per machine, ships in multiple aircraft or dozens of trucks, and contains hundreds of thousands of parts.

The chain behind one chip spans the planet: silicon refined in one country, wafers grown in another, lithography from the Netherlands, chemistry from Japan, fabrication in Taiwan, packaging in Malaysia, design in California. It is the long chain of Chapter 2, rebuilt at nanometre scale, with the same property: every link adds cost, energy, and a point of failure. The world discovered the failure points in 2021, when a shortage of chips idled car factories on three continents for over a year.

Meanwhile the historical engine of the industry has slowed. For decades, transistor counts doubled about every two years and every doubling came with cheaper, cooler transistors. The cost-per-transistor decline has flattened. Power density has forced designers to leave growing fractions of the chip switched off at any moment — the industry term is dark silicon. The response has been to specialize: graphics processors, tensor processors, dedicated accelerators — more different kinds of chips, each with its own software, rather than one kind that improves.

4. The software stack

On top of the hardware sits a software stack that has grown by three to four orders of magnitude within one working lifetime, without a corresponding growth in what individual programs do for their users.

The numbers. The software that guided the Apollo lunar lander held in the order of a hundred thousand lines of code. Windows contains around 50 million. A modern premium car carries roughly 100 million lines — ten times more than a decade earlier, and more than the operating systems it connects to. Google’s company-wide repository passed 2 billion lines years ago.

The web shows the same curve at document scale. An average web page in the mid-1990s weighed a few tens of kilobytes. The average page now weighs over 2 megabytes — a growth of two orders of magnitude — while delivering, in most cases, text, images and forms, the same as before. The additional weight is machinery: frameworks, trackers, advertising infrastructure, compatibility layers.

Programs are no longer written so much as assembled. A routine JavaScript application pulls in hundreds to thousands of external packages, each with its own dependencies, few of them read by anyone who uses them. In 2016 the removal of one eleven-line package, left-pad, broke build systems across the industry for a day. In 2021 a flaw in one ubiquitous logging library, Log4j, created an emergency in effectively every large organization on earth at once, because the same unexamined component sat inside all of them.

The layers bought something real. Abstraction is why millions of people can build software without being chip architects; the trade was machine efficiency for human accessibility, and it was accepted one layer at a time.

The pattern was named decades ago. Wirth’s law, from 1995: software gets slower faster than hardware gets faster. Every gain of Moore’s law was absorbed by the growth of the stack above it.

5. The maintenance economy

Most of computing’s labour is not construction. It is upkeep.

Industry surveys have consistently found that organizations spend the majority of their IT budgets — commonly estimated between half and three quarters — on maintaining and operating existing systems rather than building anything new.

The Consortium for Information and Software Quality put numbers on the American case. Its 2022 report estimates the annual cost of poor software quality in the United States at 2.41 trillion dollars — about a tenth of GDP — from operational failures, vulnerabilities and failed projects. Separately, it estimates the accumulated technical debt — the backlog of rework that existing systems are known to need — at roughly 1.52 trillion dollars, about equal to the total annual wage bill of the entire US IT workforce.

Underneath the modern surface, the old layers keep running. An estimated 200–800 billion lines of COBOL, a language from 1959, remain in production, carrying banking, insurance and government systems. They are not maintained because they are good. They are maintained because nobody can safely replace them: the systems around them have grown so entangled that replacement projects routinely fail. Several governments discovered during the 2020 pandemic that emergency payments were delayed because the sixty-year-old systems beneath them could not be modified in time.

6. Error correction everywhere

Because no layer of the stack is reliable on its own, every layer checks the one below it.

Memory ships with error-correcting codes. Storage adds checksums. Networks retransmit corrupted packets. File systems verify their own structures. Databases log every change twice so that failures can be replayed. Applications validate their inputs, then validate them again at the next layer, because no layer trusts the previous one. Each check costs computation, storage and energy; together they form a permanent tax on every operation, paid to compensate for the fragility of the register of section 1.

The frontier of computing extends the pattern rather than escaping it. In current quantum computing designs, one reliable logical qubit is projected to require on the order of a thousand physical qubits, the surplus devoted entirely to detecting and correcting the errors of the rest. The most advanced machine on the horizon would consist, by count, of roughly 99.9% error-correction apparatus and 0.1% computation.

7. The record

The inventory of this chapter, in one place.

A register that cannot rest: memory refreshed fifteen times a second, clocks switching billions of times a second, idle servers drawing up to half their peak power. An energy cost per operation five to seven orders of magnitude above the physical floor, spent mostly on moving data across a bottleneck named and unresolved since 1945. A hardware chain of thousand-step processes, 20-billion-dollar factories and 400-million-euro machines, whose founding engine — cheaper transistors every two years — has stalled. A software stack grown ten-thousand-fold, assembled from unread components, in which one eleven-line package or one logging library can halt the world. A maintenance economy of 2.4 trillion dollars per year in one country, running on 1959-vintage code that nobody can remove. And at every layer, machinery that exists only to check the layer below.

Each figure is documented, mainstream, and undisputed in its own literature. No architect chose this stack. Each layer was a reasonable answer to the failures of the layer beneath it. The result is a structure that spends most of its energy remaining what it is, most of its labour repairing what it was, and a growing share of both checking whether either succeeded.

That is what is there.

Annotated References

Landauer, R. (1961), “Irreversibility and Heat Generation in the Computing Process,” IBM Journal of Research and Development. Why read? The paper that established the physical floor under computation: erasing one bit costs kT ln 2. Four pages. Everything in section 2 is measured against it.

JEDEC, DDR SDRAM standards (current editions). Why read? The industry’s own specification of the DRAM refresh requirement — the 64-millisecond cycle that makes binary memory a continuous consumer. Dry, but it is the primary source for the claim that the register cannot rest.

Horowitz, M. (2014), “Computing’s Energy Problem (and what we can do about it),” IEEE ISSCC. Why read? The standard reference on where computation’s energy actually goes: data movement dominates, arithmetic is nearly free. The clearest quantitative account of the von Neumann bottleneck as an energy problem.

Koomey, J. & Taylor, J. (2015), “New data supports finding that 30 percent of servers are ‘comatose’,” Anthesis Group. Why read? The measurement behind the idle-server figures: roughly 30% of surveyed enterprise servers powered and cooled while doing no identifiable work.

Esmaeilzadeh, H. et al. (2011), “Dark Silicon and the End of Multicore Scaling,” ISCA. Why read? The paper that named dark silicon and quantified why growing fractions of every chip must stay switched off. Marks the end of the era in which hardware improvement was automatic.

CISQ / Krasner, H. (2022), “The Cost of Poor Software Quality in the US: A 2022 Report.” Why read? The source of the 2.41-trillion-dollar annual figure and the 1.52-trillion-dollar technical-debt estimate. Read the methodology section; the two numbers measure different things and should not be added.

Wirth, N. (1995), "A Plea for Lean Software," IEEE Computer. Why read? The origin of Wirth's law, written by the designer of Pascal thirty years ago. Every complaint in it has since become measurably worse, which is precisely the point.

HTTP Archive, "State of the Web" / Page Weight reports (ongoing). Why read? The continuous public measurement of web-page weight — the two-orders-of-magnitude growth of section 4, tracked month by month across millions of pages.

Baldwin, C. & Clark, K. (2000), Design Rules: The Power of Modularity. Why read? The standard account of why systems are assembled from modules and what that assembly costs and hides. Background for the dependency economy of npm, left-pad and Log4j.

Fowler, T. (2020), reporting on state unemployment systems and COBOL (various outlets); and Reuters coverage of the 2021 chip shortage. Why read? Two documented demonstrations of the stack's rigidity: sixty-year-old code that could not be changed during an emergency, and car factories idled worldwide by the absence of chips. The abstract fragility of sections 3 and 5, observed in public.

Chapter 4 — Economy and Money: The Layer Above the Work

Part I of What Is There *— J. Konstapel, Leiden, 2026*

This chapter records the current state of the monetary and financial system: how large it is, what it consists of, what it costs, and how often it fails. It contains no theory and no proposal. Every figure can be checked against the sources listed at the end.

1. The debt

Global debt reached a record 348 trillion dollars at the end of 2025. Nearly 29 trillion was added in that single year — the fastest build-up since the pandemic, and an annual increase larger than the entire GDP of the United States. The total stands at roughly 308% of world GDP: for every euro of yearly production, three euros of outstanding obligation.

The composition: governments owe about 107 trillion, non-financial companies about 101 trillion, households about 65 trillion, and the financial sector the remainder. Governments drove the latest rise, with the United States, China and the euro area responsible for three quarters of it.

The direction is one-way. In 2016 global debt stood at 215 trillion. In 2018, 247 trillion. In 2023, 307 trillion. In 2025, 348 trillion. Debt has grown faster than the world economy in almost every year on record, which means the ratio between the obligation and the production that must service it keeps rising. The refinancing calendar shows what this implies: in 2026 alone, emerging markets must roll over more than 9 trillion dollars of maturing debt and mature markets more than 20 trillion — not new investment, but old borrowing replaced by new borrowing to remain outstanding.

2. What money is

The nature of the money itself is documented by its issuers.

In 2014 the Bank of England published "Money creation in the modern economy," a plain-language description of how money comes into existence. Its finding: the majority of money in circulation is created by commercial banks at the moment they make loans. A bank does not lend out deposits it has collected. It creates the deposit by lending. In modern economies, roughly 97% of the money supply consists of such bank deposits; physical cash is the small remainder.

Two properties follow directly, and both are stated in the central banks' own literature. First, virtually all money is somebody's debt; the money supply and the debt of section 1 are two views of the same ledger. Second, money is extinguished when loans are repaid — so an economy that collectively paid down its debt would shrink its own money supply. The system therefore requires the debt to grow for the money to grow.

3. The trading layer and the real layer

Above the production of goods and services sits a trading layer whose volumes have detached from what they nominally serve.

Foreign exchange. The Bank for International Settlements measures global currency trading every three years. The latest figure: about 7.5 trillion dollars per day. World trade in goods and services totals roughly 32 trillion per year — about 90 billion per day. Currency trading thus exceeds the trade it nominally facilitates by a factor of roughly eighty to one hundred. More currency changes hands every five days than the world exports in a year.

Derivatives. The BIS records the outstanding notional value of over-the-counter derivatives — contracts whose value derives from other assets — at around 700 trillion dollars. That is roughly seven times world GDP: the contracts written on top of the economy are several times larger than the economy.

Speed. In the major stock markets, around half of all trading volume is executed by high-frequency algorithms. Typical holding periods run from seconds down to microseconds. The exchanges sell co-location — server space physically adjacent to the matching engine — because at these speeds the length of a cable is a competitive advantage. Firms have laid dedicated fibre and microwave links between Chicago and New York to save fractions of a millisecond. This is the infrastructure of section 3 of Chapter 2 — energy, silicon, transmission — deployed so that a position may exist for less time than a human eye-blink.

4. The cost of the layer

The financial sector describes itself as an intermediary: it moves savings to investment and charges for the service. The charge has been measured.

Thomas Philippon reconstructed the unit cost of financial intermediation in the United States — what society pays the financial sector per dollar of intermediated asset — across 130 years. The result: about 2%, essentially unchanged from the 1880s to today. The telegraph, the telephone, the computer, and the entire digital revolution of Chapter 3 did not lower the toll. The sector absorbed every efficiency gain as income.

The income is visible in the aggregate statistics. The financial industry's share of US corporate profits, around 10–15% for most of the twentieth century, rose to roughly 40% just before 2008 and has remained far above its historical level since. Meanwhile a growing share of corporate cash flow returns to the financial layer rather than to production: S&P 500 companies now spend on the order of 900 billion to 1 trillion dollars per year buying back their own shares — in many years more than they invest in research and development.

5. The crises

The system's failures are periodic and documented.

Financial historians count, depending on definition, several hundred banking crises, currency crises and sovereign defaults over the past two centuries. The frequency did not decline as the system grew more sophisticated; the era of deregulated, globalized finance after 1980 produced crises at a higher rate than the tightly regulated decades before it — the Latin American debt crisis, the savings-and-loan collapse, the Asian crisis, the dot-com crash, and 2008.

The cost of 2008 alone: US government analyses put the output lost by the American economy in the range of 10 to 20 trillion dollars. Millions of homes foreclosed, roughly nine million US jobs lost, and equivalent damage across Europe, where the crisis ran on into a sovereign-debt crisis that lasted years.

The response is itself part of the inventory. To stabilize the system, central banks bought assets on a scale without precedent. The US Federal Reserve's balance sheet grew from under 1 trillion dollars in 2007 to a peak near 9 trillion. The European Central Bank's reached nearly 9 trillion euros. The Bank of Japan's exceeded the entire GDP of Japan. The rescue instrument was the creation of new money against new debt — the mechanism of section 2, applied at maximum scale. The debt record of section 1 is, in part, the accumulated residue of the rescues.

6. The rules about the rules

Each crisis added a regulatory layer, and the layers are countable.

The first Basel accord on bank capital, from 1988, ran to about 30 pages. Basel II, from 2004, ran to 347 pages. Basel III and its amendments run to thousands of pages, supplemented by national implementations that multiply the count again. The American Dodd-Frank Act of 2010 was 848 pages of law that instructed regulators to write roughly 400 separate rules; the implementing regulations run to tens of thousands of pages.

The banks answered with staff. Large international banks employ compliance departments of tens of thousands of people; several of the largest have publicly reported compliance and risk headcounts exceeding 20,000 — more people checking the bank than many banks employ in total. Annual compliance spending across the financial industry is estimated in the hundreds of billions of dollars.

The structure mirrors Chapter 3 exactly: each layer checks the layer below, each failure justifies a new layer, and the checking itself becomes one of the system's largest activities.

7. The record

The inventory of this chapter, in one place.

Obligations of 348 trillion dollars, three times world output, growing faster than the economy in almost every recorded year, with 29 trillion of refinancing due in 2026 merely to stand still. A money supply that is 97% bank-created debt and can only grow if the debt grows. A trading layer that turns over eighty to one hundred times the underlying trade, carries derivative contracts worth seven times world GDP, and holds positions for microseconds. An intermediation toll of 2%, unmoved by 130 years of technology. Crises arriving faster in the sophisticated era than before it, the largest of which cost one country 10–20 trillion dollars and was resolved by creating trillions more of the same instrument that failed. And a rulebook grown from 30 pages to tens of thousands, enforced by compliance departments larger than banks.

Each figure is documented, mainstream, and undisputed in its own literature — most of them published by the system's own institutions: the IIF, the BIS, the central banks. No one designed a system in which the claims on the work would grow permanently faster than the work. It grew, one defensible instrument at a time, and each repair added a layer to what it repaired.

That is what is there.

Annotated References

Institute of International Finance, "Global Debt Monitor" (February 2026 edition). Why read? The standard quarterly accounting of world debt: the 348-trillion total, the 308% ratio, the sector breakdown and the 2026 refinancing calendar. Published by the banking industry's own association, which is precisely why the figures are hard to dispute.

McLeay, M., Radia, A. & Thomas, R. (2014), "Money creation in the modern economy," Bank of England Quarterly Bulletin. Why read? The central bank's own correction of the textbook story: banks create money by lending; they do not lend out savings. Fourteen readable pages. The single most clarifying document about money in print.

Bank for International Settlements, "Triennial Central Bank Survey of Foreign Exchange and OTC Derivatives Markets" (latest edition). Why read? The source of the 7.5-trillion-per-day currency figure and the OTC derivatives notionals. The BIS is the bank of the central banks; these are the system's own measurements of its own volume.

Philippon, T. (2015), "Has the US Finance Industry Become Less Efficient?," American Economic Review. Why read? The 130-year measurement of the intermediation toll: about 2%, flat through every technological revolution. Short, careful, and never refuted — only ignored.

Reinhart, C. & Rogoff, K. (2009), This Time Is Different: Eight Centuries of Financial Folly. Why read? The standard historical database of financial crises across 66 countries and 800 years. Read it for the frequency tables; the pattern of section 5 is theirs.

US Government Accountability Office (2013), "Financial Crisis Losses and Potential Impacts of the Dodd-Frank Act." Why read? The official attempt to price the 2008 crisis for the United States: output losses estimated in the tens of trillions. A government document, not an activist one.

Federal Reserve, ECB and Bank of Japan, balance sheet statistics (public data series). Why read? The primary record of the rescue: central bank balance sheets growing roughly tenfold after 2008. The raw series make the point without commentary.

Basel Committee on Banking Supervision, Basel I (1988), Basel II (2004), Basel III framework (2010–2023). Why read? Read none of them in full; compare their page counts. The growth from 30 pages to thousands is itself the finding of section 6.

Lewis, M. (2014), Flash Boys: A Wall Street Revolt. Why read? The accessible account of the high-frequency trading infrastructure: the dedicated fibre line built to save three milliseconds between Chicago and New Jersey. Journalism, but the physical facts it reports are undisputed.

Lazonick, W. (2014), "Profits Without Prosperity," Harvard Business Review. Why read? The documentation of the buyback economy: corporate cash returned to the financial layer at a scale rivalling investment in production. The origin of the figures in section 4.

Chapter 5 — Institutions: Control on Control

Part I of What Is There — *J. Konstapel, Leiden, 2026*

This chapter records the current state of the large institutions: healthcare, education, government, and the organized workplace. It measures one thing throughout — the ratio between the work and the apparatus around the work. It contains no theory and no proposal. Every figure can be checked against the sources listed at the end.

1. Healthcare: the size

Healthcare is the largest institution of the modern state. The United States spends about 17–18% of GDP on it; Western European countries 10–12%. In the Netherlands, roughly one in six or seven working people works in care. The government's own scientific council, the WRR, projected in 2021 that on current trends this becomes one in four by 2040 — a quarter of the entire workforce absorbed by one sector.

The spending buys less health per euro over time. The United States spends nearly twice the rich-country average per person and has lower life expectancy than almost all of its peers. Life expectancy in several Western countries has plateaued or declined in recent years — the first sustained reversals since the world wars.

2. Healthcare: the apparatus

Where the money and the hours go has been measured.

In the United States, administration consumes about 34% of total health spending, against about 17% in neighbouring Canada — same medicine, same machines, double the overhead. The difference alone exceeds 600 billion dollars per year: billing, coding, insurance verification, prior authorization, claims disputes.

The staffing curve tells the story in one line. Between 1970 and the 2010s, the number of physicians in the United States roughly doubled. The number of healthcare administrators grew by an estimated 3,200%. The apparatus grew more than an order of magnitude faster than the profession it administers.

The hours of the professionals themselves follow. A time-motion study in Annals of Internal Medicine found American physicians spending two hours on desk work and electronic records for every one hour with a patient — and additional record work at home after hours. The electronic health record, introduced to reduce paperwork, is the primary instrument of it.

The Netherlands runs a milder version of the same ratio. Surveys among Dutch care professionals consistently report 25–40% of working time spent on administration and registration, depending on the sector. A national deregulation programme — "(Ont)Regel de Zorg" — has existed since 2018 specifically to remove rules. After years of scrapping lists, the reported administrative burden had not materially declined; new registration requirements arrived as fast as old ones were removed.

3. Youth care: the market

Dutch youth care shows what happens when the control apparatus and a market meet.

Since the 2015 decentralization, spending on youth care roughly doubled, passing 6 billion euros per year. The number of children in some form of youth care rose to about one in seven — in some municipalities one in five. Waiting lists for the heaviest cases grew at the same time.

The provider side exploded. Thousands of youth-care providers now operate in the market, from large institutions to single-person practices; municipalities contract hundreds of them each, and each contract carries its own tendering, accounting and audit obligations. Investigations found private-equity-owned providers achieving double-digit margins on light care while complex care ran waiting lists — the market selecting the profitable cases, the apparatus unable to see it happening despite measuring everything.

More money, more providers, more registration, more children in the system, longer waits for the children who needed it most. Every number in this paragraph moved in the same decade, in the direction opposite to the system’s stated purpose.

4. Education: the same curve

Education repeats the pattern at every level.

In American universities, the number of administrators and professional staff has grown several times faster than the number of teachers for four decades. Between the mid-1970s and today, faculty grew modestly while administrative positions multiplied; at many universities, non-teaching professionals now outnumber the teaching staff. Over the same period, tuition roughly tripled in real terms, and American student debt passed 1.7 trillion dollars — a sum administered, serviced and securitized by the financial layer of Chapter 4.

In Dutch primary and secondary education, teachers report spending roughly a third of their time on non-teaching tasks: registration, pupil-tracking systems, reporting, meetings about the reporting. The sector simultaneously carries a structural teacher shortage. The country cannot find people to stand in front of classes and occupies a third of those it has with administration of the standing.

The measurable output moves accordingly. Dutch 15-year-olds’ reading performance in the international PISA measurements has declined steadily for two decades; by the latest rounds, a quarter to a third risk leaving school functionally illiterate — in one of the richest, most schooled countries on earth, spending more per pupil than ever.

5. Government: the rules and one case

The rulebook itself is countable. The US Code of Federal Regulations has grown to roughly 190,000 pages. Dutch ministries collectively spend well over a billion euros per year hiring external consultants — capacity bought to manage complexity the apparatus itself generates.

One Dutch case shows the full machinery end to end: the childcare-benefits affair. The tax authority’s fraud-control system flagged tens of thousands of parents — disproportionately with dual nationality — as fraudsters, on grounds later ruled unlawful. About 68,000 parents registered as victims. Families were driven into debt collection; over 2,000 children were removed from their homes in the period concerned. The scandal brought down a government in 2021.

Then the repair repeated the pattern. The recovery operation built its own apparatus: assessment frameworks, files per parent running to thousands of pages, appeals on the assessments, appeals on the appeals. Costs of the recovery rose toward and past 7 billion euros — several times the amount originally recovered from the parents — while, years in, a large share of victims still awaited final settlement. The control system caused the harm; the control system for repairing the harm reproduced it.

6. The workplace

The organized workplace closes the inventory, and the Dutch metal and electronics industry — 38,000 companies, 350,000 jobs, 18% of Dutch exports — has been measured precisely.

TNO’s 2026 report for the sector, drawing on the national work surveys, records the following series. The share of companies engaged in process innovation — improving how the work itself is organized — fell from 46% in 2010 to 27.5% in 2021. Training effort declined 13% across industry over fifteen years. Burnout complaints in the sector rose from 14.5% in 2010 to 18.3% in 2022. Reported automation of work did not rise between 2017 and 2022 despite acute labour shortage. Over the same period, the share of companies with improved financial results grew to over 40%, and the gap between what labour produces and what it is paid — the markdown on labour — has widened since 2007.

Read as one table: the organizations stopped improving the work, trained less, automated no more, burned out more of their people — and grew more profitable while doing so. TNO’s own conclusion was that adding people and adding technology will not raise productivity within five years; only reorganizing the work itself would. That is the one measure the surveys show declining.

7. The record

The inventory of this chapter, in one place.

A care sector heading toward employing one working person in four, spending a third of its money and up to 40% of its professionals' time on administration, with administrators multiplying thirty times faster than physicians. A youth-care system that doubled its budget and its registrations while the waiting lists for serious cases grew. Universities where the non-teaching staff outnumber the teachers and the students carry 1.7 trillion dollars of debt. Schools that occupy a third of scarce teachers' time with registration while reading ability falls for twenty years. A state with 190,000 pages of rules whose fraud-control system harmed 68,000 families and whose repair operation cost 7 billion euros and reproduced the harm. And an industrial sector that stopped improving its own work, measured itself doing so, and grew more profitable.

Each figure is documented, mainstream, and in most cases produced by the institutions themselves — the WRR, TNO, the ministries, the universities. Every control layer was added for a reason, in response to a real failure. Nobody chose a system in which the apparatus grows faster than the work in every sector at once. It grew, one defensible rule at a time.

That is what is there.

Annotated References

WRR (2021), "Kiezen voor houdbare zorg" (Choosing Sustainable Care). Why read? The Dutch government's own scientific council projecting care toward one in four workers by 2040. The report that establishes the size problem from inside the state, not from its critics.

Himmelstein, D., Campbell, T. & Woolhandler, S. (2020), "Health Care Administrative Costs in the United States and Canada," Annals of Internal Medicine. Why read? The cleanest natural experiment on administrative overhead: 34% versus 17%, same continent, same medicine. The 600-billion-dollar difference of section 2 is computed here.

Sinsky, C. et al. (2016), "Allocation of Physician Time in Ambulatory Practice," Annals of Internal Medicine. Why read? The time-motion study behind the two-hours-desk-per-hour-of-patient figure. Direct observation, not survey — the physicians were followed with stopwatches.

Woolhandler, S. & Himmelstein, D., analyses of US physician and administrator growth (based on Bureau of Labor Statistics series). Why read? The source of the 3,200% administrator-growth figure against a doubling of physicians. Check the underlying BLS occupational series; the curve survives every reasonable definition.

Ministerie van VWS, "(Ont)Regel de Zorg" programme reports (2018–present). Why read? The state's own deregulation programme documenting, year by year, that removing rules did not reduce the reported burden. An institution measuring its own inability to shrink itself.

Jeugdautoriteit / CBS, youth-care statistics; and follow-the-money investigations into private equity in Dutch youth care. Why read? The doubled spending, the one-in-seven figure, and the margin data on light versus complex care. Together they document a market and a control apparatus failing in opposite directions simultaneously.

Ginsberg, B. (2011), The Fall of the Faculty: The Rise of the All-Administrative University. Why read? The standard account of administrative growth in American universities, written by a professor from inside one. Polemical in tone, but its staffing numbers are drawn from federal data and have been repeatedly confirmed.

OECD, PISA results for the Netherlands (2003–2022 rounds). Why read? The twenty-year decline in Dutch reading performance, measured by the same international instrument throughout. The output side of section 4, in one downward series.

Parlementaire ondervragingscommissie Kinderopvangtoeslag (2020), "Ongekend onrecht" (Unprecedented Injustice); and subsequent Rekenkamer reports on the recovery operation. Why read? The official record of the benefits affair and — in the audit office's follow-ups — of the recovery operation that cost multiples of the original harm. The complete control-loop failure, documented by the

state about itself.

Dhondt, S. & Oeij, P. (2026), "Industriebeleid met aandacht voor kwaliteit van werk," TNO 2026 17508, commissioned by FNV Metaal. Why read? The measured series of section 6: process innovation halved, training down, burnout up, profits up. Thirty-six pages, national survey data, and the sector's own conclusion that only reorganizing the work will help — the one thing the data show declining.

Chapter 6 — The Human Being: The Measured Life

Part I of What Is There *— J. Konstapel, Leiden, 2026*

The previous chapters inventoried the systems. This chapter inventories the people inside them: what the systems' own statistics record about the state of the human beings for whom all of it nominally exists. It contains no theory and no proposal. Every figure can be checked against the sources listed at the end.

1. The exit from work

The Netherlands keeps precise records of who can no longer work. The records show the youngest leaving fastest.

WIA applications — the benefit for those still unable to work after two full years of illness — reached 93,000 in 2024, the highest number since the scheme existed. The fastest-growing group is people in their twenties. Between 2018 and 2025, annual WIA registrations by men aged twenty to thirty rose 84%; by women of the same age, 121%. The younger the applicant, the larger the share leaving with mental complaints — depression, anxiety, burnout. Nationally, WIA inflow for psychological causes rose 29% in three years.

Two properties make these numbers heavier than they look. First, the WIA counts only the end of a long funnel: everyone in it has already been ill for two years, so the visible group is a fraction of a much larger one still inside the funnel. Second, the exit is nearly permanent: of those entering with mental complaints, one in seven returns to any work within five years. A twenty-five-year-old entering the WIA has, statistically, left the workforce for good — with forty unworked years ahead.

The assessing institution cannot keep up. The backlog of pending WIA assessments is projected to grow toward 100,000–150,000 by 2030. The apparatus of Chapter 5 measures the outflow precisely and cannot process it.

2. The mind

The wider mental-health statistics frame the WIA numbers.

Roughly one in five Dutch employees reports burnout complaints — about 1.7 million working people, a share that has risen for over a decade. In the metal industry of Chapter 5, the series ran from 14.5% in 2010 to 18.3% in 2022. Professors in the field note that the prevalence of psychological problems has risen for fifteen years straight, and that no accepted scientific explanation for the rise exists.

Over a million Dutch people use antidepressants. In the United States, about 13% of adults do — among women over forty, roughly one in four. Prescriptions for ADHD medication have multiplied within a generation, in children and increasingly in adults.

The treatment system has queues. Dutch mental healthcare carries waiting lists on the order of 80,000–100,000 people; for the heaviest categories, waiting times exceed the legal norms structurally, year after year. The same pattern as youth care in Chapter 5: rising demand, rising spending, rising waits — all three at once.

3. The young

The steepest curves belong to the youngest, and they bend at a datable point.

Across the Western world, measures of adolescent anxiety, depression, self-harm and psychiatric admission run roughly flat through the 1990s and 2000s, then turn sharply upward in the early 2010s. In the United States, rates of major depression among teenagers rose on the order of 150% between 2010 and 2020; emergency-room admissions for self-harm among girls doubled and, for the youngest girls, nearly tripled. The Dutch student monitors report half of university students experiencing psychological complaints, a quarter to a third severe.

The timing coincides with one documented change in daily life: between roughly 2010 and 2015, the smartphone and social media became universal among the young. Teenagers now report seven to eight hours of recreational screen time per day. Whether the correlation is causal is genuinely contested in the literature: researchers such as Candice Odgers find the measured screen effects small and point to economic insecurity and inequality as the larger factors, while others, such as Twenge and Haidt, read the timing as causal. That the two curves — screen adoption and psychological distress — rose together, in every country measured, is contested by no one.

The same cohort is the one entering the WIA funnel of section 1 ten years later. The curves connect: the distressed fifteen-year-old of 2015 is the twenty-five-year-old leaving the workforce in 2025.

4. Alone

The statistics of human contact form their own series, and they all point the same way.

Time spent with friends, measured by the American time-use surveys: for young people, from about 150 minutes per day in the early 2000s to roughly 40 minutes by the early 2020s — a decline of 70%, most of it before the pandemic. Time alone rose by the equivalent of a full day per month.

In 2023 the US Surgeon General issued a formal advisory declaring an “epidemic of loneliness,” attaching a health cost: sustained loneliness raises mortality risk comparably to smoking up to fifteen cigarettes a day. The United Kingdom appointed a Minister for Loneliness in 2018 — a cabinet post for the absence of human contact.

The household statistics agree. Nearly 40% of Dutch households now consist of one person; in the large cities, half. Marriage rates, cohabitation rates and reported friendship counts have declined across the West; the share of Americans reporting no close friends at all quadrupled in three decades. Birth rates followed: the Netherlands stands near 1.4 children per woman, most of Europe and East Asia between 0.7 and 1.5 — far below replacement, the lowest levels ever recorded in peacetime. The standard demographic account attributes the long decline to prosperity, contraception and the education of women — documented drivers, and achievements in their own right. What the record adds is only the level now reached, and that it is still falling.

5. The deaths

At the far end of the series stand the deaths.

Economists Anne Case and Angus Deaton documented the rise of what they named deaths of despair — from drugs, alcohol and suicide — among working-age Americans without a college degree. The category now claims on the order of 150,000–200,000 American lives per year; drug overdoses alone passed 100,000 in a single year, more than American combat deaths in Vietnam, Korea and every war since, combined, per year.

The consequence is visible in the master statistic. US life expectancy stopped rising in the mid-2010s and fell for several years running — before the pandemic — the first sustained peacetime decline in a century. Several European countries have seen their own curves flatten. For the first time since the statistic exists, children in a number of rich countries are not expected to live longer than their parents.

6. The days

What the days themselves consist of has been measured, and can be stated without commentary.

The average adult in a Western country spends about seven hours per day looking at screens, roughly nine to ten hours sitting, and — outside work and household — under twenty minutes in face-to-face social contact. Commuting takes close to an hour. Time in nature, tracked in national surveys, averages minutes per day; a large share of children play outside less than an hour daily, a threshold below the exercise requirement for prisoners.

The physical surroundings changed in the same direction. In the Netherlands, villages and neighbourhoods lost their fixed meeting points one category at a time — documented in the retail and services statistics: bank branches down over 80% in two decades, post offices effectively abolished, village shops, cafés, churches and libraries closing by the hundreds per year. Each closure was individually rational: the transactions moved online, where they are cheaper. What disappeared with them was not the transaction but the encounter that came with it — the queue, the counter, the known face. The statistics of section 4 record the aggregate result.

None of this was chosen either. Every screen hour, every closed counter, every online transaction was individually convenient. The days rearranged themselves one convenience at a time.

7. The record

The inventory of this chapter, in one place.

A record 93,000 disability applications in one small country in one year, with twenty-somethings the fastest-growing group — up 84% and 121% in seven years — leaving with mental complaints, one in seven ever returning. A fifth of the workforce reporting burnout, a million-plus on antidepressants, treatment queues of 100,000. Adolescent distress curves bending sharply upward in the early 2010s in every measured country, in step with seven-to-eight-hour screen days. Time with friends down 70% among the young; a cabinet minister for loneliness; 40% single households; birth rates of 1.4 and falling. Deaths of despair by the hundred thousand and life expectancy declining in peacetime for the first time in a century. And days consisting, on average, of seven screen hours, ten sitting hours, and twenty minutes of human company.

Each figure is documented, mainstream, and produced by the systems' own statistical offices — the UWV, the CBS, the surgeons general, the time-use surveys. The systems of Chapters 1 through 5 measure the people inside them with great precision. The measurements say the people are doing worse on every axis that involves other people, meaning, or rest — and no accepted explanation exists in the literature of any single field.

That is what is there.

Annotated References

UWV, WIA statistics and trend analyses (2024–2026), with the FD analysis of the 2018–2025 series. Why read? The primary record of section 1: 93,000 applications, +84% and +121% among twenty-somethings, the 29% rise in psychological inflow, and the one-in-seven return rate. The institution's own numbers, broken down by the newspaper.

TNO/CBS, Nationale Enquête Arbeidsomstandigheden (annual). Why read? The long-running national measurement of burnout complaints — the source of the one-in-five figure and the sector series used in Chapter 5. The same instrument, year after year, which is what makes the trend hard to dismiss.

Case, A. & Deaton, A. (2020), Deaths of Despair and the Future of Capitalism. Why read? The documentation of section 5: mortality from drugs, alcohol and suicide rising for an entire class of working-age Americans, and the resulting fall in life expectancy. Careful, data-driven, and written by a Nobel laureate in economics.

Office of the US Surgeon General (2023), "Our Epidemic of Loneliness and Isolation." Why read? A state medical authority formally classifying the absence of human contact as a public-health epidemic, with the fifteen-cigarettes mortality comparison. The document's existence is itself a datum.

Twenge, J. (2023), Generations; and Haidt, J. (2024), The Anxious Generation. Why read? The assembled international evidence for the post-2010 bend in adolescent mental-health curves and its coincidence with smartphone adoption. Read them for the data series; the causal debate around them is part of the record too.

American Time Use Survey (BLS), analyses of social time 2003–2023. Why read? The source of the 70% decline in time with friends among the young and the rise of time alone. Time-use data measure behaviour, not opinion — people report what they did yesterday, hour by hour.

Trimbos Instituut / RIVM, Monitor Mentale gezondheid en Middelengebruik Studenten (recent rounds). Why read? The Dutch student figures: half with psychological complaints, a quarter to a third severe. The national measurement of the cohort now entering the workforce — and, per section 1, beginning to leave it.

CBS, household statistics and birth statistics (StatLine, current series). Why read? The 40% single-person households and the 1.4 fertility rate, from the national statistical office. Two series that quantify, more directly than any survey, how the society is physically arranged.

Nederlandse Zorgautoriteit, GGZ waiting-list reports (annual). Why read? The official record of the mental-healthcare queues and the structural exceedance of legal waiting norms. Chapter 5's institutional pattern, applied to the subject of this chapter.

Vriens, I. et al. / retail and services statistics on bank branches, post offices and village amenities (CBS, DNB, Locatus series). Why read? The countable disappearance of the daily meeting points: bank branches down over 80%, counters and shops closing by the hundreds per year. The physical substrate of the loneliness statistics, documented one closure at a time.

Chapter 7 — The Earth: The Ground Under the Chain

Part I of What Is There *— J. Konstapel, Leiden, 2026*

This chapter records the state of the physical ground on which everything in the previous chapters runs: the atmosphere, the water, the soil, the living world — and the answers currently on offer. It contains no theory and no proposal. Every figure can be checked against the sources listed at the end.

1. The heat

The atmospheric record is the simplest series in this book. Carbon dioxide stood at about 280 parts per million before industry, 315 when systematic measurement began in 1958, and passes 425 now. The curve has never bent; the annual increase is larger today than in any earlier decade.

The temperature followed. The last two years rank as the warmest ever measured, and the global average has touched the 1.5-degree line that the 2015 Paris agreement was written to avoid — reached not in 2100 but within ten years of the signing. Every year since 2015 stands among the warmest on record.

The warming is not distributed evenly. Central Hungary, in the middle of Europe, warms roughly 50% faster than the continental average. The poles warm three to four times faster than the globe. The averages in the headlines understate every place where people actually live through the extremes.

2. The water

The water statistics divide into two: the water that falls and the water that stays. The second is the one collapsing.

Groundwater is being withdrawn faster than it recharges across every major agricultural region: the North China Plain, northwest India, the American High Plains and Central Valley, Iran, the Mediterranean basins. The GRACE satellites, which weigh the water in the ground from orbit, show the depletion directly. The land answers by sinking: parts of California's Central Valley have subsided nine metres; Jakarta sinks so fast that Indonesia is building a replacement capital.

Europe now has its own case. On the Hungarian plain between the Danube and the Tisza, the groundwater near Kiskunmajsa fell from 6.5 to 8 metres depth in a single year. Scientists of the Hungarian Hydrological Society describe the region's agriculture as dying and the sand ridge as past its tipping point; the word in use is desertification — in the heart of Europe. The 2022 drought cut the Hungarian grain harvest by up to half, the largest loss in the EU, around 6 billion euros of damage. The Paks nuclear plant, supplying 40% of the country's electricity, had to reduce output because the Danube ran too low to cool it — the energy chain of Chapter 2 throttled by the water of this one.

The origin of the Hungarian case is documented engineering. In the regulation works begun in 1846, the Tisza was shortened by hundreds of kilometres and its floodplains — the sponge of the plain — were drained for agriculture. The water that once stayed for months now leaves for the Black Sea in days. The Netherlands runs a gentler version of the same design: a drainage system built to remove water as fast as possible, and a countryside now officially registered as structurally desiccating, with drought damage to nature and agriculture in repeated summers.

3. The soil and the living

The UN's food organization estimates that a third of the world's soils are degraded — eroded, compacted, salinized or stripped of organic matter — and that topsoil is being lost many times faster than it forms. The organic-matter content of intensively farmed soils has declined for decades; each lost percentage point reduces the water the soil can hold, which feeds back into section 2.

The living world above the soil is measured by the Living Planet Index: monitored wildlife populations have declined on average 73% since 1970 — within one human lifetime. The German Krefeld study found flying-insect biomass in protected reserves down about 75% in 27 years — in the protected areas. Bird populations followed the insects; European farmland birds are down by more than half.

One aggregate figure locates the whole picture. In 2020, researchers calculated that the total mass of human-made material — concrete, asphalt, metal, plastic — had come to exceed the mass of all living things on earth. The built chain now outweighs the biosphere it runs on, and it doubles roughly every twenty years.

4. The rivers, the sand, the ground

The planet's plumbing has been rebuilt. Only about a third of the world's long rivers still flow freely; the rest are fragmented by tens of thousands of large dams. The dams trap the sediment that once built the deltas, so the deltas — home to hundreds of millions — sink and shrink while the sea rises toward them.

Sand and gravel are the most extracted solid materials on earth: roughly 50 billion tonnes per year, mostly for the concrete of section 3's anthropogenic mass. Riverbeds and coastlines are mined at rates that lower rivers, undermine bridges and erase islands. Six of the nine planetary boundaries that earth-system scientists use to define a safe operating space — climate, biosphere integrity, land use, fresh water, nutrient flows, novel entities — are assessed as already crossed.

5. The answers on offer

Two counter-series belong in this record before the current remedies are counted. The ozone layer is recovering: the 1987 Montreal Protocol banned one class of substances, the ban was implemented worldwide, and the measured ozone hole is now on course to close within decades — the largest documented repair of a planetary system, achieved by one treaty and one removal. And Europe's forests have expanded for decades; forest area on the continent has grown by an area the size of a mid-sized country since 1990, largely by withdrawal of marginal farmland. Repair at planetary scale is in the record. Both entries did it by taking something out.

The remedies currently financed follow the opposite pattern: each adds a layer to the chain rather than shortening it.

Geo-engineering has moved from taboo to portfolio. Investors have put on the order of 120 million dollars into climate-modification ventures: a startup with some 60 million euros for stratospheric aerosols, another with 29 million for orbiting mirrors to sell sunlight, a British programme of about 60 million pounds, and a company selling “cooling credits” — one dollar per gram of sulphur dioxide released, claimed to offset a tonne of CO_2, without published evidence. The proposal, at scale, would require injecting millions of tonnes of sulphur into the stratosphere every year, indefinitely; stopping would release the suppressed warming within years. A group of several hundred scientists has called for an international non-use agreement — a treaty to prevent the remedy.

The energy transition, inventoried in Chapter 2, is the larger answer: a second extraction-conversion-storage chain built beside the first, drawing on the mines of section 4 and the falling ore grades of Chapter 2. Carbon markets add a financial layer in the style of Chapter 4: over a hundred billion dollars of traded certificates, with repeated investigations finding that a large share of offset projects — in some audits the great majority — delivered no measurable additional carbon reduction. The claims trade grew; the verified effect did not.

6. The Dutch case: the lock

The Netherlands closes the chapter with a demonstration of what happens when the control apparatus of Chapter 5 meets the earth of this one.

Dutch agriculture, the second-largest agricultural exporter in the world by value on a territory the size of a small American state, deposits more reactive nitrogen per hectare than almost any country on earth. The deposition degrades the protected nature areas the state itself designated. In 2019 the highest administrative court ruled the accounting system that had permitted this — a scheme that booked future reductions against present emissions — unlawful.

The consequence was immediate and structural: construction permits, road projects and farm expansions froze nationwide, because each required nitrogen accounting that could no longer be produced. Years of buyout schemes, emission models, court cases about the models, and a political crisis followed; a farmers' protest party briefly became the largest in the country. The state now spends tens of billions buying out the production capacity it previously subsidized into existence, at a pace the courts repeatedly rule insufficient. A country locked by the interaction of its own export chain, its own nature registry, and its own accounting rules — every component functioning as designed.

7. The record

The inventory of this chapter, in one place.

Carbon dioxide at 425 ppm and the 1.5-degree line touched a decade after it was set as the limit. Groundwater falling under every major food region, land subsiding metres, and desertification underway in central Europe, where one year moved the water table a metre and a half and drought halved a national harvest. A third of soils degraded, wild populations down 73%, insects down 75% in protected reserves, and the built mass now heavier than everything alive. Two thirds of long rivers dammed, 50 billion tonnes of sand a year, six of nine planetary boundaries crossed. Remedies consisting of sulphur in the stratosphere, mirrors in orbit, a parallel extraction chain, and offset markets that mostly offset nothing. And one country frozen solid by its own environmental accounting.

Each figure is documented, mainstream, and drawn from the monitoring systems the chain built to watch itself — satellites, statistical offices, court records. The pattern is the same as in every previous chapter, now written into the ground: the chain grew one defensible step at a time, the water and the living world absorbed the cost, and the proposed repairs are new links.

That is what is there. This closes the inventory.

Annotated References

NOAA Global Monitoring Laboratory, Mauna Loa CO2 record; and WMO, State of the Global Climate (annual). Why read? The two primary series of section 1: the unbroken CO2 curve since 1958 and the temperature record including the 1.5-degree exceedance. Raw measurement, no interpretation required.

Rodell, M. et al. / NASA GRACE and GRACE-FO mission results on groundwater storage. Why read? The satellites that weigh groundwater from orbit. The maps show every major depletion region in one view — the fastest way to grasp section 2 globally.

Kabel, D. (2026), "Een woestijn in het hart van Europa," Het Financieele Dagblad; with OECD drought-damage assessments for Hungary. Why read? The documented European desertification case: the metre-and-a-half groundwater drop at Kiskunmajsa, the 50%-faster warming, the halved harvest, the throttled nuclear plant. Field reporting anchored in named scientists and official damage figures.

FAO and ITPS (2015, with updates), "Status of the World's Soil Resources." Why read? The source of the one-third-degraded figure and the erosion-versus-formation rates. The UN's own soil inventory, written by soil scientists rather than campaigners.

WWF/ZSL (2024), Living Planet Report; and Hallmann, C. et al. (2017), "More than 75 percent decline in flying insect biomass," PLoS ONE. Why read? The 73% wildlife decline and the Krefeld insect measurement — the latter notable because it was taken inside protected reserves, removing the usual habitat-loss explanation.

Elhacham, E. et al. (2020), "Global human-made mass exceeds all living biomass," Nature. Why read? The single calculation that locates the whole chapter: the built chain now outweighs the biosphere. Four pages; the crossover year is approximately 2020.

Grill, G. et al. (2019), "Mapping the world's free-flowing rivers," Nature; and UNEP (2022), "Sand and Sustainability." Why read? The two-thirds fragmentation figure and the 50-billion-tonne sand economy. Together they document the rebuilt plumbing and the mined ground of section 4.

Richardson, K. et al. (2023), "Earth beyond six of nine planetary boundaries," Science Advances. Why read? The current formal assessment of the safe operating space: six boundaries crossed. The framework paper for reading every other figure in this chapter as one system.

Selderbeek, E. (2026), “Investeerders steken geld in manipulatie van weer en klimaat,” Het Financieele Dagblad; and Biermann, F. et al., “Solar geoengineering: The case for an international non-use agreement,” WIREs Climate Change. Why read? The investor inventory of section 5 — the 120 million dollars, the cooling credits sold without evidence — and the scientists’ formal call for a treaty against the remedy. Read them together; the pair is the finding.

Raad van State (2019), PAS ruling; and Rekenkamer/planbureau assessments of the Dutch nitrogen approach. Why read? The court decision that froze a country and the audit trail of the buyout programmes since. The complete documentation of section 6: a lock built entirely from functioning components.

Appendix A — What Already Works

Appendix to What Is There *— J. Konstapel, Leiden, 2026*

The seven chapters recorded what is there. This appendix records something else that is there: working answers. Every entry below exists, runs today or ran for decades, and has been measured. None is a proposal. The selection rule is strict: an entry qualifies only if it demonstrably delivers the same service as the standard system with fewer links, less apparatus, or less permanent activity. The entries are ordered by the chapters they answer.

A.1 Science (Chapter 1)

Registered Reports. A journal format in which the study design is peer-reviewed and accepted *before* the results exist. Publication no longer depends on the outcome, so the incentive to produce positive results disappears. The measured effect: in conventional journals, 5–20% of published findings are null results; in Registered Reports, around 44–60% — close to what honest science statistically produces. Over 300 journals now offer the format. One change in the order of steps removes the mechanism behind a large share of the replication failure of Chapter 1.

Preprints. Physics and mathematics have run on arXiv since 1991: results published directly, free, within days, read by the whole field, with the journal version following later or never. The multi-billion-euro tollgate of Chapter 1, section 6 is bypassed by a server that costs a few million per year. Biology and medicine followed with bioRxiv and medRxiv; during the pandemic, the fastest reliable channel of scientific communication on earth was the free one.

A.2 Energy (Chapter 2)

The heat pump. Combustion converts fuel to heat at best at 100%. A heat pump moves heat instead of making it and delivers 300–500% — three to five units of warmth per unit of electricity. It is the LED logic of Chapter 2 applied to heating: abandon the flame as intermediate step. It is also the rare case where the short chain is already winning commercially, tens of millions of units per year.

The passive house. Roughly 40,000 certified buildings, three decades of measurement: insulation, orientation and heat recovery reduce heating demand by about 90% against conventional stock. Not by adding a generation-storage-backup chain, but by removing the need the chain was meant to serve. The measured extra construction cost runs 5–10%; the energy link largely disappears.

Regenerative water farming, measured in one field. On the drying Hungarian plain of Chapter 7, the regenerative farmer Csaba Szabó at Ferencszállás lost about 15% of his harvest in the drought year that cost his conventional neighbours over 50%. The difference was organic matter: each percentage point holds roughly 8–17 millimetres of extra rain where it fell. Beside him, the volunteer Waterhoeders block drainage canals to keep winter water on the meadows — reversing, ditch by ditch, the 1846 engineering that drains the plain. Both exist, both are measured, neither requires a single new link.

A.3 Computing (Chapter 3)

Setun. In 1958 Moscow State University built a computer on balanced ternary — the three states −1, 0, +1, with a genuine neutral. About fifty were produced and ran reliably for years; the engineers reported it simpler and cheaper than its binary contemporaries. It was discontinued by administrative decision, not by technical failure. It stands as the documented proof that the binary register of Chapter 3 was a choice, not a necessity.

Event-driven and in-memory hardware. Neuromorphic chips such as Intel's Loihi abandon the two structures that dominate Chapter 3's energy bill: the global clock and the memory-processor separation. Circuits act only when a signal arrives — no clock, no refresh, rest as the default state — and computation happens where the data sits. Measured gains on suitable workloads run to two, three, four orders of magnitude less energy. The brain, running the same principles, does its work on about 20 watts.

Sparse distributed memory. Pentti Kanerva's architecture from 1988 stores information by address in a vast sparse space rather than in numbered slots: similar things land near each other, retrieval works from partial or damaged cues, and errors degrade the answer gracefully instead of destroying it. It is a worked-out, published alternative to the addressing scheme every current computer uses — and the design basis on which the MAZE engineering line builds.

A.4 Money (Chapter 4)

WIR. In 1934, in the deflation crisis, Swiss businesses founded a mutual-credit circle: members pay each other in a unit created at the moment of the transaction and extinguished when balanced — money without a bank creating it as interest-bearing debt. WIR still operates, with tens of thousands of Swiss SMEs; nine decades of data show its turnover rising in recessions and falling in booms, cushioning the cycle that Chapter 4 documents. Sardinia's Sardex ran the same design from 2010 and carried thousands of firms through the euro crisis.

Agent-centric ledgers. Holochain implements accounting without global consensus: each participant keeps a signed chain of their own transactions, validated by peers, with no world-state that everyone must compute and store. No mining, no proof-of-work energy bill, no 700-trillion abstraction layer — the ledger scales with the participants instead of with the trading volume of Chapter 4, section 3. It exists, runs, and is the stated reference point for the MAZE-Bank design line.

A.5 Institutions (Chapter 5)

Buurtzorg. In 2006 a Dutch nurse founded a home-care organization on one rule: self-managing teams of at most twelve nurses, a full back office of under 1% overhead, no managers. It grew to some 15,000 employees and became the largest home-care provider in the country. The measurements, repeated by accountants and ministries: overhead around 8% against 25% sector average, roughly 40% fewer care hours per client for equal or better outcomes, the highest client satisfaction in the sector, and repeatedly the highest employee satisfaction of any large Dutch employer. The apparatus of Chapter 5, removed, and the work improved by its removal — measured for two decades.

The citizens' assembly. Ireland handed its most deadlocked questions — marriage, abortion — to assemblies of citizens drawn by lot, briefed by experts, deliberating over months. Their recommendations went to referendum and passed. France, Belgium's German-speaking community (with a permanent citizens' council since 2019) and dozens of cities followed. The mechanism replaces the party-media-lobby chain with the shortest possible one: the population, sampled, informed, asked. It exists in statute and has settled questions that thirty years of the standard apparatus could not.

Agora. A publicly funded Dutch secondary school in Roermond with no classes, no subjects and no timetable: each pupil pursues self-chosen challenges with a coach, meeting the national exam requirements by their own route. Inspected, accredited, expanding to dozens of schools. It removes the registration-and-timetable apparatus of Chapter 5, section 4 and is permitted to exist inside the standard system — which makes it a measured demonstration rather than a thought experiment.

A.6 The human being (Chapter 6)

The mutual-aid group. The oldest running answer to the loneliness and distress of Chapter 6 is the peer circle: Alcoholics Anonymous, founded 1935, now some two million members in 180 countries, cost per member near zero, no professional apparatus. A 2020 Cochrane review — the strictest instrument evidence-based medicine has — found AA and twelve-step facilitation performing as well as or better than professional psychotherapies for abstinence, at a fraction of the cost. The format has been copied for grief, debt, illness and bereavement. It is the documented case of care that consists entirely of the encounter the statistics of Chapter 6 show disappearing.

Positive Health. Machteld Huber's Dutch framework redefines health as the ability to adapt and self-manage across six life dimensions, replacing the defect-repair model. In the Limburg region where general practitioners adopted it, consultations changed character, referrals to second-line care dropped measurably, and the approach spread through the national health agreements. A working, measured reduction of the funnel into the apparatus of Chapter 5.

Life-pattern instruments. In the engineering line of this book's author, several instruments run today: SWARP derives a person's profile and next step from birth data and supports matching between people; a living-environment application watches public data sources and alerts residents; MAZE reads market tension as a state variable. They are listed here on the same terms as every other entry — as running engineering, measured by whether their output is useful, claiming nothing further in this appendix.

A.7 The earth (Chapter 7)

The johads of Rajasthan. From 1985, Rajendra Singh and village communities built thousands of small earthen check dams — johads — across dry Rajasthan. The measured result over three decades: groundwater rose metres, several rivers that had been dry for decades flow again year-round, and villages that had emptied repopulated. Total apparatus: earth, local labour, and the rule that the water stays where it falls. It is the Hungarian Waterhoeders' method, thirty years further along, at the scale of a region.

The Loess Plateau. From 1994, China restored some 35,000 square kilometres of eroded plateau — terracing, planting, closing slopes to grazing. Satellite-documented outcome: vegetation cover multiplied, sediment flowing into the Yellow River cut dramatically, local incomes doubled. One of the largest ecosystem restorations ever undertaken, and its core operation was subtraction: stop the draining, stop the grazing, let the sponge rebuild.

The Montreal Protocol. In 1987 the world's governments banned one class of industrial substances, the CFCs, to stop the destruction of the ozone layer. The ban was implemented globally, substitutes were found, and the measured ozone hole is now closing — full recovery is projected within decades. It is the largest repair of a planetary system ever carried out, and its entire mechanism was subtraction: no offsetting market, no compensating technology layer, one substance class removed. It stands in this list as the demonstration that the pattern of the other entries scales to the planet.

The beaver. Where beavers are reintroduced, their dams re-wet valleys, raise water tables, filter sediment and halve flood peaks — measured in British and American trials. The animal performs, for nothing, continuously, the retention engineering that water boards buy in concrete. Several countries now formally deploy reintroduction as water policy.

A.8 The common property

This appendix has stated twenty-odd entries without connecting them. One observation belongs in the record before the next appendix, because it is itself a fact about the list.

Every entry works by removal. The Registered Report removes the outcome from the publication decision. The heat pump and the passive house remove the flame and the demand. Setun and the neuromorphic chip remove the clock, the refresh and the shuttle between memory and processor. WIR removes the bank from the creation of money. Buurtzorg removes the management layer; the assembly removes the intermediary chain; Agora removes the timetable. The peer circle removes the professional apparatus from the encounter. The johad, the terrace and the beaver remove the drainage.

Not one entry works by adding a link, a layer, a control, or a conversion. The standard system of Chapters 1–7 grew by addition; everything in this appendix that outperforms it does so by subtraction.

Two boundaries of this list belong in the record. Additive successes exist — the vaccine, artificial fertilizer, the connected computer — and they stand on the delivery side stated in the introduction; the observation here concerns answers that outperform a standing system at its own service. And the entries have measured scopes: Buurtzorg is home care, not intensive care; the assemblies have settled moral questions, not yet tax codes; WIR carries on the order of one percent of Swiss SME turnover beside the franc, not instead of it. The regularity claimed is exactly as wide as the list: within their measured scopes, the entries that beat the standard system all beat it the same way.

That regularity — one shape, appearing independently in physics publishing, heating, hardware, money, nursing, politics, schooling, therapy and hydrology — is the last fact of the inventory. The next appendix gives it a name.

Annotated References

Chambers, C. & Tzavella, L. (2022), "The past, present and future of Registered Reports," Nature Human Behaviour. Why read? The evidence on what happens when publication precedes results: null-result rates rising toward honest levels. The measured repair of Chapter 1's incentive.

Feist, T. (various), documentation of arXiv operations and costs (arXiv.org annual reports). Why read? The running proof that global scientific communication costs millions, not billions. Compare its budget with the publisher margins of Chapter 1.

Passive House Institute, measured-performance databases; and IEA heat-pump programme reports. Why read? Three decades of measured 90% heating reductions, and the field performance of moving heat instead of making it. Both are subtraction engineering with long datasets.

Kabel, D. (2026), FD reportage on the Hungarian plain (Szabó, De Waterhoeders). Why read? The two field-scale demonstrations of section A.2 — the 15%-versus-50% harvest and the blocked canals — reported with named people and places.

Brusentsov, N. & Ramil Alvarez, J. (2011), "Ternary Computers: The Setun and the Setun 70," in Perspectives on Soviet and Russian Computing. Why read? The first-hand account of the built ternary computer by its designer: fifty machines, years of reliable service, ended by administration. Chapter 3's register, historically optional.

Davies, M. et al. (2018), "Loihi: A Neuromorphic Manycore Processor," IEEE Micro; and Kanerva, P. (1988), Sparse Distributed Memory. Why read? The clockless, refreshless hardware and the address-by-content memory — the two published, working departures from the stack that cannot rest.

Stodder, J. (2009), "Complementary Credit Networks and Macroeconomic Stability: Switzerland's Wirtschaftsring," Journal of Economic Behavior & Organization. Why read? The econometric study of nine decades of WIR: counter-cyclical turnover, measured. Money without bank-created debt, tested through every crisis since 1934.

KPMG (2015), "The added value of Buurtzorg"; and subsequent ministry evaluations. Why read? The audited numbers behind the Buurtzorg case: 8% overhead, ~40% fewer hours, best satisfaction scores. The accountancy profession confirming the removal of its own kind of apparatus.

Farrell, D., Suiter, J. et al., studies of the Irish Citizens' Assembly (Irish Political Studies and book-length treatments). Why read? The documented mechanics and outcomes of sortition deliberation on questions the party system could not settle. The shortest political chain, in statute.

Kelly, J. F. et al. (2020), "Alcoholics Anonymous and other 12-step programs for alcohol use disorder," Cochrane Database of Systematic Reviews. Why read? The strictest review format in medicine finding the zero-cost peer circle equal or superior to professional treatment. The measured value of the bare encounter.

Huber, M. et al. (2011), "How should we define health?," BMJ; with the Limburg 'Positieve Gezondheid' evaluations. Why read? The redefinition of health as capacity rather than absence of defect, and the field data showing shorter funnels into care where it is applied.

Singh, R. / Tarun Bharat Sangh documentation; and Liu, J. D. (documentary and papers), "The Loess Plateau Rehabilitation." Why read? The two large-scale water restorations of A.7: rivers returned in Rajasthan, a plateau re-greened in China — both by retention and subtraction, both satellite-verifiable.

Puttock, A. et al. (2017–2021), studies of the River Otter and Devon beaver trials, Science of the Total Environment and follow-ups. Why read? The measured hydrology of the beaver: flood peaks halved, water tables raised, at zero operating cost. Retention engineering by an animal, peer-reviewed.

Appendix B — The Fusion: One Description

Appendix to What Is There — *J. Konstapel, Leiden, 2026*

Part I recorded seven systems. Each chapter ended the same way: the system grows more complex, loses more, and nobody chose it. Appendix A recorded the working answers and ended with one fact about them: every one works by removal, none by addition.

This appendix names the structure in which both facts are the same fact. It is a description, not a defence. Nothing in Part I depends on it. Everything in Part I fits it.

B.1 The net in eight lines

Picture a fishing net, but continuous: one strand, nowhere cut.

1. There is one connected strand. Everything is a configuration of it.
2. The strand winds. Windings exist at every depth: very small and fast, very large and slow, and every scale between.
3. A configuration persists only if it closes — if it returns onto itself in phase. What closes, stays. What does not close, passes.
4. This is the only selection there is. Nothing is caused to exist; what fits, remains.
5. Closed windings at one depth form the meshes of the next depth. The knots of the fine net are the rope of the coarse net.
6. A closure carries a state: surplus, rest, or deficit — written +1, 0, −1.
7. What cannot close carries through the net as tension. Tension travels until it finds a depth where it closes, or it discharges.
8. There is nothing else. Matter, fields, money, institutions, organisms: configurations, closures, meshes, tension — at different depths of the same strand.

Everything below is these eight lines, applied.

B.2 The register: rest is a state

Line 6 deserves its own section, because Chapter 3 turned on it.

A closure in the net has three conditions, not two. It can hold more than its balance, exactly its balance, or less: +1, 0, −1. The middle is not the absence of a state. It is a state — the one in which nothing needs to be done.

Written as arithmetic, this is balanced ternary. Each position takes −1, 0 or +1. Every number and its negative use the same digits with flipped signs, so subtraction is not a separate operation. Zero sits in the middle of the number line, not at its edge. And the sum of a closed configuration is what line 6 requires of it: (+1) + 0 + (−1) = 0. A closure is internally balanced — its total is zero while its structure is not nothing.

The binary register of Chapter 3 has no middle. On and off, with off meaning gone. A system written in that register cannot represent rest; it can only simulate rest by staying active — the refresh every 64 milliseconds, the clock that never stops, the server that idles at half power. Chapter 3 measured the cost of a missing digit. Setun, in Appendix A, is the built demonstration that the digit was available in 1958.

B.3 The comma: closure is never finished

Line 3 says configurations persist by closing. Here is why closing is never complete, with the numbers.

The net has two elementary moves. A winding can double — fold onto itself, factor 2. And a closure can carry the three states of B.2 — factor 3. Growth by doubling builds the ladder 2, 4, 8, 16, 32, … Growth by the ternary state builds the ladder 3, 9, 27, 81, 243, …

Now ask when the two ladders meet: when does 2^a = 3^b?

Never. A power of 2 is even; a power of 3 is odd. No positive integers a and b satisfy the equation. The two ladders approach each other and never touch.

How close they come is exactly computable. The question "how many doublings match b triplings" is the question of approximating log 3 / log 2 = 1.58496… by fractions. The continued-fraction convergents give the best possible approximations: 2/1, 3/2, 8/5, 19/12, 65/41, 84/53, …

Take the fourth: 19 doublings against 12 triplings.

3^12 = 531,441. 2^19 = 524,288. 3^12 / 2^19 = 531,441 / 524,288 = 1.013643…

The two ladders miss by 1.36%. This number is ancient; musicians met it first, tuning twelve fifths against seven octaves, and named it the Pythagorean comma. In the net it is structural: at every depth where the doubling ladder and the ternary ladder are laid against each other, closure succeeds to within a comma and no further.

Three consequences follow, and they organize everything in Part I.

First: **depths are quantized.** Good near-closures exist only at the convergents — at 1, 2, 5, 12, 41, 53 triplings — not in between. The net does not have a continuum of stable scales. It has a ladder of them. This is why Part I found the same shapes at discrete levels — atom, cell, person, team, firm, nation — rather than a smooth continuum.

Second: **a remainder always stays.** Every closure leaves its comma — a residue that did not fit. The residue is tension (line 7). It travels to the next depth and drives the winding there. The net is therefore never finished and never silent; a fully closed universe would be a dead one, and the comma is why it is not.

Third: **the remainder can be handled in two ways.** It can be passed to the depth where it closes — that is Appendix A, every entry. Or it can be held where it does not fit, by force, with apparatus — and that is Part I, every chapter.

B.4 Constants as state

One more piece, then the rereading.

Each mesh, at each depth, has a state — call it χ: how tight it is drawn, how it is phased against its neighbours. What is measured *within* a mesh as a fixed number — a constant, a rate, a price level, a norm — is a property *of* the mesh: the value at which that mesh currently closes.

This is a constitutive relation, g(χ): the "constant" and the state are two readings of one condition, simultaneous, with no arrow between them. Change the state and the constant reads differently — not because one caused the other, but because they were never two things.

Standing physics already contains the surface reading of this. A coupling that "runs" with energy scale is a constant behaving as a state function; renormalisation is the bookkeeping of it. The net's statement is only that this is not a peculiarity of particle physics. It is what constants are.

B.5 The seven chapters, reread

Now walk back through Part I. Each chapter closed with a riddle: growth without a chooser, loss without an explanation. In the eight lines, each riddle is a description.

Chapter 1 — Science. The method cuts the one strand into separate objects: particles apart from fields, fields apart from space, mind apart from matter, discipline apart from discipline. Every cut creates two edges that were never edges, and the edges must then be glued: 26 inserted parameters, renormalisation, coupling constants, interaction terms, interdisciplinary committees. The dark sector is the plainest entry in the ledger — 95% of the universe booked as glue, because the cut pieces do not re-add to the whole. The complexity of Chapter 1 is not a property of the world. It is the price of the cutting, paid in apparatus. And the replication crisis is line 3 suspended: publication rewards claims that never have to return in phase — that are never selected by fitting — so the literature fills with configurations that would not survive one closure test. The Registered Report of Appendix A simply restores the test.

Chapter 2 — Energy. Combustion opens closures — molecular windings — at a depth far below the depth of use, and then forces the release through mesh after mesh that does not fit it: flame to steam to rotation to current to wire to motor. At every misfit, what does not pass leaves as heat. The two thirds rejected is not inefficiency to be engineered away; it is the comma of every forced passage, summed along the chain. The heat pump and the passive house work because they do not open the deep closure at all — they move, or retain, at the depth where the service lives. Chapter 2's exception, the LED, is the same move: light without the detour through fire.

Chapter 3 — Computing. The stack that cannot rest is a system built in a register without the 0 of B.2. Every pathology of the chapter follows: refresh, because a bit cannot hold itself; the clock, because nothing may settle; the idle server at half power, because off means gone; error correction on every layer, because closure is maintained by checking from outside instead of by form from inside. Kanerva's memory in Appendix A closes by form — an address is found by resemblance, and damage degrades the answer instead of destroying it — which is line 3 implemented in silicon. Setun is B.2 implemented. The MAZE engineering line builds on both, and this appendix is the frame it builds in.

Chapter 4 — Money. Money that is created as interest-bearing debt is tension issued without a closure to carry it: every unit enters with a claim attached that exceeds it, so the total claims must grow faster than the total work — Chapter 4's one-way series, 215 to 348 trillion, is this arithmetic and nothing else. The trading layer is winding without closing: positions of microseconds, currency turnover at a hundred times the trade it serves — configurations that never return onto anything and therefore only conduct tension. The crises are the discharges: tension held past the depth where it could close, releasing at the depth where it must. WIR, in Appendix A, is money as closure — created in the transaction, extinguished in the return, sum zero, structure not nothing. Nine decades of counter-cyclical data show what a closing loop does that an open one cannot: it damps instead of amplifying.

Chapter 5 — Institutions. Work, in the net, is a small loop: the nurse and the client, the teacher and the pupil, the maker and the made — need met, return received, closure. The institutional history of Chapter 5 is the cutting of these loops into functions — intake apart from care apart from billing apart from oversight — and the apparatus is the glue: registration, reporting, control on control, each layer checking from outside a closure that no longer closes inside. The finding that the apparatus grows faster than the work in every sector at once is then expected: the number of cut edges grows with every reorganization, and every edge needs glue. Buurtzorg removed the cuts — one team, one client, whole loop — and the measured 8% overhead against 25% is the glue bill of the cuts, refunded. The toeslagenaffaire is the inverse demonstration: a control apparatus with no closing loop to the people it judged, running to its own discharge, and the repair — built the same way — reproducing it.

Chapter 6 — The human being. A person is a closure among closures: held in form by returning loops — the daily encounter, the known face, the work that answers back, the rest state in which nothing need be done. Chapter 6 measured the removal of the loops, one series at a time: the counter closed, the friend hours down 70%, the encounter replaced by a screen — signal in, nothing returning, winding without closure, hour after hour. It measured the removal of the rest state: the always-on register of Chapter 3, worn as a life. And it measured the discharge column: burnout, the WIA funnel, the twentysomethings leaving — loading held past its depth. The chapter noted that no single field has an accepted explanation. In the net the reason is visible: each field holds one cut piece of the person — the psychiatric, the economic, the social, the physiological — and the curves live in the loops between the pieces, exactly where no discipline looks. The mutual-aid circle of Appendix A works because it is nothing but the loop: people who return, in phase, at no depth but the human one.

Chapter 7 — The earth. Water, soil and the living world are the retention meshes of the surface: closures that hold what falls where it falls. The 1846 Tisza regulation cut the mesh; the metre-and-a-half groundwater drop at Kiskunmajsa, 180 years later, is the same cut, still draining. The johad, the terrace, the blocked canal and the beaver are the mesh reset — and the measured returns (rivers flowing after decades, harvest loss 15% against 50%) are what a restored closure does. Geo-engineering, by contrast, proposes to act on the flux while leaving the mesh cut: to screen the incoming energy rather than to hold the water that regulates it. The relevant magnitudes stand in the record: the global forcing to be screened is 2.7 W/m²; the latent-heat lever of one millimetre per day of evapotranspiration from a re-wetted surface is 28.4 W/m² — ten times larger, locally, where people live, from setting the mesh instead of shading the net. The remedy on offer works on the wrong line of the eight.

B.6 One description

Set the two halves of the book side by side.

Part I: systems that grew by cutting and adding — cutting the strand into pieces, adding glue, holding remainders by force, in a register without rest. Their common signature, measured seven times: apparatus growing faster than work, loss growing with every link, discharge arriving on schedule, and no chooser anywhere — because addition is what cut structures do; nobody has to choose it.

Appendix A: answers that work by removal — restoring the loop, moving at the right depth, letting the mesh close, giving rest its digit. Their common signature: less apparatus, less loss, and results that hold without enforcement — because what closes, stays. That is line 3, and it needs no maintenance budget.

One structure, two ways of handling its comma. That is the entire content of this book.

The description has a name: the Vacuum.Net theory. It has ancestors, and they are named in the foundational papers: monism, process philosophy, general systems theory, cybernetics, the laws of form. What the present form adds to that lineage is the arithmetic — the register, the comma, the depth ladder — and the running implementations. Its full form — the derivation of the closure condition, the depth ladder and its convergents, the state function χ and the constants it carries, the ternary register and its arithmetic, and the applications from matter to markets — is developed step by step in the foundational papers and the working notes at constable.blog, and runs today in the engineering line of Appendix A: SWARP, the living-environment instruments, MAZE.

Part I stood without this appendix; every figure in it belongs to the institutions that produced it. This appendix adds no figure. It adds the net in which the figures were never separate facts — and in which the way out was already visible, in every chapter, as the one thing that worked.

Annotated References

Konstapel, J. (2026), "The Vacuum.Net Theory: Foundational Paper" and "The Complete Derivation of the Vacuum.Net Theory," constable.blog. Why read? The full development of the eight lines of B.1: the strand, the closure condition, the depth structure and the state function. This appendix is the compact form; the papers carry the derivations.

Konstapel, J. (2026), "The Comma in the Net — Why Closure Depth Is Quantized," constable.blog. Why read? The complete treatment of B.3: the two ladders, the continued-fraction convergents, and why stable depths sit at 1, 2, 5, 12, 41, 53. The arithmetic of this appendix, taken to full depth.

Knuth, D. (1997), The Art of Computer Programming, Vol. 2, section on balanced ternary. Why read? The standard technical description of the register of B.2, including Knuth's judgement of balanced ternary as perhaps the most elegant number system. Independent of any theory; the arithmetic stands on its own.

Brusentsov, N. & Ramil Alvarez, J. (2011), "Ternary Computers: The Setun and the Setun 70." Why read? The built existence proof that the register with rest was engineering, not speculation — fifty machines, years of service, ended by decision.

Kanerva, P. (1988), Sparse Distributed Memory. MIT Press. Why read? Closure by form in working mathematics: retrieval by resemblance, graceful degradation, addresses as places in a space rather than numbers in a row. The memory architecture of B.5's Chapter-3 reading, and a design basis of MAZE.

Spencer-Brown, G. (1969), Laws of Form. Why read? The distinction as the first act, and re-entry as the first closure — the formal ancestry of lines 1 and 3, written before any of the present applications existed.

Barbour, J. V. (2004 and editions), "Pythagorean comma" treatments in Tuning and Temperament: A Historical Survey. Why read? The musical history of the 1.36% — twelve fifths against seven octaves — where the comma was first met, measured and lived with. The oldest documented encounter with the arithmetic of B.3.

Konstapel, J. (2026), "Setting the State, Not the Screen" and "Closing the Mesh," constable.blog. Why read? The two applied papers behind the Chapter-7 rereading: the screen arithmetic (2.7 W/m^2 forcing, 18 Mt SO_2 per year steady state) against the latent-heat lever (28.4 W/m^2 per mm/day), and the Hungarian storage calculation. Both self-contained, both priced.

**McLeay, M., Radia, A. & Thomas, R. (2014), "Money creation in the modern economy," Bank of

England Quarterly Bulletin. Why read? Re-read after B.5: the central bank's own mechanics of money as issued tension, which the WIR loop closes and the standard loop cannot. The same fourteen pages as in Chapter 4, now legible as net arithmetic.

KPMG (2015), "The added value of Buurtzorg." Why read? Re-read after B.5: the audited refund of the glue bill when the loop is restored. The single cleanest measured case of the entire description.