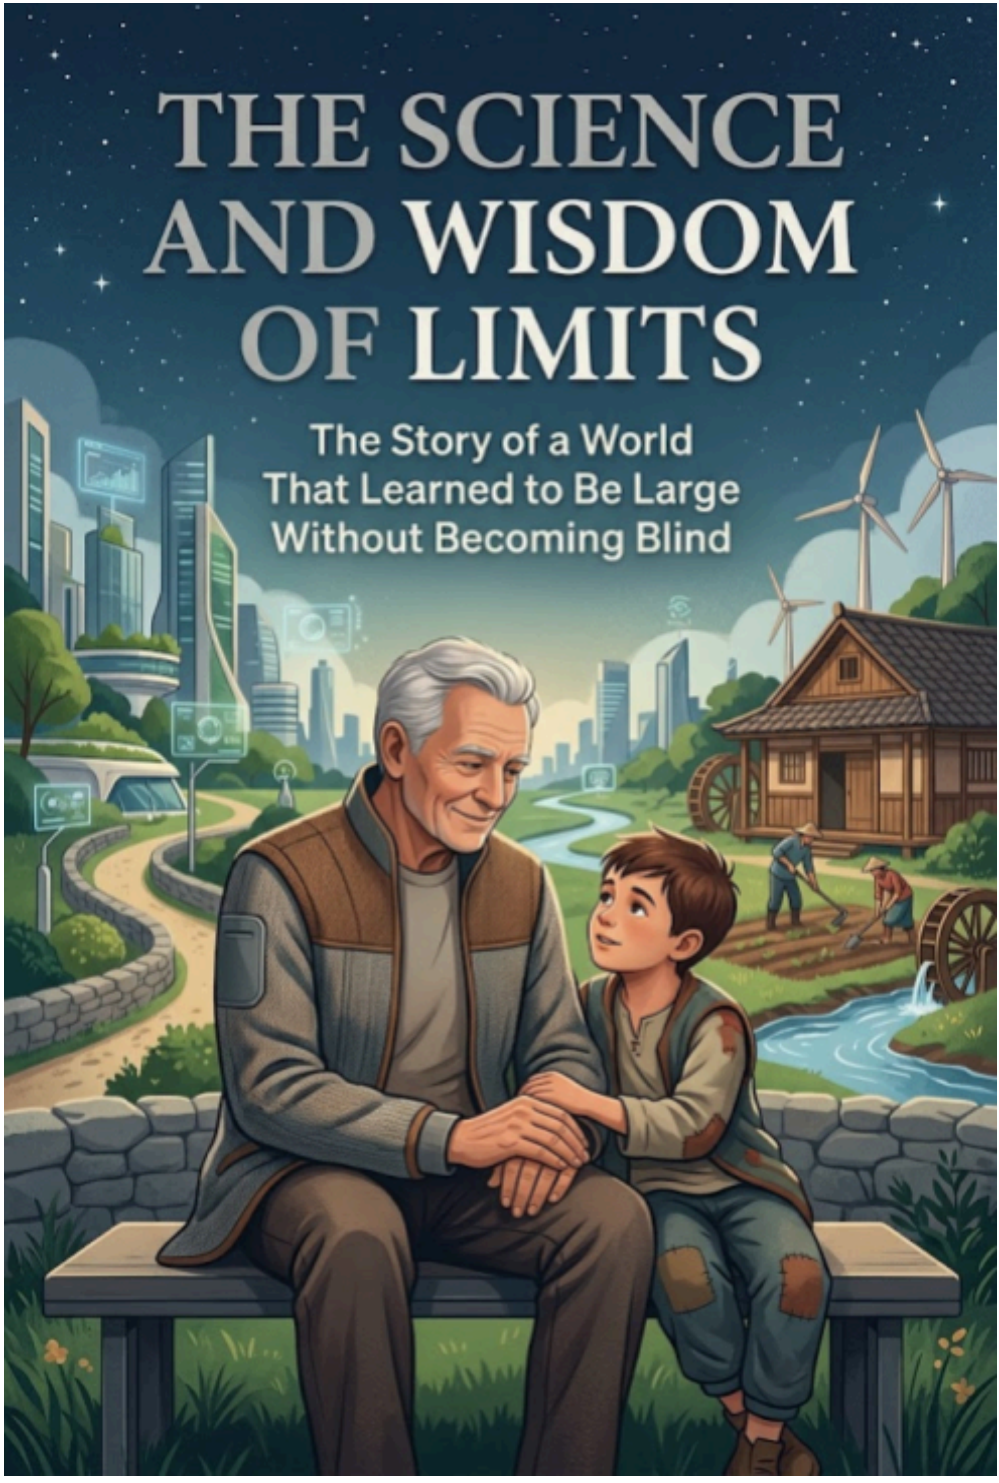
THE SCIENCE AND WISDOM OF LIMITS
The Story of a World That Learned to Be Large Without Becoming Blind

THE SCIENCE AND WISDOM OF LIMITS

The Story of a World That Learned to Be Large Without Becoming Blind

Author's Note

This work was created under the umbrella of Axiologic Research as part of two interconnected research programs, one dedicated to Artificial Intelligence and the other to Outfinitism, both led by Sînică Alboaie. To explore the details of how these two programs intersect and build upon each other, we invite readers to visit the Outfinity Venture Studio page at www.outfinity.ch.

While Artificial Intelligence language models were used to assist with text generation, the foundational philosophy, thematic structure, and analytical approach derive exclusively from the author's decades of dedicated study of social phenomena, social technologies, and genuine hard technologies resulting from various European and self funded research projects. Recently, within the Achilles research project, we have been deeply studying AI generated literature and its automated verification using neuro symbolic approaches; all the free books made available to the public are part of the suite of works we use to test and popularise our latest ideas and technologies.

A Note to the Reader

This story takes place in a possible future. Its characters, places, and future institutions are fictional. The ideas concerning complexity, bounded rationality, polycentric governance, collective intelligence, interdependent infrastructures, planetary boundaries, and collaboration between humans and artificial intelligence draw on real research, discussed in the scientific note and bibliography at the end of the book.

This book does not argue for a return to a small world, nor for placing the planet under a single intelligence. Its question is harder: how can we build systems large enough for the problems we face, yet intelligible, contestable, and repairable enough for the people who must live within them?

Contents

PROLOGUE

An Assignment for Monday

In the spring of 2094, three days before the solar storm was due to reach Earth, my grandson walked into the greenhouse holding a sheet of paper as if it were a warrant for his arrest.

"I have a problem," Matei said.

He was twelve, and at that age problems came in two sizes. Some were absolute catastrophes, such as a school project. The others were not even worth looking up for, such as the warning issued that morning by the Solar Observatory.

I was tying the tomato plants to their strings. The greenhouse was small, attached to our house, with an irrigation system simple enough to repair with a screwdriver and patience. That was precisely why I liked it. Outside, almost every object had been built from layers of technology that even engineers no longer followed all the way down. Inside, if a plant wilted, the cause was usually less than two meters away.

"What kind of problem?"

"I have to choose one rule for how the world should be organized over the next hundred years. One rule. And I have to explain it on Monday."

"Sounds like a modest assignment."

"That is exactly what I said."

He handed me the sheet. His teacher had phrased the task more carefully than he had summarized it. The students had to choose a principle of social organization, show what problem it solved, what risks it created, and under what conditions it should be abandoned. It was not bad. When I was a child, assignments like this usually demanded certainty and three arguments. In 2094, school required children to include reasons why they might be wrong.

"And what rule did you choose?"

“I haven’t yet. But I thought it would be simpler if artificial intelligences calculated what was best, and people voted only when the models disagreed.”

I clipped a tomato stem to its support and looked at him.

“Why only when the models disagree?”

“Because people argue all the time. Models don’t get tired, they don’t forget, and they can see all the data.”

“Can they see all the data?”

“Almost all of it.”

“Can they know which data are missing?”

“Sometimes.”

“Can they decide for themselves what is worth optimizing?”

Matei rolled his eyes.

“If you start with grandfather questions, I’ll finish the project after the storm.”

A short chime came from the speaker in the corner. The house system displayed the same notice every home in the region had received: solar activity had increased, and the estimated severity of the geomagnetic storm had been raised by one level. The energy grid would enter modular mode as a precaution. High-speed rail service might be reduced, some orbital links would be shut down, and cities were to verify their ability to operate temporarily as electrical islands.

Matei looked at the message and shrugged.

“The grid will handle it.”

“Which grid?”

“The whole grid.”

“There is no single grid.”

“You know what I mean.”

I did. Children of his generation had grown up with infrastructures that presented themselves as continuous services. Water came, electricity came, transport came, information came. The fact that thousands of organizations, protocols, and people stood behind each of those verbs was about as visible to them as a tree’s roots are to someone looking at its crown.

“That is exactly where your assignment begins,” I told him. “The world did not become hard to govern because people became less intelligent. It became hard to govern because we built things whose size was no longer measured only

in kilometers or numbers of employees. It was measured by how many consequences they could produce before anyone understood them."

Matei set the paper on the workbench.

"That sounds like something a professor says before a two-hour lecture."

"You don't have to listen to me for two hours. Come with me tomorrow to the Systems Archive. Then we'll go to the energy cooperative. On Thursday I'm attending the water meeting, and on Friday we can visit the civic modeling center. If, at the end, you still think a single intelligence should run everything, at least you'll know what you're asking it to do."

"And the storm?"

"The storm will provide the practical demonstration."

That interested him. At twelve, an idea became more respectable if there was some chance the lights might go out during the explanation.

We went inside. On the kitchen table, the local system projected a list of recommended checks. It was not telling people to stock up for the apocalypse. The storm was not expected to cause a disaster, and the regional grid had been designed to fragment in a controlled way. Still, backup batteries had to be charged, radios checked, lists of people dependent on medical equipment updated, and water reserves in high-rise buildings confirmed.

"Why doesn't the system do all of this by itself?" Matei asked.

"Because the system doesn't know that Mr. Banu moved his breathing machine into the workshop yesterday. The database still has it in the bedroom."

"He should have updated it."

"He should have. But he didn't. The real world is full of people who should have updated something."

I handed him the portable radio and asked him to check the battery. He opened the cover, looked inside, and smiled.

"It's almost dead."

"According to the system, it was charged."

"Then the system was wrong."

"No. The system correctly reported the last state it knew. We were wrong when we confused the report with reality."

He sat at the table and opened his notebook. On the first page he wrote: "A large world needs good data." Then he stopped.

"That isn't enough, is it?"

"No. Sometimes it has good data and bad questions. Other times it has good questions, but the information arrives too late. Or it reaches someone who is not allowed to stop the process. Or people see the problem, but each person sees only one piece and assumes somebody else understands the whole."

"So what changed? How did we end up with the systems we have now?"

I looked at the date in the corner of the window. More than sixty years had passed since I had worked on the first infrastructures in which artificial agents made decisions affecting millions of people. It had not been an age of evil robots, nor an age of geniuses saving the world with an app. It had been more ordinary than that, and for that very reason harder to correct. Every organization had been trying to solve a real problem. Each optimized its own part. Together, we built systems that moved faster than our ability to judge them.

We had not experienced one great collapse. We had experienced thousands of things that worked well enough to scale and badly enough to create costs nobody noticed in time.

"We learned slowly," I said. "First we thought the problem was a lack of information. Then we gathered more information than anyone could use. We thought the problem was a lack of coordination. Then we connected everything and discovered that an error could travel faster than its explanation. We thought the problem was weak institutions. We gave them stronger tools and sometimes got more efficient mistakes. Eventually, we began asking not only how powerful a system should be, but how large it could become before it could no longer understand itself."

Matei wrote that down too, although he shortened it.

"When does a system become too large to understand itself?"

"That is a better question than your rule."

"The assignment asks for a rule."

"Good questions have the inconvenient habit of coming before good rules."

That evening we checked the house, the street's shared battery, and the neighborhood communication system. None of it was spectacular. That was the point. Resilience looks boring before a crisis and intelligent afterward.

At bedtime, Matei left his notebook on the hall table. Under the first question he had added another:

"If nobody understands the whole, who is responsible for what the whole does?"

That was where we would begin the next day.

CHAPTER 1

The World That Had Grown Beyond Sight

The Systems Archive occupied a former railway shed on the edge of the city. Freight cars had once passed through it. Now it received school groups, civil servants in training, and people who wanted to remember why modern infrastructures had been built with visible switches.

Matei expected a museum of old computers. Instead he found a grocery store, a control room, a hospital room, a recruitment office, and a wall covered in cables. All were working reconstructions from the first decades of the century.

“Where are the robots?” he asked.

“In many of the things you’re looking at, except that back then they didn’t have arms and eyes. They calculated prices, selected candidates, planned routes, recommended treatments, classified risks, and decided which message a person would see next. When people said ‘robot,’ they imagined a machine walking through the door. The most important ones came in through forms.”

The first installation was a luminous map of a European electricity network in 2003. A virtual guide could replay events in real time. Matei started the simulation. Green lines represented normal flows. Then one line in Switzerland turned orange, then red. Another took over the load. A few seconds later the colors began changing across a much larger area. In less than half an hour, much of Italy had lost power.

“Because of a tree?” he asked.

“Because vegetation came into contact with a line, but not because of that alone. The tree was the visible beginning. The real cause was the combination of network loading, coordination delays, automatic protections, and dependencies among components. If you cut power to a house with a pair of shears, the shears look like the cause. If the same disturbance starts a continental cascade, you also have to ask why the system allowed it to travel.”

“So the whole thing was badly designed?”

“No. The grid had supplied electricity to many millions of people. A system can be very good at its ordinary job and weak against a particular kind of surprise. That is what makes serious problems harder to see. If a bridge collapses on opening day, everyone recognizes the failure. If it works for thirty years and then a rare combination of conditions destroys it, people argue over whether it was badly designed, badly maintained, or simply unlucky.”

On the screen, the simulation returned to the beginning. Matei started it again, this time slowing down each step. At that speed, everything looked easy to prevent. One operator should have reduced the flow. Another should have understood the risk. Someone should have called sooner.

“Now it looks obvious,” he said.

“Once you know what happened, the past acquires an intelligence the people living through it did not have. That is one of the traps of explanation. A good story erases all the futures that might have followed and keeps only the road events actually took.”

We moved to the next installation. It was a table covered with ordinary objects: a mug, a phone, a shirt, a battery, a box of medicine. Touch an object and the walls filled with maps. The mug summoned mines, ports, factories, shipping companies, warehouses, contracts, and payment systems. The shirt opened an even longer network. The phone turned the hall into a sky of connected points.

“How many people had to know the whole route?” I asked him.

“Someone at the company.”

“Which company?”

He looked at the map again. Hundreds of firms appeared for a single phone. Some produced materials, some designed components, some wrote software, provided transport, certified standards, or financed inventory.

“Then maybe nobody.”

“Exactly. Civilization became capable of producing things no single person understood in full. That was not an accidental defect. It was the source of its power.”

I told him about distributed knowledge. A large ship was not navigated by one person holding every measurement, chart, and procedure in mind. A team, together with its instruments, could perform a cognitive operation more

complex than any one member could manage alone. The same thing happened in a hospital, a university, a factory, or a transport network.

"So the group becomes smarter than the people?"

"Sometimes. Other times it becomes stupid in a way no single person would have chosen."

"How?"

"Information can exist in the system without reaching the place where it changes the decision. The technician knows a component is overheating. Procurement knows the supplier was changed. The manager knows the deadline cannot move. The customer reports a strange sound. Every piece of information is real. Nobody sees the configuration."

"But an AI could put those pieces together."

"It could. If the data are recorded, if the systems can communicate, if the terms mean the same thing, if nobody has a reason to hide information, and if the model knows which combination matters. That is a lot of conditions packed into one 'could.'"

At the far end of the hall stood a reconstruction of a command center from the 2030s. The walls were covered with charts. Some moved so fast they were almost decorative. There were indicators for deliveries, satisfaction, consumption, risk, productivity, emissions, and reputation. When all of them were green, the room felt reassuring.

"Is that what your systems looked like?" Matei asked.

"Some of them. We liked dashboards back then. We thought that if something was important, it should be measured, and if it was measured, it should go into a colored rectangle."

"Isn't that logical?"

"It is, until people start working for the rectangle."

I activated a teaching scenario. A patient support service had been evaluated by average call duration. At first the metric helped: pointless calls got shorter, and people received answers faster. After a few months, operators learned to close complicated cases and send them to another department. Average call time fell. Patients waited longer. On the screen, efficiency turned green at exactly the moment the service became worse.

Matei laughed.

"That's cheating."

“Sometimes. Other times it is ordinary adaptation. Once you turn a measure into a target, people discover how to hit the target. They do not have to be malicious. An employee protects a job. A manager protects a budget. A department protects its reputation. The system selects behavior that improves the number.”

“Then you change the number.”

“And people learn the new number. Or you add ten indicators and the work turns into managing indicators. The problem is not that measurement is useless. Without it we would be blind. The problem is that every measure is a window, and we tend to forget the wall around it.”

I brought up an old image of Herbert Simon, an economist, political scientist, and one of the founders of artificial intelligence. Matei knew the name from a course, though he associated Simon mainly with early problem-solving programs.

“Simon insisted that people do not make decisions with complete information, infinite time, and unlimited computing power,” I told him. “We look for good-enough solutions using partial models of the world. That is not a character flaw. It is the condition of any finite mind.”

“But if AI has more memory and computes faster, the limit disappears.”

“It moves. A more powerful system can consider more possibilities, but it still has to decide which variables matter, what objective to pursue, and when to stop. Give it the wrong description of the problem and it can be wrong with impressive precision.”

I showed him an old example from urban routing. A city had optimized traffic to reduce average travel time. The model found efficient solutions, but some of them shifted traffic through neighborhoods inhabited mainly by lower-income residents. The average improved. Noise, accidents, and pollution became more concentrated. It had not been a conspiracy against those neighborhoods. The model simply did not treat the distribution of costs as a primary objective.

“Who chose the objective?”

“A committee, working from a mandate, some data, and a budget. Perhaps nobody on the committee thought of that effect. Perhaps some did but lacked enough evidence. Perhaps they had evidence, but the political timetable demanded a quick result. When the chain becomes long, you no longer find

one person who decided everything. You find many incomplete decisions that fit into one another."

"So nobody is to blame?"

"No. It means responsibility cannot be understood only as a hunt for a single culprit. People are responsible for their part, leaders for the architecture, and the institution for the mechanisms by which it detects and corrects effects. If everyone can say, 'I followed the procedure,' while the procedure repeatedly causes harm, then the procedure itself becomes the object of judgment."

I called this situation ghost agency. A system can appear to pursue a goal even though no participant chose that goal in full. A platform seems to prefer outrage because outraged messages keep people engaged. A company seems obsessed with growth because every internal mechanism rewards expansion. A bureaucracy seems to want more forms because every exception creates a new rule and every rule acquires a protector.

Matei grimaced.

"Ghost agency makes it sound as though the system is haunted."

"It's just a name. The system does not need ghosts. It needs incentives that repeat."

"Then why doesn't someone stop it?"

"Because in many systems, the power to notice and the power to stop are in different places."

That was the sentence he wrote down.

We walked through a corridor projecting messages from the social networks of the early century. Some were funny. Others were scandals that had lasted a day. Some were lies shared millions of times. The exhibit let visitors modify the recommendation algorithm. Increase the weight given to reactions and conflict-driven content rose. Increase the weight given to time spent and longer material appeared, but so did more attention traps. Reduce personalization and diversity increased, but many users left.

"Wasn't there a button for truth?" Matei asked.

"If there had been, the dispute would simply have moved to who defined truth, using which sources, and how quickly. Some things can be verified. Others are judgments, predictions, or conflicts of value. Platforms tried to treat all of them as classification problems. Sometimes they succeeded. Other times

they turned social arbitration into a hidden function inside a commercial model."

"But people could leave."

"In theory. In practice, their friends, customers, reputation, professional groups, and public memory were there. A door truly exists only if you can walk through it without leaving too much of your life behind."

We reached the section devoted to old supply chains. A game asked the visitor to manufacture a simple medicine using suppliers from around the world. At first Matei always chose the cheapest supplier. His score rose. Then a drought reduced one raw material, a port closed, and a manufacturer halted its line over a quality problem. Within two minutes, his virtual system could no longer deliver the medicine.

"The game is cheating," he said.

"It isn't cheating. It is showing you the price of efficiency when efficiency is calculated only for ordinary conditions."

"Then I keep more suppliers."

"That costs money."

"I build up inventories."

"That costs money."

"I produce locally."

"That usually costs more, and sometimes you lack the resources or expertise."

He sighed.

"So every option is bad."

"No. Every option has a cost. One of the great mistakes of my era was to call something a cost only if it appeared immediately in the accounts. Redundancy looked like waste. Unused capacity looked like laziness. The alternative supplier looked unnecessary until the day it became the only one able to deliver."

I explained that large systems are not necessarily fragile. Some become more robust precisely because of their size: they can distribute load, maintain reserves, and mobilize specialists. Fragility appears when connections are so tight that failure propagates and actors are so specialized that nobody can improvise outside a role.

"A body is a large system," Matei said. "And it works."

“Good example. But a body is not a sack full of mixed cells. It has organs, membranes, signals, redundancies, and reflexes. The liver does not ask the brain for permission for every chemical reaction. Your hand pulls away from a hot surface before you finish explaining why. Some functions are local and fast, others global and slow.”

“Then society should be built like a body.”

“Up to a point. The metaphor becomes dangerous if we assume society has a single brain and people are merely cells that must obey. A body does not ask its cells for consent. A society has to preserve persons with purposes of their own, rights, and the ability to challenge the whole.”

We left the main hall and stepped into a courtyard where an old locomotive had been preserved. Matei climbed onto the platform and looked at tracks disappearing beneath vegetation.

“How do you know when a system has become too large?”

“There is no universal number. A community of fifty can be too large for an intimate discussion and too small for a transplant hospital. An organization with ten employees can be opaque if it depends on a model nobody understands. A planetary asteroid-observation system can be quite clear if it has a narrow purpose and public procedures.”

“Then what do you measure?”

I took his notebook from the backpack and drew two circles. In the first I wrote “consequences.” In the second, “what we can observe and correct.”

“When the first circle grows far beyond the second, what I call scale debt appears. The system produces effects faster and farther than it can understand and repair them.”

“Like financial debt?”

“Something like it. You can operate for a long time with it. It may even help you grow. But you pay interest in the form of errors, delays, mistrust, and fragility. And when a surprise comes, you discover that much of your apparent capacity had been borrowed from the future.”

“Will the solar storm test our scale debt?”

“It will test whether our grid can become smaller before the disturbance breaks it into pieces.”

On the way out, we passed a display case containing a phone from 2025. Matei looked at it like a prehistoric tool.

“You did everything with that?”

“Almost. We talked, shopped, worked, followed the news, argued with strangers, and received notifications telling us we ought to sleep more, sent by the device that was keeping us awake.”

“Sounds stupid.”

“It was also extraordinary. Don’t forget that. Serious criticism begins by acknowledging what a system made possible. Otherwise you cannot explain why people adopted it.”

On the tram home, Matei stayed quiet. Then he opened his notebook and wrote three sentences:

“A large system is not bad simply because it is large.”

“It becomes dangerous when its effects move beyond the reach of those who can understand and stop them.”

“Sometimes nobody decides the whole, but the whole decides for everyone.”

He circled the last sentence.

“That one’s good,” he said.

“Good because it makes you uneasy, or because you can defend it?”

He erased the circle.

“Tomorrow we test it.”

CHAPTER 2

Decisions Without a Decider

The next morning, Matei showed me a message from school. His teacher had read the questions the students uploaded to the class space and asked them to identify a case in which an injustice could arise without anyone explicitly wanting it.

"I already answered," he said. "I wrote 'an algorithm.'"

"An algorithm does not want anything."

"Exactly."

"And yet you blamed it."

Matei stopped halfway through putting on one boot.

"Then who do I blame?"

"Today we're going to find out."

The Center for Administrative History was two stops from the Systems Archive, but it looked as though it had been built by another civilization. Yesterday's museum had displayed machines and networks. This one preserved decisions: forms, evaluation criteria, contracts, scores, and appeal procedures. In the entrance hall a simple sentence was projected on the wall: "Every classification divides the world into those who pass and those who remain outside."

Matei read it twice.

"That makes it sound as though all rules are bad."

"No. Without classifications you could not allocate places in a school, prioritize a patient, or check whether a bridge is safe. The question is what you classify, using what data, for what purpose, and what the classified person can do if the model is wrong."

The first room reconstructed a recruitment process from the 2030s. Visitors had to apply for an imaginary job. Matei filled in the profile of a strong candidate but left a two-year gap in the employment history. The system rejected him.

"Why?"

A polite answer appeared on the screen: the profile did not correspond closely enough to patterns of success identified in historical data.

"That doesn't say anything," he protested.

"It says that people who were hired and promoted in the past were less likely to have gaps like that."

"Maybe the candidate was caring for someone who was ill."

"The model doesn't know."

"Then it should ask."

"Questions cost time. The company had received fifty thousand applications. The system had been introduced to reduce human workload."

Matei requested the detailed explanation. He learned that the employment gap, frequent changes of address, and the absence of certain words from the cover letter had lowered the score. No single criterion had been declared decisive. Together, they had pushed him below the threshold.

"Can the candidate appeal?"

He pressed the button. In the historical version of the procedure, an appeal was available only if the person could demonstrate a factual error. The statistical model itself could not be challenged because the company regarded it as a trade secret.

"So nobody tells him why he lost, and he can't show that the explanation is bad."

"That sometimes happened. The system was efficient for the company and impossible for the individual to interrogate."

"Who decided that?"

"The recruitment team chose the objective. The vendor built the model. The legal department limited the explanations. Management set the budget. The data came from decisions made by managers in the past. Everyone added something. The result belonged completely to no one."

"Still, the company is responsible."

"Yes. That is the difference between an explanation and an excuse. Distributed harm does not cease to be harm. It means accountability has to be designed at the same level at which the effect appears."

We entered a room filled with dozens of transparent boxes. Each contained an object and a label: work wristband, credit contract, delivery sensor, medical

form, housing-rental reputation system. When we brought a hand close, the box projected the network of decisions connected to the object.

Matei chose the work wristband. It had been used in a warehouse to track employees' movement and calculate an optimal pace for tasks. The system reduced unnecessary walking and coordinated inventory better. At the same time, it could turn every pause, hesitation, or deviation into a performance signal.

"Why did people accept it?"

"Because they needed the job. Because the pay might be better. Because some preferred clear instructions. Because the company promised greater safety and more objective evaluations. Not all the effects were negative."

"But what if the pace was too high?"

"Managers could say the system had calculated it from the data. The vendor could say the company had chosen the parameters. The company could say the workers had accepted the contract. The worker felt the pain, but the explanation moved through the organization without stopping anywhere."

"That's ghost agency."

"It is one example. But there is something else as well: governance distance."

I explained that a decision can be near or far away in more than a geographic sense. Distance grows when those who set the rules do not meet those affected by them, when effects appear much later, when the mechanism is hard to understand, or when no accessible person has the authority to change the outcome.

A manager in the same building can be institutionally farther away than a judge in another city if the judge can review the decision while the manager merely executes a score.

"So distance can be measured by how many people you have to ask before you find someone who can say yes or no?"

"That is one measure. It also matters whether the person understands the case, can see the consequence, and bears some cost when the decision is wrong."

In the next room, visitors ran an imaginary company. Matei had to choose objectives for procurement, sales, logistics, legal, and human resources. He set targets for lower costs, higher sales, and lower risk. After several rounds, every

department had met its metrics. Profit was up. At the same time, suppliers had become more fragile, employees were leaving faster, and complaints were rising.

"The game wants me to lose," he said.

"The game wants you to notice that local optimization does not automatically produce a good whole."

"Then I'll add a metric for employee satisfaction."

He added it. After a few more rounds, the human resources department had learned to send the survey at favorable times and exclude people who had already left from the calculation. Reported satisfaction went up.

"That really is cheating."

"It is also institutional pressure. If the department's budget depends on the metric, people will find a definition that can be improved. They do not have to lie outright. It is enough to decide who counts."

We talked about the old observation that when a measure becomes a target, it often stops being a good measure. Not because numbers are enemies of reality, but because people respond to the rules by which they are evaluated. A social system is not a passive object. It is made of participants who observe the observer.

"Then how do you measure something without damaging what you're measuring?"

"You use several sources, check effects outside the metric, sometimes change the measures, preserve human judgment, and above all refuse to confuse the score with the purpose. But there is no perfect protection. Every evaluation system creates behavior."

We reached an installation the museum called "The Table of Invisible Costs." On it lay ordinary products from the past: very cheap meat, clothes delivered in a day, electronics almost impossible to repair, free digital services. Visitors could lift the displayed price and see what had been left outside it: pollution, soil depletion, accidents, attention time, poorly paid labor, waste, risks borne by suppliers or communities.

Matei picked up a pair of shoes.

"If all these costs are real, why weren't they in the price?"

"Because some were hard to measure, some appeared in another country or twenty years later, and some were borne by people who were not part of the transaction. Economists called them externalities."

"So the bill goes to someone else."

"Sometimes, yes. Or it comes back to the same person in another form. They buy cheaply, but pay through dirtier air, cleanup taxes, health costs, or the loss of an ecosystem."

In his notebook I drew an arrow from action to effect and called the distance between them the externality radius. Some choices had a radius of a few meters. Others crossed continents and generations. The farther the radius extended beyond the community making the decision, the easier it became for the benefit to stay local while the cost was pushed outward.

Matei looked at the table again.

"Is that parasitism?"

"Sometimes. But we have to use the word carefully. A biological parasite is an organism. In society, we can have an actor who exploits consciously, or a structure in which benefits systematically flow in one direction and costs in another. I call the second situation a parasitic gradient."

"An excellent term if you want children to close the book."

"You're probably right. Call it a 'cost slope': the good things roll toward whoever has power, and the bad things toward whoever cannot defend themselves."

He liked that phrasing better.

I then took him to a screen simulating a free digital platform. The platform offered communication, entertainment, and excellent work tools. It charged most users nothing. It made money from advertising and from its ability to predict what kept people connected.

Matei adjusted the algorithm. When he increased time spent, revenue rose. After many rounds, users became more tired, more irritable, and less able to control their attention. No single session looked serious. The effect emerged from billions of small minutes.

"You can't say the platform stole a minute at a time," he said. "People chose to stay."

"The choice existed. But the platform tested millions of variants to discover what reduced the probability of leaving. When one participant has one mind and the other has an infrastructure continuously experimenting on that mind's behavior, formal freedom does not describe the whole relationship."

"Then you ban optimization."

“And lose many useful things: good recommendations, fraud detection, translation, accessibility, adaptive services. The problem was not optimization itself but a narrow objective and an asymmetry. The platform could learn about the user faster than the user could learn about the platform.”

“What did people do about it?”

“Late in the process, we began requiring portability, audits, limits on behavioral experiments, and the ability to use alternative interfaces over the same social network. It became harder for one company to own identity, relationships, history, and access rules at the same time.”

“So you could leave without losing everyone.”

“Exactly. A real exit is not a ‘delete account’ button. It is the ability to continue your life somewhere else without being socially or economically amputated.”

That was when I told him about Albert Hirschman and the relationship between exit and voice. When a system works badly, people can try to leave it or change it from within. If exit is impossible and voice has no effect, loyalty becomes captivity.

“But some systems can’t let everyone take everything with them,” Matei said. “You can’t bring part of the railway home.”

“Correct. That is why, where technical exit is difficult, the right to challenge must be stronger. The more you depend on a system, the clearer the path of appeal that system owes you.”

We stopped at the museum café. The serving robot brought two teas and displayed the source of the ingredients, energy use, and price. Matei kept opening details until he reached the contract of the cooperative that had produced the honey.

“Now we can see the whole route.”

“We can see a representation of it. Don’t forget the difference.”

“There you go again.”

“A transparent system can display false, incomplete, or irrelevant data. Transparency does not mean dumping thousands of documents on a table. It means important information can be traced to its source and connected to a possibility for action.”

“So if I discover that the honey comes from somewhere the bees are treated badly, I should be able to change suppliers.”

"Or flag the problem, request verification, choose another product. Information without options is sometimes just ethical decoration."

After the break, we entered the last room. Four sentences were written on one wall, formulated after the administrative reforms of the 2060s:

An important decision must leave a verifiable trace.

The affected person must be able to discover what mattered.

Someone must have the authority to correct the outcome.

The cost of error must not be left entirely to the person harmed by it.

Matei read them and said they seemed obvious.

"The most expensive ideas in history tend to look obvious after they have been paid for," I replied.

"How do you make the system bear the cost of error?"

"Through compensation, audits, professional liability, model revision, automatic suspension when patterns of error emerge, and an obligation to keep resources available for repair. If an organization profits from automating a million decisions, it cannot treat every mistake as the user's private misfortune."

"And if nobody intended the harm?"

"Intent matters when judging an individual's culpability. It is not enough when judging an architecture. A bridge that collapses without malicious intent still has to be redesigned."

As we left, rain had begun to fall. We walked beneath the covered walkway toward the tram stop. Matei watched the water gather in gutters and flow into the city's reservoirs.

"I think I understand why 'an algorithm' is not a good answer," he said. "The algorithm is one part. You have to look at who gives it its objective, who builds it, who can stop it, and who pays when it is wrong."

"Yes."

"But my question is still there. If everyone only made one piece, who is responsible for the whole?"

"The organization that put the pieces together. The leaders who accepted the architecture. The authorities that allowed it. Professionals for their decisions. Sometimes users as well, if they had real power and enough information. Responsibility can be shared without being evaporated."

"That sounds complicated."

“It is. Simple solutions to complex problems are often just elegant ways of moving the cost outside the sentence.”

On the tram he opened his notebook again. Beneath yesterday’s questions he wrote:

“A decision is legitimate only if its effects can travel back to someone who has both the power and the duty to correct it.”

Then, in parentheses, he added:

“It is not enough to have an appeal form.”

This time I changed nothing.

CHAPTER 3

Why We Did Not All Move to Villages

On the third day, Matei came to breakfast with a solution.

"I understand now. Large systems are hard to control. So we make them small."

"All of them?"

"As much as possible. Small communities, local production, decisions made by the people affected. Then nobody can hide behind procedure."

"Today we're going to the market," I said.

"To prove that local tomatoes solve governance?"

"To show you why your idea is good and insufficient."

Our market did not look like the markets of my childhood. The roof produced energy, rainwater was collected, and every producer could display the route taken by the goods. Yet the essential part had remained the same: people knew one another. The woman selling apples knew which variety Matei liked. The baker knew who had forgotten to pay the previous week. A retired carpenter repaired his neighbors' crates for free, and the neighbors brought him food when his back hurt.

"It works here," Matei said after a few minutes. "If someone sells something bad, everyone finds out."

"Reputation is an old and efficient institution."

"And cheap."

"Very cheap, unless you are the person a false story is being told about."

Behind the stalls was a community workshop. That was where we usually found Mara, who had once been a magistrate and, after retiring, had become one of the best repairers of old machines in the city. That day she was taking apart a fan manufactured almost seventy years earlier.

"You brought the student?" she asked me.

"He has decided the world should be divided into small communities."

“Excellent,” she said. “Then he can begin by spending a month in one where his aunt controls the shop, the mayor is his godfather, and the only employer plays chess with the police chief.”

Matei sat down on a chair.

“Why do you assume a small community is corrupt?”

“I don’t. I’m saying small size does not vaccinate people against power. Sometimes it makes power more visible. Other times it makes power harder to escape.”

Mara had grown up in a mountain town where almost everyone depended on a single enterprise. People helped one another, left their keys in the door, and organized celebrations together. At the same time, anyone who came into conflict with the director could lose not only a job but access to housing, a reputation, and the chance that any relative would ever be hired.

“Wasn’t there an internet?” Matei asked.

“There was. But life does not move somewhere else just because you can send messages. My parents were there, the house was there, and my father’s skills were useful in that factory. Formally, we could leave. In practice, the cost was enormous.”

“That is domination without exit,” he said.

I looked at Mara.

“He does his homework.”

She closed the fan and showed him an almost invisible crack in the housing.

“In small systems, power is like this crack. You can see it more easily, but it runs through the objects you use every day. You cannot simply say, ‘I won’t buy from power anymore.’ Power may be your landlord, the doctor’s cousin, and the man who runs mountain rescue.”

Matei asked whether that was not still better than an anonymous corporation.

“Depends,” Mara said. “An anonymous corporation can treat you like a number. A village can treat you like the story it knows about your grandfather. Sometimes you need the right not to be fully known.”

We left the workshop and sat at a table near the bakery. The market was a good place to explain why people romanticize small scale. Consequences returned quickly. If the baker made bad bread, he heard about it before noon. If someone needed help, the neighbors did not wait for an agency’s approval.

Local knowledge was rich, and rules could be adapted without months of negotiation.

"Elinor Ostrom studied communities that managed forests, irrigation systems, fisheries, and other common resources," I told him. "She showed that people can build robust local rules without everything being privatized or commanded by a center. But she never said that proximity alone solves the problem."

"What else was needed?"

"Clear boundaries, rules suited to local conditions, participation by those affected, monitoring, graduated sanctions, conflict-resolution mechanisms, and, for larger systems, nested levels of organization. Community did not work by magic. It worked through institutions."

"And reputation."

"Yes, but reputation alone can become revenge. Graduated sanctions mean you do not destroy a person for a first mistake. Conflict resolution means you do not leave a dispute to be judged solely by the stronger side."

Matei looked across the market. At one table, two producers were discussing the use of a shared cold room. One wanted to expand it; the other said the energy should be reserved for the medical store during peak periods.

"Who decides here?"

"The cooperative, under rules set by its members. But if the decision affects the region's food security or grid consumption, other levels become involved as well."

"So it isn't completely local."

"Nothing important is completely local or completely global. The question is which part belongs at which level."

I told him about subsidiarity: the idea that decisions should be made as close as possible to the people affected, but moved upward when lower levels cannot solve the problem. Matei said the rule was obvious, then spotted the difficulty himself.

"Who decides that the local level can't handle it?"

"Exactly. A central authority can declare too quickly that a problem exceeds the community. A community can insist too long that the effects stop at its border. That is why we need criteria, not just goodwill."

In his notebook I drew a small area with arrows pointing out of it.

“If the relevant effects stay inside, the decision can stay local. If the arrows cross the boundary, you have to involve the people who receive the consequences. The radius of the problem determines the radius of cooperation.”

“But people disagree about what ‘relevant’ means.”

“Correct. Science can measure some of the effects. Politics decides what weight to give them. No formula eliminates conflict completely.”

By lunchtime the market was crowded. A group of children was coming back from school, and several older people had taken their usual tables. Matei greeted almost everyone.

“See?” he said. “Knowing one another still helps.”

“A great deal. Trust lowers the cost of cooperation. You do not have to verify every gesture and write contracts for everything. But local trust comes in two forms. It can be a bridge: people help one another and include newcomers. Or it can be a wall: we help our own and treat everyone else as a threat.”

“How do you tell the difference?”

“Watch what happens to the stranger, the minority, and the person who wants to leave.”

In the past there had been much discussion about a possible natural limit on the number of stable social relationships a person could maintain. A figure of roughly one hundred and fifty became famous, although later research showed that it was not a precise constant. Social networks depended on culture, technology, time, and the type of relationship.

“There is no magic number at which people become a community,” I explained. “You can have a team of twenty where nobody trusts anyone and a city of a million where neighborhoods, professions, and associations form stable networks. Human scale is not just a number of people. It is the ability to understand the rules, have real relationships, and influence what is close to you.”

“Then large cities can have human scale?”

“If they are organized into many spaces of belonging and decision, yes. A large city also gives you something a village cannot always offer: the possibility of not being reduced to your past.”

I told him how cities allowed people to change their social circles, professions, and public identities. Anonymity could encourage indifference,

but it could also offer freedom. Someone judged for their family, faith, orientation, or an old mistake could begin again.

"So anonymity is good?"

"It is a resource, like light or fire. It can protect, conceal, and destroy. The question is where you need it and where you need accountability."

After lunch we walked to the Transition House, an unobtrusive institution near the market. People leaving highly closed communities could live there for a few months, find work, and rebuild documents, contacts, or skills. It was not a secret shelter, but the exact location of residents was not public.

The coordinator explained that many autonomous communities were peaceful and prosperous. Others had become dependent on a leader, a family, or a rigid interpretation of tradition. The problem was not cultural difference but the absence of exit and appeal.

"If a group wants to live by its own rules, why should anyone from outside intervene?" Matei asked.

"As long as adults can genuinely leave, children are protected, and the group does not impose serious costs on others, the space for autonomy should be broad," I said. "But the autonomy of the community cannot entirely cancel the autonomy of the person."

"And who sets the limit?"

"Basic rights, enforced by broader institutions that can themselves be challenged. We call them common floors. Communities can build many kinds of houses, but the floor cannot sink below protection of bodily integrity, freedom of conscience, access to justice, and the right to leave."

"And the ceilings?"

"Limits on the costs they can send outside. A community can choose how it produces energy, but it cannot pollute the region's water. It can have its own calendar and schools, but it cannot prevent its members from learning enough to exist outside it as well."

Matei wrote the phrase "common floors, freedom above them." Then he crossed it out.

"Sounds too nice. In reality everyone argues over how high the floor is."

"Very good. Keep that too. A serious rule has to include the place where the dispute begins."

On the way home we stopped on a pedestrian bridge. Below us, the river fed several channels: one for irrigation, one for a restored wetland, and one for the small hydroelectric plant. Every user had good reasons. Farmers wanted stability. Ecologists wanted natural flows. The city wanted energy and flood protection.

"Who should decide here?" I asked him.

"The people who live near the river."

"And those downstream?"

"Them too."

"The region that depends on the energy?"

"Yes."

"The biodiversity authority?"

"Probably."

"Then the decision is no longer small."

Matei leaned on the railing.

"Maybe the point isn't to make the institution as small as possible. You have to make it as large as the effects."

"That is a better beginning."

"But if the effects are global, we end up with a global center again."

"Not necessarily. You can coordinate a global function without moving every decision to the same center. A body has shared circulation, but it does not ask the liver to write memories or the lungs to decide where the legs go."

"And the body metaphor becomes dangerous if it assumes a single brain," he recited.

"You learn fast."

By the time we reached home, the solar warning had been updated again. The storm would probably be stronger than first estimated. The energy cooperative had asked volunteers to join the next day's test of neighborhood-islanding procedures.

Matei looked at the message.

"Our neighborhood will become a small community."

"For a few hours, yes."

"And if it can't produce enough energy?"

"It negotiates with the other islands. Some functions separate; others remain coordinated."

“So even the grid doesn’t choose between small and large.”

“Well-designed networks do not choose one scale. They can reorganize when the size of the problem changes.”

In his notebook, Matei wrote the day’s conclusion:

“A small community can see consequences better, but it can see too little of the world and too much of each person’s life.”

He paused, then added:

“Human scale does not mean knowing everyone. It means being able to understand the rule, challenge power, and leave without losing your life.”

This time he circled the sentence before asking whether he could defend it.

CHAPTER 4

The Art of Seams

The energy cooperative operated from a low building beside the old transformer station. It did not look like the city's nervous center. There were bicycles in the yard, toolboxes, and two cats asleep beside a warm pipe. Inside, one wall displayed the state of the local grid: solar generation, batteries, building consumption, connections to neighboring districts, and reserves for essential services.

The storm would arrive the following evening. The models could not predict exactly how intense the induced fields in electrical networks would be, so the operators were preparing several configurations. Some lines would be disconnected as a precaution. High-speed trains would run less frequently. The most vulnerable satellites had already been placed in protective mode.

Matei studied the diagram.

"If we disconnect the neighborhoods, don't we destroy the very advantage of a large grid?"

"For a while, yes. A large grid can move energy from places with a surplus to places with a shortage. But the same connection can carry a disturbance. Sometimes you preserve the whole precisely because you can divide it in a controlled way."

Sorin, the engineer on duty, gave us reflective vests and led us into the testing room. I had known him for years. He had started as an electrician, later studied control systems, and refused to let a model operate any piece of equipment he could not stop manually.

"Has your grandfather told you about seams?" Sorin asked Matei.

"Not yet."

"Then he will probably tell you they are the most important invention after the screw."

"I wouldn't go that far," I said.

"That is because you haven't spent enough time repairing systems glued into one piece."

Sorin opened a cabinet where the cables were arranged in labeled modules. Each could be isolated, tested, and replaced without shutting down the entire installation.

“A good seam is not a weakness,” he explained. “It is a place where we know how two parts meet and how to separate them if one starts behaving badly.”

“Like a fire door.”

“Exactly. The door lets you pass normally and stops the fire under abnormal conditions.”

We entered a room where volunteers were simulating different failures. A grid model generated scenarios, but important commands had to be confirmed by two people. Not because people were always better, but because models and operators failed in different ways.

In the first scenario, a regional line became overloaded. The system recommended isolating two districts and temporarily shutting down fast vehicle chargers. The operators accepted. In the second, the model recommended reducing power to an industrial area. Sorin refused.

“Why?” Matei asked. “Consumption is high there.”

“Because the water-treatment plant appears in the main registry as an industrial facility. The model sees the category; I know the process. If we cut its power without preparation, we do not merely lose production. We can damage the filters and interrupt the water supply.”

“Why isn’t it classified correctly?”

“The registry migration isn’t finished. Somebody should have updated something,” Sorin said, glancing at me.

Matei smiled. He had started to recognize the phrase.

“Then the model doesn’t know reality.”

“Neither do I, not completely. The difference is that we have different sources of knowledge and can check each other. If the model and the human always say the same thing, you do not have cognitive redundancy. You have one error speaking in two voices.”

We spent some time beside the operators. Every decision was recorded together with its reason, the data used, and the degree of confidence. After the event, the system would compare predictions with actual effects. It was not spectacular technology. It was a discipline of memory.

“Why record all the justifications?” Matei asked. “Don’t you waste time during a storm?”

“Some are generated automatically,” Sorin replied. “For rapid decisions we use short protocols. But if we do not preserve why we did something, after the crisis we will invent explanations that make us look smarter than we were. Without traces, we do not learn. We just tell stories about ourselves.”

I took Matei to the physical model of the grid. It consisted of small buildings, lines, and switches. You could connect all the districts into a single network or separate them into islands. Some had their own generation, while others depended more heavily on imports. The hospital, water center, and emergency communications were connected through redundant routes.

“Redundancy looks expensive,” he said.

“It is expensive,” I replied. “That is why systems eliminate it when they are evaluated only by efficiency on ordinary days.”

“But if you keep it everywhere, you waste resources.”

“Correct. You do not build twenty bridges over every stream. You identify which functions are critical, what types of failures are plausible, and what interruption would cost. Resilience does not mean duplicating everything. It means keeping enough alternative paths that one failure does not become a verdict.”

We talked about nearly decomposable systems, an idea Herbert Simon had identified in many complex structures. Within a subsystem, interactions are dense and fast. Between subsystems, connections are fewer and more controlled. A cell has membranes. An organism has organs. Well-designed software has modules. A resilient city has districts and services that can continue to function partly even when wider connections are disrupted.

“Does nearly decomposable mean you can almost take it completely apart?”

“It means the parts have enough autonomy to be understood and repaired locally, while remaining connected for shared functions. They are not independent. They are dependent in a disciplined way.”

“Like people in a good family,” Sorin said. “Close enough to help one another, separate enough not to share a toothbrush.”

Matei wrote down the comparison, probably because it was the first one that week that did not sound like an exam term.

In a side room stood an old distribution panel. Cables had been added in dozens of stages. Some crossed the housing without labels. Others supplied more functions than they had originally been intended to support.

"That is what many organizations looked like," I told him. "Every new problem added a process, a department, an exception, or an information system. After a few decades, nobody could remove anything without fearing that something else would stop."

"Why didn't they rebuild from scratch?"

"Because the system had to keep operating. People lived in houses, received salaries, took medicines. You cannot shut down a country for an update. Real complexity accumulates over time, on top of old choices, compromises, and pieces that still do something useful."

"Then how do you repair it?"

"Slowly. You define interfaces. Isolate functions. Replace hidden dependencies. Introduce boundaries. Sometimes you accept that one part will remain imperfect and build around it."

Sorin touched the old casing.

"Young engineers want clean systems. Old engineers want systems where you know what happens when cleanliness fails."

Then came the islanding test. The operators opened the main connection to the regional grid. On the screen, our district became an island. Some nonessential loads were reduced. The batteries covered the difference. The solar panels were producing little because of the clouds, and the system asked buildings to postpone heating water for two hours.

Matei checked his device.

"I didn't notice anything."

"That is the aim of a good test," Sorin said. "But do not confuse your comfort with an absence of cost. The bakery moved a baking cycle, the ceramics workshop shut down its kiln, and three buildings reduced ventilation in unoccupied spaces. Their flexibility protected you."

"Do they get paid?"

"Yes. Otherwise we would turn cooperation into an obligation imposed on those with less power. The grid buys flexibility; it does not demand it for free."

That sentence took the discussion in a direction Matei had not anticipated. Until then, he had seen resilience as a technical property: batteries, cables,

switches. Sorin showed him that it was also a distribution of costs. If a wealthy district protected its comfort by shifting interruptions toward an industrial area where poorer people worked, the grid could be technically robust and socially unfair.

"Then the algorithm has to optimize fairness too."

"It can include fairness criteria," Sorin said. "But who defines what is fair? Equal numbers of minutes without power? Protecting vulnerable people? Compensating those who give something up? Priority for essential services? The model can calculate the consequences of a definition. It cannot make that definition legitimate by itself."

"So people vote."

"Sometimes. Other times we use rules established through public deliberation, professional standards, and rights that are not put to a vote during every storm. If you decide everything in real time, a crisis favors whoever can shout the loudest."

After the test, we went up to the roof. From there we could see the solar panels, thermal reservoirs, and the link to the neighboring district. The city looked calm. Only the flows on the screen revealed how much coordination was required to produce that normality.

"The old grid tried to be seamless," I said. "One market, one control system, rapid transfers over long distances. It was extraordinarily efficient. But when a problem entered it, sometimes it found no closed doors."

"And now we have lots of small grids."

"We have a large grid capable of becoming many small grids and then becoming whole again. Permanent fragmentation is not the solution. The solution is the ability to change scale without chaos."

Matei looked at the cables running toward the other districts.

"Would that work for institutions too?"

"Yes. Teams with real autonomy, clear responsibility, and common interfaces. Cities that decide locally but share standards and infrastructure. States that retain their own powers but cooperate on cross-border problems. The important thing is for one part to be able to fail without its failure instantly becoming everyone's failure."

"But it also has to be able to ask for help."

“Exactly. A membrane is not a wall. It lets some things through and stops others.”

On our way out, Sorin showed us a red cabinet containing manual controls.

“If the automated systems contradict one another or we lose communications, we can operate locally,” he said.

“Isn’t it dangerous to let people intervene?” Matei asked. “They can make mistakes.”

“They can. That is why they train, work in pairs, and have limits. But a system nobody can stop locally is not autonomous. It is merely far away.”

By the time we left, the solar storm had been raised to the level at which preventive islanding was likely. The public display showed not only the forecast but its uncertainty. Some models predicted moderate effects. Others indicated higher risk for northern latitudes and long transmission networks.

Matei noticed that the system did not display a single prediction.

“Why doesn’t it combine them into one answer?”

“It combines them for some decisions, but it preserves the disagreement too. An average can hide the fact that one model sees a rare and dangerous possibility.”

“So models need seams too.”

“Especially models. If all of them are trained on the same data and depend on the same assumptions, ten predictions do not mean ten perspectives.”

At home, Matei drew the neighborhood grid as a group of islands connected by bridges. Beside the bridges he wrote what should pass through: energy, emergency signals, payments, aid. Then he wrote what should be stopped: overloads, control errors, attacks, and unverified commands.

Under the drawing he added:

“A resilient system is not one that never breaks. It is one that knows where it can break without losing everything.”

He read the sentence aloud.

“That sounds a little dramatic.”

“Then make it more precise.”

He crossed out “without losing everything” and wrote “without a local failure automatically becoming a general failure.”

“Less memorable,” he said.

“More useful.”

CHAPTER 5

The Machines That Helped Us See

The Civic Modeling Center occupied the upper floors of the library. The building had been designed before artificial agents became an ordinary part of government and had been modified several times. There were no enormous screens with a digital face at the entrance, and no name for "the intelligence of the city." Instead, the wall displayed a map of the systems in use, the organizations that administered them, and the rights citizens had in relation to them.

"How many artificial intelligences are here?" Matei asked.

"Depends on what you call an intelligence. There are models for energy, water, transport, public health, budgets, translation, forecasting, and deliberation. Some are small and specialized. Others coordinate agents. None is 'the mind of the city.'"

"Why not unify them?"

"For the same reason we do not fuse all the organs of the body into one enormous organ. The functions have different data, risks, and rules. A model that estimates water consumption does not need residents' medical histories."

"But it might discover useful connections."

"And it might create unnecessary power. The ability to combine data is both a benefit and a risk. Not every correlation deserves a right of access."

At reception we were met by Lia, the center's coordinator. I had sent her Matei's questions in advance. She had prepared an exercise for him: the city had to choose among three projects for reducing temperature during very hot summers. One model proposed rapid tree planting, another reflective surfaces, and a third a combination of green corridors, shade, and transport changes.

"Which solution is correct?" Matei asked.

"That is the first wrong question," Lia said. "Correct for what?"

Maps appeared on the table showing temperature, water consumption, biodiversity, costs, mobility, and income distribution. Tree planting offered

large benefits but required water and time. Reflective surfaces worked quickly, but performance depended on maintenance and could affect visual comfort. Green corridors improved several objectives at once but required reducing space for cars and making contested changes to urban planning.

"Can't the model calculate an optimal combination?" Matei asked.

"It can, if you tell it how to compare one degree of temperature, one hectare of habitat, ten minutes of travel, and the cost of removing parking spaces. But those weights are a social judgment, not a physical result."

"Then people choose the weights."

"Yes, after models help them understand the consequences. In the past, many systems hid value judgments inside the objective function. It looked as though the machine had found the answer, even though somebody had already decided what 'better' meant."

Lia showed him the provenance panel. For every prediction, you could see the data, source, age, uncertainty, and conditions under which the model had performed poorly. You could compare competing models and see where they reached similar conclusions for different reasons.

"Why show the flaws too?" Matei asked.

"Because mature trust does not mean believing a tool never fails. It means knowing how it fails and when not to use it."

He selected a new neighborhood built after the last major data update. The temperature map became less certain, and the system marked the area explicitly.

"Early commercial models often avoided messages like that," I said. "Products sold better when they looked certain."

"People actually trusted a model that didn't say how sure it was?"

"People trust a clean interface very easily. Uncertainty looks unprofessional. But reality does not become more certain when you remove the error bars."

Lia opened an archive of examples from the first decades of generative AI. Models could write, translate, program, summarize, and converse in ways that had seemed impossible only a few years before. They increased productivity in many activities and made skills accessible that had previously required specialists. At the same time, they invented sources, produced persuasive but wrong answers, amplified biases, and could conceal a lack of understanding beneath fluent language.

"Still, they became much better," Matei said.

“Much better,” I replied. “But progress in capability does not eliminate the question of role. An excellent autopilot does not also have to be everyone’s court, teacher, doctor, employer, and friend.”

“Why not, if it is better?”

“Because ‘better’ is not a general property. It may be better at prediction, worse at explaining an unusual case, faster, less sensitive to local context, or dependent on data controlled by a single organization. And even when it is more accurate, the question remains: who sets its objective, and who can challenge its use?”

I told him about research from the mid-2020s on human-AI teams. Public enthusiasm often assumed that a human and a model together would automatically outperform either one alone. Meta-analyses showed a more complicated picture. In many tasks, the combination helped relative to the human alone but did not necessarily beat the best-performing component. In some decisions, collaboration could even reduce performance.

“How can you get worse by adding something intelligent?” Matei asked.

“If the human trusts the model exactly when the model is wrong. If the human intervenes exactly when the model is right. If an explanation creates false confidence. If nobody knows who has the final word. Two good tools do not automatically make a good team.”

“Then you have to train the team.”

“Yes. The interface, the distribution of authority, the way uncertainty is presented, and the type of task are all part of the system. You cannot add a model to an old procedure and assume you have created hybrid intelligence.”

Lia took us to the Public Memory Room. The name was slightly grand for a room with tables and shelves, but the function mattered. There, artificial agents helped the administration track claims made in debates, assumptions behind decisions, minority warnings, and later outcomes.

On one screen was a transport-line project approved seven years earlier. The system could reconstruct what had been known at the time, which models had been used, what objections had existed, and what effects had actually appeared.

“Why not keep only the final decision?” Matei asked.

“Because the verdict erases the path,” Lia said. “When conditions change, we need to know not only what we decided but why. Sometimes a policy was good only under an assumption that is no longer true.”

She searched for an objection raised by a small group of residents. They had warned that the new line would raise rents around the stations and push some residents toward the outskirts. The prediction had been considered possible but secondary. Seven years later, the system showed that the effect had appeared in two areas.

“Did the AI discover that those people had been right?”

“It helped connect the warning to the new data. People interpreted the result and changed housing policy. The machine did not become the judge of the past. It kept the past from losing its witnesses.”

To my mind, that was one of the most valuable functions of artificial intelligence. Institutions had never forgotten for lack of archives. They forgot because relevant information failed to return at the right moment. A document could exist on a server and still be absent from a decision.

“So AI is a kind of memory that knows when to remember,” Matei said.

“Yes, if it is designed well. But memory can be manipulated too. Who decides what enters the archive? Which sources count as credible? How long do we keep the data? How do we protect people from a memory that never forgets personal mistakes?”

“You can’t keep everything.”

“Nor should we. A society needs institutional memory and personal forgetting. If every mistake you make at sixteen remains available to every employer when you are forty, transparency becomes permanent punishment.”

Lia then opened the contradiction module. For every proposal, the center used several agents with different roles. One searched for missing data. Another identified groups that might bear costs. Another compared the proposal with historical experience. One tried to construct the strongest argument against the preferred solution.

“Don’t they get lost in discussion?” Matei asked.

“Sometimes. That is why we do not use maximum contradiction for every decision. If a pipe needs repairing, we do not convene a parliament of models. The intensity of scrutiny should match the stakes and the irreversibility.”

“Irreversibility?”

"How hard it is to repair the result after you are wrong. For an easily reversible decision, you can experiment quickly. For one that changes an ecosystem or closes off access for a generation, you demand more evidence and more deliberation."

I added a new term to his notebook: correctability. A good system was not merely accurate. It was designed so that it could discover and repair its own mistakes without disproportionate cost.

"Then the best AI isn't one that never makes mistakes," Matei said. "It's one that recognizes quickly when it has made one."

"And that does not have the power to hide the mistake," Lia added.

We reached the part I had been most eager to show him: the architecture of civic agents. Instead of one central system with access to all data, the city used separate services. Medical data stayed in health infrastructure. Transport models received only the information they needed, often in aggregated form. For analyses requiring correlation, computations could be performed without moving all the data, and access was logged and limited in time.

"That sounds more complicated than putting everything in one place," Matei said.

"It is more complicated," Lia admitted. "Centralization is often efficient. That is precisely why you need a serious reason before concentrating power. A perfect database is also a perfect target."

"But what if security is good?"

"Security reduces risk. It does not eliminate formally legitimate abuse. The most dangerous access is not always theft. It may be the legal use of data for a purpose society did not anticipate when it collected them."

In my era, many organizations accumulated information because storage was cheap and the data might become useful. Later, models found new uses for it. Data collected for convenience became tools of evaluation. Data collected for safety became tools of control. Each use could have a justification. Together, they changed the relationship between person and institution.

"That is why we have purpose limitations," I told him. "A system must explain not only what data it uses but what for. A technical capability does not automatically become permission."

Matei sat at one of the consoles and asked the agents to compare the three urban-cooling projects. Instead of a single ranking, the system produced three

sets of recommendations, each beginning from different priorities. It highlighted areas where all models agreed and areas where the outcome depended on values or uncertain data.

"It isn't giving me the answer," he said.

"It is giving you a better-formulated problem," Lia replied. "Sometimes that is the real progress."

"But people want answers."

"Yes. And that is exactly why a system that speaks fluently can receive more authority than it deserves. A clear answer produces relief. A map of uncertainty produces work."

Before we left, Lia showed Matei the center's legal kill switch. It was not a physical button but a set of rules under which a model could be suspended if systematic errors appeared, if its sources became invalid, or if its use exceeded the approved mandate. Suspension did not depend only on the technical team. It could be requested by an auditor, a court, the civic commission, or, in some cases, triggered automatically by risk indicators.

"So even the creators don't fully control it."

"Correct. A public system should not depend on the goodness of its creators. It needs limits that survive changes in personnel."

On the way home, Matei asked why my generation had been so enthusiastic about artificial intelligence if the risks were visible.

"Because it did astonishing things. It helped us work, research, communicate, and discover. And because our problems had become so large that any tool promising a wider view seemed inevitable."

"Was it a mistake?"

"It would have been a mistake to reject the tool. The mistake was believing that more computing power automatically solved the absence of legitimacy, conflicts of value, and concentrations of power."

"What do you call the good role for AI now?"

"A cognitive exoskeleton. It enlarges what we can observe, compare, and preserve, without pretending to become the person or the community."

"An exoskeleton can still move your body."

"Exactly why it needs controls, limits, and the ability to be removed."

That evening, the storm forecast became more worrying. There was no reason for panic, but there was a real chance the regional grid would separate

for several hours. The civic center published, at the same time, the predictions of four models, the arguments for the proposed measures, and the conditions that would trigger additional restrictions.

Matei read everything, then opened his notebook.

“AI should not be the consciousness of the world,” he wrote. “It can be the instrument through which many communities see their own consequences more clearly and explain them to others.”

Then he added:

“An answer is not good merely because it comes quickly and sounds certain.”

“Did you learn that at the center?” I asked.

“No. I learned that from you.”

I was not sure whether it was a compliment.

CHAPTER 6

A Mind Made of Disagreements

The water meeting took place on Thursday afternoon in the regional deliberation hall. The river basin crossed four towns, two agricultural zones, a reserve, and an industrial region. The droughts of recent years had not caused a serious crisis, but flows had become more variable. The council had to decide how to allocate reserves during very dry periods and what investments to make before the next decade.

Matei received an observer badge and instructions not to intervene. He immediately asked whether he could submit questions through the room's system. The secretary said yes, but only after he first tried to formulate the strongest argument against his own question.

"Why?" he whispered.

"Because it is easy to ask other people to defend their position and harder to show that you have understood it," I replied.

Around the table sat farmers, city representatives, engineers, doctors, ecologists, industrial operators, and citizens selected by lot. They did not have the same formal power at every stage. Technical data were validated by experts, but objectives and trade-offs were not left to experts alone. Every participant could consult several analytical agents, but the room prohibited secret personalized recommendations. If a model influenced the deliberation, its argument had to be visible and open to criticism by the others.

The first presentation came from the hydrology team. The models showed a higher probability of summers with low river flows, but the uncertainty ranges remained wide. Three strategies followed. The first expanded reservoirs and transfer infrastructure. The second invested heavily in reducing consumption and reuse. The third combined moderate measures with flexible rules adjusted each year.

After the presentation, every participant answered several questions individually before the common discussion began.

"Why don't they just talk?" Matei asked.

"To preserve some independent thinking. If the first speakers set the tone, others can end up repeating the dominant ideas even without explicit pressure."

Research on collective intelligence had shown that groups can have a general capacity for solving tasks that is not simply reducible to the intelligence of their best member. The distribution of participation mattered, as did social sensitivity and the way information circulated. But constant connection also had costs. When people saw other people's solutions too early, they explored fewer alternatives. Sometimes alternating between separate work and collaboration produced better results than uninterrupted communication.

"So in order to think together, they have to be apart for a while," Matei said.

"Yes. A permanently connected crowd can arrive quickly at an answer and slowly at the idea that the answer is wrong."

The discussion began with agriculture. Farmers' representatives argued that reduced reserves would damage investment and food security. Ecologists showed that minimum flows were necessary for the ecosystem and water quality. Cities could reduce consumption, but not immediately and not without cost. Industry offered funding for infrastructure in exchange for higher priority during dry periods.

At first, every position seemed reasonable within its own frame. Farmers spoke about production. Ecologists about biological thresholds. Cities about public services. Industry about jobs and contracts. The conflict came primarily not from a lack of data, but from the fact that the same data entered different stories.

The deliberation system projected an argument map on the wall. It did not merely summarize conclusions. It showed premises, sources, and points of disagreement. When two sides used the same expression in different senses, the system marked the ambiguity. When an argument depended on an uncertain prediction, it displayed the range. When someone repeated a claim that had already been answered, it did not delete the claim but linked it to the earlier response.

"Is that collective intelligence?" Matei asked.

"It is part of the infrastructure. Collective intelligence does not appear just because you put people in the same room. Relevant information has to be able

to enter, differences must not be erased, and the group has to be able to observe its own errors."

After the first hour, the mediation agent generated a common formulation: all participants accepted that no sector should be protected by transferring unlimited risk to the others and that the minimum ecological reserve could not be treated as an ordinary consumer.

Matei read the text.

"Looks like they agree."

"They agree with a general sentence. We do not yet know whether they mean the same thing by it."

That was exactly what the next stage tested. The system presented four concrete scenarios and asked every participant to apply the principle. Consensus immediately shrank. Some believed a high-value agricultural crop deserved priority. Others argued that production could be moved, whereas a lost ecosystem could not easily be restored. The cities wanted to protect basic consumption but did not define "basic" in the same way.

"So the common formula was false," Matei said.

"No. It was incomplete. Sometimes a general agreement is the first bridge. Other times it is a curtain hiding the conflict. That is why it has to be tested on cases."

In the first decades of AI-mediated deliberation, researchers had shown that language models could help groups with divergent views formulate common statements and include minority perspectives more successfully. The result had been promising, but it had not turned consensus into an automatic property. A text acceptable to everyone could be useful. It could also package real differences in elegant ambiguity.

"A good mediator should not maximize agreement," I said. "It should maximize clarity about where agreement exists and where it does not."

"Then why call it a mediator?"

"Because it helps people move from misunderstanding to a well-formulated disagreement. Sometimes that is major progress."

In the second part of the meeting, a contradiction agent was assigned to attack the strategy leading in the preliminary preferences: the flexible combination of moderate investments. It showed that flexibility could become a polite name for postponement. If every year was treated separately, the region

could avoid expensive investments until the remaining options became much more costly.

Then it attacked the large-infrastructure strategy. It showed the risk of locking the region into an expensive system built on uncertain forecasts and of creating new demand precisely because more capacity became available.

Finally, it attacked the consumption-reduction strategy. It showed that the burden could fall disproportionately on small farmers and households that already consumed little, while wealthier actors bought efficient technologies and preserved their comfort.

"Doesn't the agent have any solution?" Matei asked.

"That is not its role. It is the organ of contradiction."

"Sounds annoying."

"A system without a critical function is pleasant until the day reality arrives from outside."

Beside us, a woman selected by lot raised an issue the models had treated superficially. Her daughter worked in a care center, and water restrictions during previous summers had affected cooling in the building and hygiene. The aggregated data had classified the center with other commercial services. Nobody had intended to reduce protection for vulnerable people. The category had simply been too broad.

The system recorded the case and checked how many similar institutions existed. It did not automatically change the policy, but it showed that the original assumption about "commercial consumption" needed revision.

Matei wrote a question: "Why didn't the experts know?" Then he formulated the counterargument himself: "They cannot anticipate every case, which is why deliberation exists." He submitted both sentences.

The moderator answered publicly. Experts were not useless because they had missed one case. They had built the models that made the discussion possible. But technical expertise did not replace the situated knowledge of people who lived with the effects.

"Does that mean every personal experience is worth as much as the data?" Matei asked later.

"No. A story can reveal a missing category, but by itself it cannot tell you how frequent the problem is. Data can show frequency while missing the mechanism. You need both, and you need to know what each can support."

During the break, participants received a report on their own process. The system showed who had spoken, which perspectives had been cited, where influence was concentrated, and how much each position had changed. The report was not used to punish people; it was used to make the group's dynamics visible.

One industrial representative had spoken far more than the others. When he saw the data, he voluntarily gave up some of his remaining time. A farmer who had spoken little had nevertheless had her arguments cited by several people. Her influence was greater than the duration of her interventions suggested.

"Why do you measure that?" Matei asked the moderator.

"Because groups easily confuse fluency with competence and repetition with importance. But perfectly equal speaking time does not guarantee a good discussion either. Sometimes someone has critical information and needs to speak more. The report does not decide. It helps us observe ourselves."

That was collective metacognition: a group's capacity to know something about the way it knows. Not merely what conclusion it had reached, but how certain it was, which sources were missing, which assumptions all the models shared, and which participants had not been heard.

A mature individual can say, "Here I'm confident. Here I'm assuming. Here fear is influencing me. Here I don't have enough data." Old institutions rarely spoke like that. They published a number, a verdict, or a plan. Uncertainty was hidden for fear of weakening authority.

"But if you keep saying you aren't sure, nobody trusts you anymore," Matei observed.

"Correct. Metacognition is not an excuse for indecision. You have to connect uncertainty to action. You can say: we have a sixty percent probability, but the cost of waiting is high, so we act. Or: we do not know enough, and the decision is irreversible, so we test at small scale."

"So the level of certainty doesn't decide by itself."

"Never. The stakes, the time available, and the possibility of correction matter too."

In the final stage, the participants built a mixed strategy. They accepted a minimum ecological reserve, special protections for essential services, investment in reuse, and a mechanism under which sectors seeking additional

guarantees financed part of the required capacity. They also set a threshold at which the plan would have to be reviewed rather than waiting for it to expire.

Not everyone was satisfied. Two farmers voted against it. An industrial representative abstained. The system preserved their objections together with the conditions under which they believed the strategy would fail.

"Why keep the arguments of the people who lost?" Matei asked.

"Because a vote decides what we do now, not what becomes true," the moderator replied. "If, three years from now, the conditions they predicted appear, we want the warning to return to the discussion."

"So the minority doesn't get the decision, but it gets memory."

"And the right to request review under criteria established in advance."

When we left, Matei looked more tired than he had after the museum.

"It would have been simpler to let a model decide."

"Of course."

"And maybe it would have chosen the same strategy."

"Possibly."

"Then why spend all these hours?"

"Because people did not merely have to find a technical plan. They had to accept how the risks were distributed, understand the reasons, and know how the plan could be revised. Legitimacy is not time wasted around an optimization problem. It is one of the conditions under which the solution can work."

"But what if people choose badly?"

"It happens. Deliberation does not guarantee wisdom. Neither expertise, markets, nor AI guarantees it. Good architecture combines different ways of detecting error and avoids giving any one of them a monopoly on reality."

On the way home, the afternoon light had an unusual color. Solar activity was already affecting high-frequency communications at northern latitudes. The observatory had updated the storm's arrival window. Our city would disconnect from the continental grid as a precaution the following evening.

Matei opened his notebook on the tram.

"Collective consciousness does not mean everyone thinks the same way," he wrote. "It means important information, including disagreement, can reach the decision and return when reality confirms it."

Then he added a shorter sentence:

“An intelligent group preserves the voices that lost.”
“That one is worth keeping,” I said.
“Because it is memorable or because I can defend it?”
“This time, both.”

CHAPTER 7

Problems with a Planetary Address

On the morning before the storm, the Planetary Observatory opened public access to its regional room. Residents did not need to follow every technical detail, but the institution had a habit of showing how warnings capable of temporarily changing the lives of millions were constructed.

The room occupied the old planetarium. The dome where I had watched projected constellations as a child was now used to represent Earth together with the systems surrounding it. The magnetic field appeared as a transparent shape. Streams of particles from the Sun were rendered in near-real time. Around the edges were satellites, long electrical grids, pipelines, and regions where the ground could amplify geomagnetic effects.

"Who runs the observatory?" Matei asked.

"No one in the sense of a planetary government," Ana, the astronomer on duty, replied. "It is a federation of agencies, universities, technical networks, and public centers. We have common protocols and reporting duties. Some decisions are local, some regional, and a few have to be coordinated globally."

"Why isn't there one organization? It would be simpler."

"Simpler for some things and riskier for others. A single center can standardize data and react quickly. It can also be wrong for the whole planet, be captured, or become a single target. We use an architecture in which observations check one another, while critical warnings have common channels."

The solar eruption appeared on the dome. It was not an explosion heading straight toward us as in a film. It was a cloud of plasma and magnetic field whose orientation would determine how strongly it interacted with Earth's magnetosphere. Probes upstream could provide a more precise warning only a relatively short time before impact.

"So we don't know exactly what will happen," Matei said.

"We know enough to prepare and not enough to guarantee the outcome," Ana replied. "That is the normal condition of serious decisions."

The observatory did not directly command every electricity network. It transmitted data, scenarios, and thresholds. Regional operators chose configurations suited to their infrastructure. Satellites had other protocols. Aviation, communications, and emergency services used the warning in different ways.

"That is global observation and distributed execution," Matei said, pleased to recognize the formula.

"Yes," I replied. "But don't turn it into an incantation. Sometimes execution itself has to be coordinated globally. Other times observation can remain local. The right scale depends on the mechanism of the problem."

Ana took us to the section devoted to planetary risks. It was not a gallery of apocalypses. Each risk was presented through its chain of observation, decision, and action. Climate, epidemics, near-Earth objects, orbital infrastructure, and certain technological risks had different architectures.

At the climate station, Earth was represented through interconnected processes: temperature, oceans, biogeochemical cycles, biodiversity, freshwater, and land use. I explained to Matei the planetary boundaries framework, proposed early in the century to describe Earth-system processes that could be pushed outside relatively safe ranges. The 2023 update estimated that six of the nine assessed boundaries had been exceeded.

"Does that mean the planet was doomed?" he asked.

"No. The boundaries were not a clock counting down to the end of the world. They were risk thresholds and orientation tools, debated and criticized scientifically. The important message was that some changes could not be judged separately. A system can look stable locally while total pressure on the planet keeps increasing."

"But emissions come from specific places. Factories, homes, cars."

"Exactly. The problem was global through accumulation, while solutions had to be implemented in real places. No planetary agreement insulates a house by itself, and no city stabilizes the atmosphere by itself."

On the screen I selected an emissions-reduction scenario. Global coordination set measurement standards, shared goals, and rules against shifting pollution from one country to another. Regions chose combinations

of energy, transport, and agriculture suited to their conditions. Cities changed buildings and mobility. Households and firms made concrete decisions.

"So every level does a part."

"Yes. When a problem is global, it does not mean all actions have to be centralized. It means the effects have to be accounted for at the scale at which they meet."

"And if one place refuses?"

"That is where difficult politics begins. You need incentives, trade rules, support for transition, and sometimes sanctions. Voluntary coordination does not always solve games in which everyone benefits if the others contribute while they do not."

Matei knew the prisoner's dilemma from strategy games. He observed that at planetary scale the problem was worse because participants did not have the same power, the same history, or the same transition costs.

"Then an identical rule for everyone can be unfair."

"Yes. Equality of formula can conceal inequality of starting point. But differentiation without criteria can become an excuse for inaction. That is why global institutions need not only objectives but credible mechanisms for distributing costs."

At the epidemics station, a map showed how a local signal could become globally relevant. A clinic noticed an unusual cluster of symptoms. Laboratories analyzed the pathogen. National authorities reported the risk, and international networks compared the information with other cases. If the signal was real, resources and recommendations flowed back.

"Why would anyone hide an epidemic?" Matei asked.

"Because reporting could trigger travel restrictions, economic losses, and political blame. If the system punishes the first place that notices the problem, it teaches everyone to notice later."

"Then reporting should be rewarded."

"And the affected place should receive support. Global cooperation does not work if the benefit of information is shared while the cost of honesty remains local."

The International Health Regulations had tried to create common obligations for detection and notification. Yet no rule could replace local

capacity. A global system without clinics, laboratories, and public trust was a radar without sensors.

"The signal has to travel faster than the disease," Matei said.

"And help has to travel fast enough to make the signal worth sending."

The next section was the children's favorite: planetary defense. The orbits of near-Earth objects appeared on the dome. Most posed no danger. The system showed how telescopes in several countries combined observations to estimate trajectories.

Ana projected images from the 2022 DART mission, when a spacecraft struck the asteroid moonlet Dimorphos and changed its orbital period. It had been the first full-scale test of a kinetic-impact deflection method.

"There is no local solution here," Matei said.

"No. A large asteroid does not become a municipal problem when it enters the atmosphere. Detection, calculation, and any eventual response have to be coordinated at planetary scale."

"Then why not give the Observatory more power over all risks?"

Ana changed the image and displayed the institution's mandate. It was narrow: observation, data standards, warning, coordination of certain responses, and public audit. It did not administer health, energy, or education.

"The fact that an institution is necessary for one function does not make it competent for every function," she said. "A fire department has the right to enter a burning building. That does not give it the right to choose the residents' furniture."

"So global power has to be very specialized."

"The harder it is to challenge, the clearer its mandate should be. Sometimes we need great power for a narrow purpose and a limited time."

I explained to Matei the idea of a globality budget. It was not a financial budget but a discipline: for every competence moved to planetary scale, the institution had to show that the problem could not be handled adequately at lower levels, that global intervention made a real contribution, that the mandate was limited, and that audit and review existed.

"Who checks all that?"

"Other centers. Courts, parliaments, scientific communities, civic organizations, states, and independent auditors. There is no final point that checks itself perfectly."

"Then the system can get stuck."

"It can. Distributed control has costs. Centralization accelerates some decisions. Polycentrism reduces the risk of a single error and permits different perspectives, but it can create delays and conflicts of jurisdiction."

"So there's no cost-free solution here either."

"That seems to disappoint you every day."

"Because adults call something a 'principle' and then attach three pages of exceptions."

"A principle without exceptions in a complex world is often just a preference that has not met reality yet."

We continued to the orbital-infrastructure section. Satellites provided communications, navigation, climate observation, and economic services. Orbits had become common resources vulnerable to debris and crowding. No country could maintain order in near-Earth space by itself, yet there was no simple world authority deciding every use.

Solutions had evolved through standards, shared registries, liability rules, insurance, and maneuver coordination. It was a good example of governance without a single sovereign. Not perfect, but capable of addressing a problem that crossed borders.

"So the planet can have organs without having one brain," Matei said.

"That is a useful formulation. One organ for observing climate, one for diseases, one for asteroids, each made up of many institutions and connected to local levels. Collective consciousness does not require a planetary person. It requires pathways by which important signals reach places capable of responding."

On the dome, the solar-storm simulation had updated. The orientation of the magnetic field looked less favorable. It was not a catastrophe, but operators recommended disconnecting long networks before the main front arrived. Our city would enter island mode at twenty-one hundred hours.

"Who made the decision?" Matei asked.

Ana displayed the chain. Observatories had provided the data. Regional models had estimated effects on infrastructure. Grid operators had made the recommendation. The energy authority had established the thresholds in advance, after public consultations and tests. Local cooperatives had confirmed that they could operate separately.

“Looks like there still isn’t one decision-maker.”

“There are distinct decisions and distinct responsibilities,” I said. “The difference from ghost agency is that the links are explicit, traces are preserved, and every level has a mandate. Distributed decision-making is not the problem. Opacity about how it is distributed is.”

As we left, Ana gave Matei a printed image of Earth seen from a distance. It had no political borders and no instructions for how the planet should be governed.

“Photographs of Earth from space made people talk about one world,” she said. “That was true and incomplete. We have one atmosphere but many communities. One planet, but not one perspective. Physical unity does not automatically produce political unity.”

On the way home, Matei wrote:

“A problem should be governed at the scale at which its consequences can no longer be avoided separately.”

Then he added:

“At the global level, raise the global function, not all of life.”

He looked at the picture of Earth.

“Do you think that could be my rule?”

“It is one half of it.”

“What is the other half?”

“What we do with the power that remains below.”

The storm would help us with that as well.

CHAPTER 8

The Night the City Detached

At twenty-one hundred hours, the city detached from the continental grid without going dark.

There was no explosion. The lights flickered once, the pumps changed source, and a small symbol appeared on our devices showing that power was now coming from the local island. For most people, the transition was less dramatic than a software update.

Matei looked out the window.

"Was that it?"

"That is what we hope."

Before the separation, the district had reduced loads that could be postponed. Batteries were charged, thermal reservoirs prepared, and public buildings had checked their plans. The hospital and water center had their own backup systems. Bakeries had shifted part of their production. Workshops with energy-intensive processes had stopped.

On the local channel, Sorin explained the state of the grid. He was not asking people to sit in darkness. He asked only that they avoid heavy consumption for the next two hours, until operators saw how the storm behaved.

"Why doesn't the system automatically reduce whatever isn't necessary?" Matei asked.

"It does where contracts allow. But it does not know what is necessary in every home. One person may not care whether the washing machine starts at midnight. Someone else may urgently need it. Good automation leaves room for context."

We went out toward the neighborhood center, where volunteers coordinated support for vulnerable residents. I was registered with the communications team, and Matei had permission to come with me. The streets were quieter than usual. Public transport was running, though less frequently. Illuminated advertisements had been switched off. Homes kept enough light that no one felt as though the city were under siege.

At the community center, one wall held a map of consumption and another a list of incidents. The first twenty minutes passed without problems. Then a message arrived from an apartment building: one of the elevators had stopped between floors.

The automated system had classified it as a minor technical incident. The building manager called and said a woman inside the elevator was using a portable oxygen concentrator.

The priority changed immediately.

"The model didn't know?" Matei asked.

"It knew she lived in the building, but not that she was in the elevator. No system has the complete state of the world."

The response team set out. Meanwhile, the medical center contacted her directly and checked the battery in her device. She was safe and was brought out after a few minutes.

The incident was small but well suited to our lesson. The technical data had been correct. For the grid, the elevator was a minor fault. Human context changed its importance.

"Then anyone can claim their problem is urgent," Matei said.

"Yes. That is why we do not eliminate assessment. We supplement it with information and judgment. Emergency systems have to be able to change a category when new context appears."

At twenty-one thirty, induced fields increased in the north of the continent. Two large networks disconnected long lines as a precaution. Some regions temporarily lost external connection, but most local islands remained stable. Information on our screen arrived with delays. Satellite communications were degraded, and terrestrial networks were using alternative routes.

Sorin sent a request: the neighboring district had a problem with one of its batteries and wanted a small energy transfer. Our system calculated that we could help without falling below the safety reserve, but recommended waiting ten minutes to confirm the storm's stability.

"What do you do?" Matei asked.

"This is where grid politics begins," I said.

Operators could refuse and protect the local reserve. They could accept immediately and increase our own risk. Or they could offer a limited amount, with an automatic right to interrupt the transfer if frequency dropped.

The cooperative chose the third option. The decision followed the mutual-aid protocol established before the crisis. It did not depend on generosity in the moment or on a regional command. The neighboring district had helped us in winter tests, but the rule would have applied even without that history.

"That is local cooperation," Matei said.

"And a regional institution. Without common standards, compatible meters, and an interruption contract, goodwill could not travel through the cable."

At twenty-two hundred hours, a different problem appeared at the food warehouse. The cooling system had received an order to reduce power, but one compartment held temperature-sensitive veterinary medicines. In the registry, the warehouse was listed only as food infrastructure. The manager noticed the change and requested an exemption.

The center approved it, then automatically opened a classification incident so the registry could be reviewed after the storm.

"Why not fix the database now?" Matei asked.

"Because right now the goal is to preserve the temperature. In a crisis, you do not confuse immediate correction with systemic reform. You do both, but at different times."

That was a lesson institutions had learned painfully. During emergencies, people discovered structural defects and sometimes tried to repair them on the spot. Rapid modifications could introduce new errors. Modern systems separated the operational loop from the learning loop. The first stabilized the situation. The second preserved traces and returned after the event.

"What if the second loop forgets?"

"Then the crisis becomes only a story. That is why the incident gets an owner, a deadline, and a closure criterion. 'We will investigate' is no longer considered an outcome."

Around twenty-three hundred hours, the aurora began. At our latitude it was rare, and under other circumstances people would have gone into the streets. The public center announced that the phenomenon was safe to watch but asked residents not to use vehicles unnecessarily. Matei persuaded me to step outside for two minutes.

Green and violet bands crossed the sky. They did not look like the perfect images in the archive. They moved irregularly and were sometimes barely visible. Neighbors watched in silence. A small child asked whether the Sun was trying to get into the city.

"No," Matei told him. "Our magnetic grid is just arguing with it."

We went back inside before I could correct the explanation.

At the center, the situation had become more complicated. One regional model estimated that the peak had passed. Another showed the possibility of a stronger second front. Averaging the two would have suggested a moderate reduction in risk, but the operators kept the scenarios separate.

"Why not choose the model that has predicted better so far?" Matei asked.

"Because previous performance matters, but the second model may represent a rare mechanism the first smooths away. In decisions with asymmetric costs, you do not ask only which prediction is most likely. You also ask what happens if you ignore the less likely one."

"Like taking an umbrella when there's only a small chance of rain, because carrying it costs little."

"Yes. Except here the umbrella means shutting down services and causing economic losses. Precaution has a cost too."

The energy council decided to maintain island mode for another two hours. Some industrial operators challenged the decision, arguing that the data did not justify the extension. Their objection was recorded and would be reviewed after the event. It could not block the measure during the risk window, but it did not disappear either.

"Is that unfair to them?" Matei asked.

"It may be. Sometimes you have to act before you can fully prove necessity. Legitimacy then comes from rules established in advance, time limits, published reasons, and later review."

"And if the review shows the decision was wrong?"

"Compensation is paid and protocols change. An authority should not be punished for every missed prediction if it followed a reasonable procedure. But neither can it repeat the same mistakes without consequences."

At midnight, the library became a charging center for medical and communication devices. Most people did not need it, but a few older buildings had weaker reserves. The bakery near the market brought bread and soup.

Nobody had asked them to. The local contract reimbursed part of the cost, but the owners had also come because they knew the people.

"That can't be completely programmed," Matei said.

"No. Nor should it be. A good system leaves room for local initiative and helps it coordinate. If every spontaneous act requires approval, you kill exactly the resource you need in unexpected situations."

"But initiative can get in the way."

"That is why there are contact points, minimum rules, and common information. Operational freedom without coordination produces chaos. Coordination without freedom produces paralysis."

Around one o'clock, a group of teenagers set out to check on older residents in two buildings. The system gave them routes and contacts but did not show medical data. In cases where a specific risk existed, the center transmitted only the relevant instruction.

"Why not tell them everything?" Matei asked.

"Because helping someone does not give you a right to their entire life. We use the minimum information necessary."

"And if the missing detail puts them in danger?"

"Then the system transmits it. Data minimization does not mean rigidity. It means every access must be justified."

Around two, the secondary front passed without the worst effects. Several northern regions had experienced outages, but modular systems had limited cascades. Communication with the continental grid stabilized. Operators began the gradual reconnection.

That was the most delicate part. Separation could be simple compared with reunion. Frequencies, loads, and protections had to be synchronized. A rushed reconnection could create the very problem islanding had prevented.

"Does that happen in society too?" Matei asked.

"What exactly?"

"It's easier to separate groups than to make them cooperate again."

"Very often. Fragmentation can limit a conflict, but it can also destroy trust and a common language. That is why seams have to be designed before rupture. If all you have are walls, reunion becomes a negotiation among strangers."

At a quarter past three, our district reconnected. The lights did not change. On the screen, the island symbol disappeared, and regional flows slowly returned.

Matei was tired but refused to leave before the short report. Sorin presented the incidents: an elevator, a misclassified warehouse, a battery in the neighboring district, two communication interruptions, and several buildings consuming more than forecast. None had been serious. All would be analyzed.

"It worked," Matei said.

"Yes," I replied.

"Then why analyze it?"

"Because success can hide errors just as well as failure. Maybe we were lucky. Maybe some protections were unnecessary. Maybe another district bore a cost that does not appear in our report. If we congratulate ourselves too quickly, the only thing we learn is that we like the outcome."

We walked home through almost empty streets. The aurora was fading. Behind the windows, people were returning to their routines. The storm would not become a legend. For most residents, it had been a night with a few notifications and an unusual sky.

That was the measure of success: a planetary risk had been observed globally, translated regionally, managed locally, and absorbed without requiring every person to understand solar physics or grid architecture.

"Now I understand why you don't want one scale," Matei said. "The Observatory saw the Sun. The region saw the big lines. The neighborhood knew about the elevator and the warehouse. None could replace the other levels."

"Yes."

"And AI connected the information, but people changed the categories when reality didn't fit inside them."

"Yes."

"Then my rule should say that every problem is solved at the smallest level that can see its effects, and larger levels intervene only for what goes beyond it."

"Almost."

"What is missing?"

“Seeing is not enough. You have to bear a relevant share of the consequences, be able to correct the decision, and not exclude the people it affects.”

Matei sighed.

“More conditions again.”

“The world has the inconvenient habit of not fitting into one sentence before it has been understood.”

When we got home, he sat at the kitchen table and drew three concentric circles: neighborhood, region, planet. Then he erased the circles and replaced them with a network of overlapping shapes.

“The circles make it look as though the levels are perfectly ordered,” he said. “But a hospital may depend directly on a global research network, and a neighborhood on a regional cooperative. It is more like a network of scales.”

“Good.”

Under the drawing he wrote:

“A conscious civilization is not one in which one center sees everything. It is one in which every important signal can reach the level capable of responding, without giving that level power over things that do not concern it.”

He reread the sentence and crossed out the word “conscious.”

“I’m still not sure what it means.”

“Then don’t use it until you are.”

CHAPTER 9

Matei's Map

On Friday morning, after only a few hours of sleep, Matei spread all the pages he had written during the week across the living-room floor. Some were orderly. Others contained arrows that began at a sentence and ended in an unexplained corner. On one sheet he had drawn the neighborhood as an electric hedgehog. On another, Earth was surrounded by nine institutions, three question marks, and an asteroid with an unhappy face.

"I have too many ideas," he said. "The assignment asks for one rule."

"Then don't begin with the rule. Begin with the problem."

In the middle of a large sheet he wrote: "The world produces effects nobody sees completely."

"That is true of any world," I said. "Even hunters ten thousand years ago did not know every effect of their actions."

"Then what is modern about it?"

"Speed, reach, and coupling. A decision made in an office could change the lives of millions before its effects were understood. A software error could cross continents. A product could depend on thousands of organizations. The problems were no longer merely unknown. They had become too large for the existing mechanisms of correction."

Matei changed the sentence:

"Our systems can produce consequences faster and farther than we can observe and correct them."

"That is the scale problem," he said.

"One of them. There is also the opposite problem: sometimes the institution is too small for the consequences it has to manage."

He added a second sentence:

"Some problems exceed the communities trying to solve them."

We began sorting the sheets around the two statements. On the left we put examples of systems too large for their own control mechanisms: opaque platforms, fragile chains, remote administrations, models that classified

without appeal. On the right we put problems too large for local solutions: climate, epidemics, asteroids, continental networks.

"Then the rule has to find the right size," Matei said.

"Yes, but 'size' can mean different things. Number of people, distance, budget, amount of data, power, or time. We need a better description."

I drew the first instrument on the map: the cognitive horizon. It was the part of the system an actor or institution could represent well enough to make a responsible decision.

"It does not have to understand every detail," I explained. "A doctor does not know the position of every molecule in the body. But the doctor needs a model suited to the clinical question. An energy operator does not know every device in every home, but needs to know enough about load, infrastructure, and vulnerabilities."

"And if AI enlarges the horizon?"

"Then the institution can manage more complexity. But the horizon does not grow merely by accumulating data. It also has to grow through good representations, verifiable sources, and the ability to recognize what is missing."

Matei drew a circle and wrote inside it: "What we can see well enough."

The second instrument was the externality radius: how far the relevant effects of a decision propagated through space, time, and social structure.

"If I decide how to arrange my room, the radius is small," Matei said.

"Usually. If you play music very loudly, it expands to the neighbors. If your device uses a lot of electricity during a shortage, it reaches the grid. No action has a fixed radius under all conditions."

"And if a factory pollutes the river, the radius follows the water."

"Yes. The institution making the decision has to include those within that radius or answer to a level that represents them."

He drew a second, larger circle with arrows pointing outward.

The third instrument was governance distance: the separation among those who set the rule, those who understand the mechanism, those who bear the effect, and those who can correct the decision.

"Can a global institution have a short governance distance?" Matei asked.

"For a well-defined function, yes. If you can see its mandate, its data, who is responsible, and the path of appeal. And a local institution can have a very long governance distance if it is captured by a group you cannot challenge."

"So 'close to people' doesn't necessarily mean geographically close."

"Correct. The relevant kind of closeness is the ability to hear, explain, and repair."

The fourth instrument was correction latency: the time between the appearance of a problem and the real possibility of changing it.

"Justice after twenty years can be legally correct and useless to a life destroyed in the meantime," I said. "A system that always recognizes the truth too late is not intelligent enough for the power it possesses."

Matei wrote beside the term: "How long until reality can change the rule?"

The fifth instrument was exit capacity. Not the formal right to leave, but the real cost of leaving: what you lose, what you can transfer, whether alternatives exist, and whether the system controls indispensable resources.

"A school cannot let students change every rule simply by leaving," Matei said.

"No. Children need continuity, and society has educational obligations. That is why, where exit is limited, voice, protection, and appeal must be stronger."

"So exit and voice compensate for each other."

"To a point. A good system tries to preserve both. If people can only leave, the community does not learn. If they can only protest but cannot leave, power may ignore the protest."

The sixth instrument was coupling: how quickly and how inevitably a change in one component is transmitted to others.

Matei pointed at the diagram of the energy grid.

"The connection helps us on normal days and can hurt us during a storm."

"Yes. Coupling is not good or bad in itself. It has to fit the function, and there have to be controlled ways to separate."

The last instrument was correctability, combining reversibility, audit, and the possibility of changing the system without destroying it.

"If a decision is easy to reverse, we can experiment faster," Matei said.

"Yes. If it is hard to reverse, we need more caution, greater diversity of models, and more participation. The wisdom of limits is not fear of action. It is matching speed to the cost of error."

When we finished, the sheet looked more like a machine plan than a school essay. Matei tried to turn the instruments into a formula. He put plus signs

beside externality radius, distance, latency, and coupling, and minus signs beside cognitive horizon, exit, and correctability.

"I can say risk rises when the first group is large and the others are small."

"As a heuristic, yes. Not as an exact law. You cannot add kilometers, months, and degrees of opacity into one universal number."

"Then why call them measures?"

"Because we can operationalize them in concrete contexts. We can measure appeal time, decision concentration, the proportion of effects outside a jurisdiction, dependence on one supplier, or the ability to move data. Science often begins with imperfect notions that become more precise through comparison."

"And wisdom?"

"Wisdom begins when you recognize that a precise measure does not decide by itself what is good."

He raised his eyebrows.

"That should be the title of your book, not my project."

I took another sheet and we tried to formulate the principle of appropriate scale. Matei's first version was:

"Decisions should be made at the smallest possible level."

We tested it against an asteroid. It failed.

The second version:

"Decisions should be made at the level that can solve the problem."

We tested it against a global government claiming it could efficiently administer the education of every child. It allowed too much.

The third version:

"A decision should be made at the smallest level capable of understanding and correcting the relevant consequences, including those who are affected; only effects that exceed that level should be moved upward."

Matei read it.

"It is correct and impossible to remember."

"First we make it correct. Then we make it short."

"What if the small level understands the problem but doesn't have the resources?"

"The larger level can provide resources without taking over all decisions. Capacity and authority do not have to be transferred together in every case."

"What if the effects are uncertain?"

"You use scenarios, monitoring, and escalation thresholds. Scale can change as the problem evolves."

"What if the people affected disagree?"

"Then the rule does not solve politics. It only tells you who has to be in the conversation and what kind of justification is required."

"What if a local level violates someone's rights?"

"Common floors allow the broader level to intervene. Autonomy does not include the right to abolish exit, appeal, or basic protection."

Matei drew a line under the principle.

"There is no rule that decides everything."

"If there were, we would need only executors, not citizens."

After lunch, we returned to the idea of collective consciousness. Matei had avoided the term after the night of the storm because it seemed vague to him. I asked him to describe what the city had done.

"It observed the storm through global networks. It compared models. It prepared the neighborhoods. It saw local incidents. It changed priorities when it learned the context. It preserved errors for analysis. It reconnected."

"Those resemble cognitive functions: perception, memory, integration, action, and correction."

"But the city doesn't feel anything."

"We do not know whether a collectivity can have subjective experience, and we do not need to assume that it can. We can speak more modestly about functional consciousness: the ability of a social system to represent its own state and effects well enough to correct itself."

"Then a thermostat is functionally conscious."

"If we define it too broadly, yes, and then the term stops helping. The city has something more: multiple perspectives, contested goals, historical memory, and the ability to change its own rules. But the term remains metaphorical. You can use 'reflexive collective intelligence' if you want greater precision."

"That sounds even worse for a twelve-year-old's project."

"Then say it simply: a society that can observe what it is doing and change its rules."

He liked that.

Then we discussed the role of artificial intelligence. In his original project, Matei had wanted models to calculate the answer while humans intervened only when they disagreed. Now he saw a different architecture.

AI could extend the cognitive horizon: combine data, compare scenarios, find connections, and preserve the memory of objections. It could reduce latency by detecting patterns of error faster. It could make the externality radius visible: who bore the costs and where. It could assist deliberation by clarifying agreements and disagreements. It could support modularity through specialized agents, data retained at the source, and verifiable interfaces.

But AI could also do the opposite. It could increase governance distance if people could no longer interrogate a decision. It could concentrate power through the accumulation of data. It could reduce exit capacity if the same platform controlled identity, reputation, and services. It could accelerate errors and manufacture artificial consensus.

"So AI does not have one political effect of its own," Matei said. "It depends on the architecture."

"It has some tendencies. The high cost of models can favor concentration. Ease of copying can favor distribution. Personalization can increase autonomy or manipulation. We do not start from a blank sheet, but institutional design matters enormously."

"Then the rule for AI would be: exoskeleton, not sovereign."

"That is a good formulation if you explain what it means. An exoskeleton has to be inspectable, limited, replaceable, and removable. You do not give it access to the whole body merely because it helps you lift a weight."

In the margin, Matei wrote five questions for any artificial system: What objective does it pursue? What data does it use? Who can stop it? Who can challenge the outcome? Who bears the cost of error?

"Those are better questions than asking whether AI is good or bad," he said.

"Much better."

Toward evening, we tried to find examples that contradicted the theory. A global library looked very large but could pose little risk if its function was limited, its data public, and exit easy. A family was small but could have a vast governance distance for a child without external protection. A decentralized market could distribute decision-making while concentrating outcomes in a

few platforms. A centralized state could administer vaccination efficiently but serve the needs of one classroom badly.

"So scale is a property of the relationship between a problem and an institution, not of the institution alone," Matei concluded.

"Exactly. There are no generally good or bad large organizations. There are functions for which they are well suited and functions for which they become blind."

"And the boundaries have to move."

"Some of them. A good boundary is not necessarily a fixed wall. It can be an escalation threshold, temporary permission, a separation that activates in a crisis, or an obligation to review."

"Then how do you stop temporary power from becoming permanent?"

"Expiration dates, public criteria, different authorities for activation and extension, audits, and the ability of courts to intervene. No solution is perfect, but power should be forced to ask for permission again."

Matei looked at all the sheets.

"I think I know what the wisdom of limits means. It doesn't mean being modest and refusing to do large things."

"No."

"It means building power together with the way you stop it, divide it, and correct it."

"Yes. Science helps us see cognitive limits, networks, externalities, and thresholds. Wisdom appears when we accept that power and knowledge do not always grow together."

He began writing the final version of his rule. After several attempts, he settled on two sentences:

"Solve each problem at the nearest level that can understand, bear, and correct its consequences. Move upward only the part that goes beyond that level, and the power moved upward must remain limited to that part."

Then he added a third sentence, even though the assignment had asked for a single rule:

"Use artificial intelligence to make effects more visible and decisions more correctable, not to hide politics behind a calculated answer."

"The teacher will say that's two rules," he said.

“Tell her the second is the technical condition for the first in a world too complex for human memory.”

“That sounds like an attempt to win an argument by making the sentence longer.”

“You’ve learned to recognize grandfather strategies too.”

Before bed, he gathered the papers and kept only the large map. At its center there was no planetary brain and no pyramid. There were several circles and nodes connected by arrows. Some arrows carried information. Others carried resources, rights of appeal, or obligations. Beside every connection he had written who could stop it.

“Why did you put a stop mechanism everywhere?” I asked.

“Because a connection nobody can stop is not just a link. It is a dependency.”

It was the best sentence he had said all week. I did not ask him to make it more precise.

EPILOGUE

The Answer That Fit on One Sheet

On Monday morning, Matei carried his map to school in a cardboard tube. His teacher had invited parents and grandparents to watch the presentations but asked us to sit at the back and not intervene. It was a wise precaution. Adults have a tendency to turn children's projects into referendums on their own intelligence.

Matei's classmates had chosen different principles. Some argued for the right of future generations to be represented in present decisions. Others spoke about universal access to education, protection of cultural diversity, or the requirement that any critical technology be able to operate in a degraded mode. One girl proposed that no important social indicator should be published without showing the distribution behind the average. I wrote the idea down.

When his turn came, Matei unrolled the map and began without the introduction he had rehearsed at home.

"At first I wanted artificial intelligences to calculate what was best and people to vote only when the models disagreed. Now I think that would move the political question into a place that is harder to see. Models can calculate consequences, but someone still has to choose what 'better' means, which costs we accept, and who has the right to challenge the decision."

He pointed to the left side of the map.

"Some systems are too large for their mechanisms of observation and correction. That does not mean they have too many people. It means their effects travel faster or farther than accountability can reach. Other problems are too large for local institutions: climate, epidemics, continental grids, and asteroids. So we cannot simply say 'everything local' or 'everything global.'"

Then he read his rule:

"Solve each problem at the nearest level that can understand, bear, and correct its consequences. Move upward only the part that goes beyond that level, and the power moved upward must remain limited to that part."

A classmate asked what "nearest" meant.

"Not only geographically," Matei replied. "Near means the people affected can find out how the decision was made, can participate or appeal, and there is someone who can change it. A local institution can be very far away if one family controls it and nobody can leave. A global institution can be relatively near for a narrow function if its mandate and data are public."

The teacher asked how he knew when an effect had to be moved to another level.

"I don't always know. You have to measure the externality radius and watch thresholds. Scale can change. An epidemic can start locally and become global. A large grid can separate into islands during a storm. The rule does not automatically produce the decision. It tells us which questions to ask."

Another student raised the issue of efficiency. If every level had rights of appeal and several models checked a decision, would the system not become too slow?

"Yes, sometimes. That is why scrutiny has to be proportional to the stakes and to how hard the decision is to repair. You do not organize planetary deliberation over a broken pipe. But you do not treat an irreversible decision like a quick test either."

At the end, the science teacher asked him to explain the role of AI in a single sentence.

Matei looked at the map.

"To enlarge the part of the world we can see and correct, without automatically receiving the right to govern the part it can see."

That was not the formulation we had prepared. It was better.

After the presentation, his classmates stuck questions around the edge of the map. Who protected minorities when the local community decided? How did you stop global centers from expanding their mandates? What happened when all the models were wrong in the same way? Who represented animals, ecosystems, or people not yet born? What did you do when a crisis moved too fast for participation?

Matei did not answer all of them. He grouped the questions and said some showed where the rule needed additional institutions, while others showed where there was not yet a good answer.

His teacher approved the project but asked him to keep the questions in the final version.

On the way home, Matei was dissatisfied.

"I thought that after a week I would have a solution."

"You have a partial solution and a map of where it can fail. That is more than many governments have before they start a program."

"But it doesn't solve conflicts."

"No architecture eliminates the fact that people have different interests and values. It can only make power more visible, costs harder to hide, and mistakes easier to correct."

We passed the market. Activity had returned to normal after the storm. At the energy cooperative, technicians were analyzing the battery that had failed. At the civic center, decisions from Friday night were being compared with the real data. The Observatory had lowered the alert level. None of this meant the world had become simple.

It meant only that it had learned not to ask a single institution to see everything, and not to let any level pretend its effects stopped at its own border.

Matei opened the tube and looked at the map once more.

"I'm going to change one more thing," he said.

"What?"

He pointed to a corner where he had drawn a small community connected to the city, the region, and global networks.

"I added arrows for data, energy, and appeal. I forgot arrows for learning. The levels shouldn't only control one another. They should be able to borrow solutions from one another without copying everything."

He added two dotted arrows.

The book could stop here with a grand conclusion. It could say that humanity had finally become conscious of itself or had found the balance between freedom and coordination. That would not be true.

In 2094, there were still opaque institutions, abusive communities, overpowered centers, and problems that even the best models could not solve. People still confused speed with competence and familiarity with truth. Some limits were used for protection, others to preserve privilege. No generation receives a final architecture.

But we had one advantage we had not always possessed: we had begun treating scale as a design problem rather than an inevitable consequence of growth.

We knew to ask how far an effect traveled, who could see it, who bore it, who could stop the mechanism, and how quickly it could be repaired. We did not always answer well. But the questions arrived less often after the disaster.

Matei rolled up the map and handed me the tube to carry.

"Do you think the rule will still be good in a hundred years?"

"Not in the same form."

"Then why use it?"

"So that we can build systems that allow your grandson to change it without having to tear down the world first."

Scientific Note

This book is speculative fiction, but it does not begin from an arbitrarily invented world. Its theoretical framework combines findings from systems science, institutional economics, cognitive psychology, organization theory, network research, and recent studies of human-AI collaboration. The future in the story is not a prediction. It is a narrative laboratory in which possible institutions and their consequences can be examined.

The starting point is Herbert A. Simon's concept of bounded rationality. Real actors do not possess complete information, unlimited time, or infinite computational capacity. They use partial representations and seek satisfactory solutions rather than absolute optima. This limitation is not an accidental imperfection but a condition of any feasible decision. Simon also showed that complex systems are often hierarchical and "nearly decomposable": interactions within subsystems are more intense than interactions between them. The idea does not require isolating components; it means organizing complexity into relatively stable units connected through intelligible interfaces.

The cybernetic concepts of W. Ross Ashby and Roger Conant provide a second foundation. The law of requisite variety suggests that a regulator needs a variety of responses commensurate with the variety of disturbances it seeks to control. The good-regulator theorem shows, within a formal framework, that an effective regulator must embody a model of the system it regulates. Applying these results directly to society requires caution: social institutions are not simple technical regulators, and citizens are not passive components. The book uses the ideas as design intuitions. An institution that can produce highly varied effects while observing the world through a few rigid indicators accumulates structural risk. Power to intervene and capacity to represent must be developed together.

Distributed cognition, studied by Edwin Hutchins, shows that cognitive processes can be performed by ensembles of people, instruments, representations, and procedures. A navigation team, hospital, or scientific

infrastructure can "know" more than any one member. Yet distributing knowledge can also produce institutional blindness when information exists in the components but fails to reach the place where it can change action. In the story, the difference between distributed intelligence and distributed stupidity depends on signal pathways, memory of justifications, rights of intervention, and the ability to revise categories.

Research on interdependent networks supports the importance of coupling and cascades. The study by Buldyrev and colleagues showed how dependencies between networks can produce abrupt transitions and cascading failures. The blackout in Italy on September 28, 2003 is used in the book as a historical example: contact between vegetation and a Swiss power line initiated a sequence in which loading, delays, and network coordination contributed to a large-scale outage. It is not fully accurate in causal terms to say that "a tree switched off Italy"; the tree was the initiating event in a system whose properties allowed the disturbance to propagate. That distinction between trigger and architecture is precisely what matters.

The polycentric governance in the story draws on the work of Elinor Ostrom. Her research showed that common resources can, under certain conditions, be managed through institutions created by users without automatically reducing the choice to privatization or central command. Principles identified by Ostrom include clearly defined boundaries, rules adapted to local conditions, participation by those affected, monitoring, graduated sanctions, accessible conflict-resolution mechanisms, and nested organization for larger systems. Polycentrism does not mean an absence of coordination. It means multiple centers of decision that can cooperate, learn, compete, and check one another. On climate change, Ostrom argued for action at multiple levels rather than waiting for one perfect global solution.

The book rejects romanticizing the small community. Research associated with "Dunbar's number" suggests constraints and recurring structures in social relationships, but it does not justify a biologically precise universal number. Later analyses challenged the possibility of extrapolating one exact limit from primate brain size. More importantly, group size alone does not determine the quality of freedom. Small communities can benefit from trust and reputation but can also produce conformity, nepotism, and high costs of exit. Albert O. Hirschman's distinction between exit and voice remains useful: participants

can respond to institutional decline by leaving or by trying to change the institution. When both options are weak, loyalty can become captivity.

Collective intelligence is treated functionally, not mystically. Studies initiated by Anita Woolley and colleagues identified a general factor associated with group performance across diverse tasks. The result does not mean every group is a mind or that consensus is intelligent. Performance depends on the distribution of participation, social sensitivity, cognitive diversity, and communication architecture. Research on intermittent collaboration has shown that periods of independent work can preserve exploration and diversity of solutions, while periodic exchange allows them to be combined. This is the basis for the stages of individual reflection before common deliberation in the story.

The role of AI is deliberately ambivalent. The meta-analysis by Michelle Vaccaro, Abdullah Almaatouq, and Thomas W. Malone of experiments with human-AI combinations found that hybrid teams can outperform humans alone but do not automatically produce synergy beyond the best-performing component. In some decision tasks, the combination can even reduce performance. This supports the emphasis on distribution of authority, calibrated trust, interfaces, and task type. On the other hand, experiments led by Michael Henry Tessler and colleagues showed that a system based on language models can help divergent groups construct common formulations appreciated by participants. The book extends that possibility cautiously: AI can mediate, summarize, and preserve the topology of disagreement, but textual consensus should not be confused with resolution of the conflict.

The planetary boundaries framework, introduced by Johan Rockström and colleagues and updated by Katherine Richardson and colleagues in 2023, provides a language for Earth-system processes that must be analyzed at planetary scale. The 2023 assessment concluded that six of the nine defined boundaries had been exceeded. The framework remains a subject of research and criticism; the boundaries are not precise moments of collapse and do not replace regional analysis. In the book, they illustrate problems for which observation and accounting have to be global while much of implementation remains local, regional, and sectoral.

The examples involving pandemics and asteroids are anchored in real institutions. The International Health Regulations coordinated by the World

Health Organization establish obligations concerning detection, assessment, notification, and response capacities. Their architecture combines local and national capacities with international information exchange. In planetary defense, the International Asteroid Warning Network and the Space Mission Planning Advisory Group provide mechanisms for observation and coordination. NASA's DART mission, which struck the asteroid moonlet Dimorphos in 2022, demonstrated an intentional change in its orbital period through kinetic impact. These examples support the principle of minimal globality: certain functions are irreducibly planetary, but their necessity does not justify unlimited expansion of the mandate of the institution performing them.

The terms "scale debt," "ghost agency," "governance distance," "externality radius," "correction latency," "globality budget," and "reflexive collective intelligence" are not presented as established concepts with standard definitions. They constitute a synthesis proposed by this book. They could become operational hypotheses. Scale debt could be studied through the gap between the complexity and speed of a system's effects and its capacity for observation, attribution, and correction. Governance distance could combine the number of layers of delegation, accessibility of the decision-maker, intelligibility of criteria, and appeal time. Exit capacity could be assessed through portability of data and reputation, switching costs, and the existence of alternatives. No measure would be universal, but comparisons across similar contexts could turn the intuition into an empirical research agenda.

The central thesis can be stated this way: the absolute size of a system predicts pathology less well than the mismatch between the reach of its consequences and the capacity of its governance mechanisms to represent, internalize, contest, and correct them. A global system with a narrow function can be more intelligible than a captured local institution. A highly connected network can be robust if it has controlled separations and redundancy; a smaller one can be fragile if it depends on a single node. A community can be autonomous without being free for its members. "Human scale," therefore, is not an ideal population size but a relationship among knowledge, effect, accountability, and the possibility of exit.

The book uses the expression "collective consciousness" only in a limited, functional sense. It does not claim that a society has phenomenal experience

comparable to that of an individual. The open question of the nature of consciousness is not resolved by organizational metaphors. What can be designed more concretely is a collective capacity for perception, memory, integration of perspectives, action, and self-correction. AI can extend that capacity, but it does not by itself provide political legitimacy or shared values. An infrastructure that sees more can become a distributed cognitive exoskeleton or a central apparatus of control. The difference is institutional, not merely technical.

Finally, the story insists on correctability. No architecture guarantees the right decision. Democratic institutions, markets, communities, expertise, and artificial models have different ways of discovering information and different ways of failing. Mature design does not seek an infallible mechanism but a combination in which one mechanism's errors can be detected by the others, power has limits, and decisions with large effects are sufficiently reversible or carefully scrutinized. In this sense, the wisdom of limits is not a philosophy of retreat. It is the discipline of building capacity together with brakes, connections together with switches, and intelligence together with the right to contradict it.

Selected Bibliography

Ashby, W. Ross. An Introduction to Cybernetics. London: Chapman & Hall, 1956.

Bernstein, Ethan, Jesse Shore, and David Lazer. "How Intermittent Breaks in Interaction Improve Collective Intelligence." Proceedings of the National Academy of Sciences 115, no. 35 (2018): 8734–8739.

Buldyrev, Sergey V., Roni Parshani, Gerald Paul, H. Eugene Stanley, and Shlomo Havlin. "Catastrophic Cascade of Failures in Interdependent Networks." Nature 464 (2010): 1025–1028.

Cheng, Andrew F. et al. "Momentum Transfer from the DART Mission Kinetic Impact on Asteroid Dimorphos." Nature 616 (2023): 457–460.

Conant, Roger C., and W. Ross Ashby. "Every Good Regulator of a System Must Be a Model of That System." International Journal of Systems Science 1, no. 2 (1970): 89–97.

Dunbar, Robin I. M. "Neocortex Size as a Constraint on Group Size in Primates." Journal of Human Evolution 22, no. 6 (1992): 469–493.

Helbing, Dirk. "Globally Networked Risks and How to Respond." Nature 497 (2013): 51–59.

Hirschman, Albert O. Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States. Cambridge, MA: Harvard University Press, 1970.

Holling, C. S. "Resilience and Stability of Ecological Systems." Annual Review of Ecology and Systematics 4 (1973): 1–23.

Hutchins, Edwin. Cognition in the Wild. Cambridge, MA: MIT Press, 1995.

Jordan, Andrew, Dave Huitema, Harro van Asselt, and Johanna Forster, eds. Governing Climate Change: Polycentricity in Action? Cambridge: Cambridge University Press, 2018.

Kellogg, Katherine C., Melissa A. Valentine, and Angèle Christin. "Algorithms at Work: The New Contested Terrain of Control." Academy of Management Annals 14, no. 1 (2020): 366–410.

Lindenfors, Patrik, Andreas Wartel, Johan Lind, and Robin I. M. Dunbar. "Dunbar's Number Deconstructed." Biology Letters 17, no. 5 (2021): 20210158.

Meadows, Donella H. Thinking in Systems: A Primer. White River Junction, VT: Chelsea Green, 2008.

Ostrom, Elinor. Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge: Cambridge University Press, 1990.

Ostrom, Elinor. "Polycentric Systems for Coping with Collective Action and Global Environmental Change." Global Environmental Change 20, no. 4 (2010): 550–557.

Perrow, Charles. Normal Accidents: Living with High-Risk Technologies. New York: Basic Books, 1984.

Richardson, Katherine et al. "Earth beyond Six of Nine Planetary Boundaries." Science Advances 9, no. 37 (2023): eadh2458.

Rockström, Johan et al. "A Safe Operating Space for Humanity." Nature 461 (2009): 472–475.

Saltzer, Jerome H., David P. Reed, and David D. Clark. "End-to-End Arguments in System Design." ACM Transactions on Computer Systems 2, no. 4 (1984): 277–288.

Scott, James C. Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed. New Haven: Yale University Press, 1998.

Selbst, Andrew D., Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. "Fairness and Abstraction in Sociotechnical Systems." Proceedings of the Conference on Fairness, Accountability, and Transparency (2019): 59–68.

Simon, Herbert A. "A Behavioral Model of Rational Choice." Quarterly Journal of Economics 69, no. 1 (1955): 99–118.

Simon, Herbert A. "The Architecture of Complexity." Proceedings of the American Philosophical Society 106, no. 6 (1962): 467–482.

Strathern, Marilyn. "Improving Ratings: Audit in the British University System." European Review 5, no. 3 (1997): 305–321.

Tessler, Michael Henry et al. "AI Can Help Humans Find Common Ground in Democratic Deliberation." Science 386 (2024): eadq2852.

Union for the Co-ordination of Transmission of Electricity. Final Report of the Investigation Committee on the 28 September 2003 Blackout in Italy. Brussels: UCTE, 2004.

Vaccaro, Michelle, Abdullah Almaatouq, and Thomas W. Malone. "When Combinations of Humans and AI Are Useful: A Systematic Review and Meta-Analysis." Nature Human Behaviour 8 (2024).

Weick, Karl E., and Kathleen M. Sutcliffe. Managing the Unexpected: Resilient Performance in an Age of Uncertainty. 2nd ed. San Francisco: Jossey-Bass, 2007.

Woolley, Anita Williams, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. "Evidence for a Collective Intelligence Factor in the Performance of Human Groups." Science 330, no. 6004 (2010): 686–688.

World Health Organization. International Health Regulations (2005). 3rd ed. Geneva: World Health Organization, 2016.