

FRONTIÈRE

FRONTIER INTELLIGENCE FOR LEADERSHIP ADVISORY

WHITE PAPER N° 1

---

# The World That Is Coming

*Intelligence is ceasing to be scarce. This is not a story of replacement  
— it is a story of opening.*

*The frontier of intelligence, in the service of your leadership.*

FULL DEVELOPMENT  
AOÛT 2026

FRONTIERE.AI  
MARSEILLE — FRANCE & INTERNATIONAL

## What this document demonstrates, and what it does not

This document is a theoretical demonstration. It argues its case, draws on established work, and states its hypotheses as such.

It invokes no engagement result, no client case, no proprietary data. The situations it describes are composite scenarios drawn from real situations. The day we have results, we will publish them anonymised — and we will say so.

What it offers is a line of reasoning you can contest, and frameworks you can apply in your own organisation.

---

### THE RED LINE, CONTRACTUAL

the machine never assesses a person · no decision is automated · the process is consented to and its findings returned · every conclusion is signed

# Contents

|    |                                                          |
|----|----------------------------------------------------------|
| 01 | The document in one page                                 |
| 02 | Tuesday, five years from now                             |
| 03 | 1. Intelligence has always been rationed                 |
| 04 | 2. What abundance really changes                         |
| 05 | 3. The lesson of the dynamo                              |
| 06 | 4. If scarcity lifts, the form is no longer justified    |
| 07 | 5. Abundance plugged into the old form                   |
| 08 | 6. Designing the interface of intelligence               |
| 09 | 7. Three ways of getting it wrong                        |
| 10 | 8. What this document demonstrates, and what it does not |
| 11 | 9. The Tuesday morning question                          |

# The document in one page

---

There are, in your organisation, questions no one has ever investigated. Often they are the heaviest ones. Investigating them would have taken six weeks and four people, and no budget line carries that. So the matter is settled by rough estimate, as it always has been, and no one finds anything to object to.

Those are the questions this document is about. Artificial intelligence is readily told as a story of replacement: what the machine will take, what will be left to us. That is the wrong way to put it. Replacement is the short story, and it will soon be told; the long story lies elsewhere, in everything that never took place. For intelligence has always been the most rationed resource in economic life — expensive, slow, lodged in a few heads and a few counted hours. We built our organisations to economise on it, and we ended up taking that economy for the natural order of things.

This document argues three things.

**Scarcity was the central fact.** The modern form of the organisation — division of labour, hierarchy, written procedure, impersonal rule — is not a bureaucratic pathology: it is a technology of reliability, invented to produce defensible decisions out of people none of whom can hold the whole file. Its rigidity is not its flaw, it is its price, and paying that price was rational. Every organisation is, in its own way, a machine for economising on a scarce resource — and that resource is thinking attention.

**The lifting of scarcity is not, in the first place, a gain in speed; it is a widening of the possible.** Industrial history has already shown this once: four decades separate the first power station from the productivity gains expected of it, and those gains did not come from the old factory made cheaper — they came from what steam had never allowed. The same holds here. The first visible effect of abundance is the acceleration of what was already being done; its real effect is the investigation of what was not — the questions no one asked, the files no one opened, the contradictions no one could afford to look for.

**This shift does not happen by itself, and its obstacle is not technical.** A new capability plugged into an old form does not deliver what it promises: it produces the plausible, faster. What has to be designed is not intelligence — it is available, identical for everyone, better next month, and no one will hold a durable advantage on that ground. What has to be designed is **its interface**: where it enters the organisation, what it opens that was not being done, what it is allowed to touch, what it must report, what it never concludes, and who answers for what it produces. That work is a profession, and it has no school yet.

**The status of this text**, stated at the outset because it governs the reading: it is a theoretical demonstration. It argues, it rests on established lineages, it states its assumptions as assumptions. It invokes no engagement result, no client case, no proprietary data — and we shall produce none for the occasion. What the reader can expect from it is a line of reasoning they can contradict, and a grid they can apply in their own organisation tomorrow.

## Tuesday, five years from now

---

It is eight in the morning, a Tuesday in 2031, and the night has been working. The file the committee will open at nine has been investigated while everyone slept: the documents read — all of them —, the figures cross-checked, the precedents summoned, the contradictions looked for and found. No one wrote that note; everyone has read it. The evening before, a question put in a single sentence mobilised, for a few minutes, more analytical power than a whole team assembled in a month ten years earlier. No one was surprised. No one is surprised by the light when they flick the switch.

And at nine, exactly what was there a century ago will still be in that room: a few people around a table, a decision that commits more than they do, and the question of who signs.

One thing alone has really changed, and it cannot be seen: ten years earlier, no one would have investigated that file. Not through negligence — through arithmetic. It would have taken six weeks and four people for a question that was not worth six; so it would have been settled by rough estimate, as it always was, and no one would have found anything to object to. Intelligence was rationed, and the rationing was so old that it had ceased to be visible.

It is that room this document sets out to explain: not the machine that worked through the night, but the decision that did not exist and now does. We shall proceed in five movements. First, what scarcity cost us without our counting it, and the form of organisation it made rational. Then what the lifting of that scarcity opens — and why industrial history tells us that the opening arrives late, and never of its own accord. Then what becomes of the form of the organisation when the constraint that justified it lifts. Then what happens when abundance is plugged into the old form with nothing redrawn. Finally, the work required for the shift to occur, and the name of that work.

# 1. Intelligence has always been rationed

---

## 1.1 — What a thought used to cost

Take the measure of what an hour of thinking attention applied to a precise problem was worth twenty years ago. You needed someone who knew — so a career, so a training, so a salary. You needed them to be available — so not doing something else. You needed them to have the documents — so someone had to assemble them. And you had to wait: serious thought is slow, it proceeds by readings, returns, contradictions.

Add it up, and you obtain the central fact of the whole economic life of the twentieth century: **analysis was expensive, and its expense governed everything else**. It governed what one decided to know. It governed how many people one employed, and for what. It governed, more deeply than one might think, the very form of institutions.

Ronald Coase put the question as early as 1937, in an article that remains the obligatory starting point: why do firms exist, rather than a market of contracts between individuals? Because resorting to the market costs something — one must search, compare, negotiate, draft, monitor. The firm exists to avoid those costs. Now, a significant part of them is *cognitive* in nature: searching is cognitive, comparing is cognitive, monitoring is cognitive. We do not attribute this reading to Coase — his question was the organisation of exchange, and his successors have shown that these costs also arise from the opportunism of the parties, which has nothing cognitive about it. We say only that a good half of his answer rests on the scarcity of thinking attention, and that the boundary of a firm runs, for that half, where economising on thought ceases to pay.

## 1.2 — The organisation as a machine for economising on intelligence

Herbert Simon named the constraint head-on. From *Administrative Behavior* (1947), and under the name of **bounded rationality** ten years later, he established that human beings do not optimise: they satisfice. They do not compare every option, they stop at the first that meets their level of aspiration; they do not process all the available information, they process what their capacity allows. This is not a moral failing, and neither is it a simple limitation of the brain: it is the meeting between a finite capacity and an environment that exceeds it — Simon spoke of two blades of a pair of scissors. The whole architecture of modern organisations is an answer to that meeting.

Look at what the division of labour does: it cuts a problem too large for one head into pieces that fit inside a head. What hierarchy does: it spares each person the burden of knowing everything, by designating who decides what. What written procedure does: it capitalises a thought already had, so that it need not be had again — a procedure is intelligence in a tin. What the impersonal rule does: it makes the decision robust to the circumstances of whoever applies it — their tiredness, their mood, their seniority — and therefore makes it possible to hold a constant level of quality across thousands of decisions, which no individual talent guarantees.

Jay Galbraith formalised the reading in the nineteen-seventies: an organisation is an **information-processing apparatus**, and its form follows from the quantity of information its task requires it to process. The greater the uncertainty, the more must be processed — and there are then only two families of solution: reduce the need for information (make units autonomous, build buffer stocks, subdivide) or increase the capacity to process it (information systems, lateral relations, committees). The whole variety of the last century's organisation charts lies in that pair.

Max Weber, finally, had described this form not as a failing but as **technically superior to any other**: the surest means ever invented of producing predictable, documented and contestable decisions at scale. This must be said plainly, because thirty years of managerial discourse have installed the opposite — and it must be said in full, for Weber was not praising it: he feared its human price, that "iron cage" in which the best-regulated organisation produces, as he put it, specialists without spirit. His reading is therefore exactly the one we need: a superiority observed, together with its bill. Slowness was not a failure of execution — no one ever chose slowness, they chose reliability, and slowness was the price. And that price was rational *precisely because* intelligence was scarce.

What these four readings say together, and none says alone: **an organisation is a machine for economising on thinking attention**. That is its hidden function, the one no one discusses at a strategy off-site, and it is the one that best explains why organisations resemble one another so closely. This synthesis is ours; the four authors did not put it in these terms.

### 1.3 — The invisible cost: what was never investigated

Here the reading turns, and this is the heart of the document.

A machine for economising on a scarce resource produces two effects. The first is visible and gets celebrated: it fits within a budget of attention what would not otherwise fit. The second is invisible, and no one counts it: **it decides, continuously, everything that will not be investigated**.

Every organisation performs that sorting, every day, without putting it into words. Analysis goes to the decisions that justify it — the largest, the most visible, the closest to the figures. Habit takes the rest: we do as we did last time, we follow the rule, we defer to the good sense of whoever does this for a living. And for whole territories there is neither analysis nor habit: there is nothing. The questions no one asks because asking them would cost more than the answer returns. The documents no one reads because no one has six weeks. The contradictions no one looks for because looking for a contradiction is a job in itself, and no budget line carries it.

The decisive feature of that sorting is that it is **invisible to whoever performs it**. One does not feel as a lack what one has never had; one does not count as a loss a decision one never knew could be illuminated. Scarcity, when it lasts long enough, ceases to appear as a constraint and disguises itself as the natural order of things: it is no longer "we cannot afford to investigate that question", it is "that question is not one that gets investigated".

Frontière has a textbook case of this mechanism in hand, and has devoted a whole paper to it. Decisions of capital — acquiring, merging, passing on, financing — are among the most heavily instrumented in economic life: accounting audits, legal, tax and employment reviews, strategic diligence, and now cyber. One variable alone escapes the apparatus: the team running the company, its dependencies, its dynamics, its fit with the project. It escapes the apparatus because it is expensive to investigate, slow to read, and no binder is waiting for it. *One does not readily look at what one cannot afford to read.* The reader will find that demonstration in our white paper no. 2, *The Blind Spot*.

What this document adds is more general, and more uncomfortable: **your organisation probably contains several blind spots of this nature, and you cannot name them, because they are exactly what your budget of attention has always excluded.** That is the first useful exercise this text can set off: to look, not for what you do badly, but for what you do not do at all.

**Scarcity does not merely limit what we do: it shapes what we believe possible.**

## 2. What abundance really changes

---

### 2.1 — From product to condition

In 1900, electricity was a product: one went to see it at exhibitions, it was sold to the bold, it was debated in drawing rooms. By 1930, it was a condition: it had disappeared into the walls. No one, since, "uses" electricity. One inhabits it. It is visible nowhere and necessary everywhere — to the point where its only perceptible event is the outage.

Artificial intelligence follows the same curve, faster. What was sold yesterday in analyst-days is produced today in machine-hours; each generation widens the perimeter — systems chain tasks together, verify themselves, hold whole workflows. You have the signs of it on your desk, and you will measure them better than we can: a note that took weeks, a document review that mobilised a team, a monitoring brief that required a subscription and two interns — the order of magnitude of time and cost has changed, and it goes on changing.

Extend the curve five years, without asking any miracle of it: intelligence ceases to be a tool one picks up and becomes a medium one inhabits. Present in every document, every exchange, every decision prepared — and as unremarkable as the current in the walls.

Does believing this require believing in the singularity — the point at which the machine improves the machine? That horizon remains disputed, and honesty requires saying so. But ubiquity does not need the marvel: it needs only a cost curve, and the curve is there. This is a deliberately low position, and we hold to it: everything that follows stands even if one rejects the strongest prophecies, and requires no more than the continuation of an observable trend. Were that continuation itself to prove false — were costs to rise again, were capabilities to plateau — the substance of this document would fall. That is the bet it takes on, and it is the only one.

### 2.2 — Replacement is the short story

Public debate has settled on the question of replacement: which occupations, which tasks, what percentage. The question is legitimate and it will have its answers. It is also, structurally, the less interesting of the two — because it reasons at constant perimeter. It takes the world as it is, with its tasks as they exist, and asks who will do them. It therefore assumes that the list of things to be done is stable, and that only the hand doing them changes.

Yet that list has never been stable. It has always been the product of what one could afford. The rationing of thinking attention shows not only in the way tasks were carried out: it shows in the very composition of the list. Change the price of the resource, and it is not the hand that changes: it is the list.

That is the shift this document asks of the reader. **An abundance of intelligence does not first displace what exists; it investigates what never was.** It makes possible and reasonable what was possible but unreasonable:

opening a file that was not worth the trouble, reading a corpus in full instead of a sample, looking systematically for contradiction instead of waiting for it to surface, investigating a question no one had ever devoted six weeks to because it was not worth six — and that is worth three hours.

The right question is therefore not *what is going to be taken from me?* but **what becomes reasonable to do, that I have never done?**

### 2.3 — What becomes scarce, and why that does not contradict the above

A second, symmetrical consequence follows from the same shift — and here we are stating our thesis, not an established fact. A resource that becomes ubiquitous ceases, at that same moment, to differentiate: no one ever built a durable advantage on being connected to the grid. The hierarchy of value therefore inverts all the way down. Data is worth almost nothing: it overflows. Analysis becomes a commodity: the machine produces it. What becomes scarce — relatively, and therefore preciously — is what the machine will not carry: judgment that commits a responsibility, consented access to people, the confidence of a decision-making table, the signature. *A world in which analysis is free is a world in which judgment is the only differentiator left.*

One clarification is called for, because it governs what follows. "Analysis becomes a commodity" does not mean "analysis becomes right". What the machine produces in abundance is a material: structured, exhaustive, coherent — and not necessarily true. It becomes an analysis at the moment when someone verifies it, weighs it and takes it on. The gap between the two is the whole subject of this document: abundance delivers the material; *design* decides whether it becomes an investigation or merely volume.

The two movements therefore do not contradict one another, they nest, and the order matters: **abundance opens the field of investigable questions; the scarcity of judgment decides who will know what to do with it.** A world in which everything can be investigated is a world in which one can drown in investigations no one concludes. Which is precisely why the shift does not happen by itself.

**The question is not what the machine will take from you. It is what it makes possible for you — and what you will do with it.**

## 3. The lesson of the dynamo

---

### 3.1 — Forty years

History has already given the measure of what is coming, and its lesson is counter-intuitive. The economist Paul David established it for electricity in an article that has become a classic — *The Dynamo and the Computer*, American Economic Review, 1990: **four decades** separate the opening of the first power station, in 1881, from the manufacturing productivity gains expected of it, which appear only at the turn of the nineteen-twenties. Four decades during which a manifestly superior technology spread without producing the promised effect, to the point where contemporaries doubted its value — this is the "productivity paradox" David uses to illuminate the computing paradox of his own time.

The reason did not lie in the machine. Work in industrial history — Warren Devine gave the detailed account of it in 1983 — shows what happened inside the factories. The steam engine had been replaced by a large electric motor, in the same place, on the same base. Everything else remained: the line shaft running below the ceiling, the belts descending from it to each workstation, the floors stacked around the drive axis, the shop floor organised not by the flow of work but by distance from the source of power. Steam had been electrified.

### 3.2 — What steam had never allowed

The gains came the day people stopped *installing* electricity and began *designing the factory around it*. One motor per machine — so no more line shaft, so no more constraint of proximity. The shop floor laid flat, on a single level. Workstations ordered according to the flow of work. Lighting, ventilation, safety — everything the forest of belts had forbidden.

And the point to hold on to is this: what the redesigned factory produced was not the old factory more cheaply. It was **what steam had never allowed** — the continuous line, the rapid reconfiguration of production lines, precision, production in series. The gain was not a gain in efficiency on unchanged tasks: it was the opening of a space of production that did not exist. Measured productivity followed only afterwards, long afterwards — because it was measuring a world that was no longer the right one.

### 3.3 — The delay is not technical

This precedent carries three lessons, and they govern everything that follows.

**The first:** the delay between a capability and its effects is not *only* a delay of technical maturation — it is also, and to a decisive extent, a delay of **organisational redesign**. David gives several causes for it, and honesty requires recalling them: the existing stock of capital, which it would have been absurd to scrap before it had been fully depreciated; the complementary innovations required to distribute current; standardisation; and the slow learning of industrial architects. Redesign is one cause among those. But it is the only one that depends on a decision rather than on time — and therefore the only one that concerns you.

**The second:** the intermediate period — the one in which the new capability is plugged into the old form — is not neutral. It is costly, disappointing, and it produces false conclusions. That is where durable scepticism is born, and it is very exactly the period we are in.

**The third,** and the most useful: one does not cross that delay by waiting. None of the factories that made the shift made it by natural ageing; all of them did deliberate, costly design work, led by people for whom it had become a profession. The shift is not a phenomenon, it is a decision.

**A new capability does not deliver what it promises for as long as it is asked to redo the old world faster; it delivers when what did not exist is designed around it.**

## 4. If scarcity lifts, the form is no longer justified

---

### 4.1 — The reasoning, and its status

We have established two things: that the modern form of the organisation is a rational answer to the scarcity of thinking attention, and that this scarcity is lifting. The conclusion presents itself as a deduction, and we give it for what it is:

#### FIGURE

If the constraint that justified the form disappears, the form is no longer justified — but it does not disappear of its own accord.

This articulation is a **practitioner's hypothesis**, resting on established lineages. It is not a research result and we shall never present it as one. It is refutable in principle — if, five years from now, the organisations that will have redesigned their relationship to intelligence are indistinguishable from those that merely deployed it, it will be false and that will have to be said. We shall not go so far as to claim we hold the protocol that would settle it: establishing the distinction would require a comparable population, a measure of what "redesigning" means, and a counterfactual, and we have none of the three. An honest hypothesis also says what would be needed to decide it.

### 4.2 — Mechanistic and organic: what theory had already named

The useful distinction has existed since 1961. Tom Burns and George Stalker, observing British firms confronted with the arrival of electronics, described two forms. The **mechanistic** form: specialised and stable tasks, authority by position, vertical communication, knowledge concentrated at the top — effective in a predictable environment. The **organic** form: tasks continuously redefined, authority by competence on the subject of the moment, lateral communication, distributed knowledge — superior when the environment moves.

Their most interesting observation is not the typology, it is the failure: the firms they watched attempt the passage from mechanistic to organic overwhelmingly failed, and they failed for reasons that had nothing technical about them. Burns and Stalker show that three systems coexist in every firm — the organisation of work, the political system, and the structure of careers and status — and that the last two get the better of the first. The mechanistic form distributes authority according to position; the organic form redistributes it according to the competence of the moment. No one willingly gives back an authority they hold from their place — above all when that place is also their career.

Paul Lawrence and Jay Lorsch established in 1967 the qualification that prevents any naive reading: there is no good form in the absolute, only forms that fit their environment or do not — and the more an organisation

differentiates, the more it must invest in its integration. Differentiation without integration does not produce agility, it produces baronies.

### 4.3 — The theorem that formalises the argument

Take a situation every leader has lived through. A client asks for something the catalogue does not provide for; a bespoke answer would be needed, but the bespoke answer requires a trade-off, the trade-off requires a meeting, and the meeting costs more than the business. So the answer is no, politely, with a form to fill in. Multiply by a thousand and you obtain a rigid organisation — without any decision to be rigid ever having been taken.

A result from cybernetics gives that experience its rigorous form. W. Ross Ashby, in 1956, established what is known as the **law of requisite variety**: only variety can destroy variety. In other words, to hold a system within the desired limits, the regulator must have at least as many distinct states as the disturbances it has to absorb. Stafford Beer later drew a genuine engineering discipline from it: when the two varieties do not equalise, they equalise all the same — downwards.

That is exactly the mechanism of the form to fill in. The world an organisation faces presents a certain variety — clients, markets, regulations, hazards. Its regulating apparatus — its rules, its procedures, its committees, its managers — must oppose to it a comparable variety. Now, when thinking attention is scarce, increasing the variety of the regulator is expensive, very expensive: every exception demands a head. So the equalising is done on the other side: the **variety of the world is reduced** to bring it within the measure of the regulator. Offerings are standardised, clients are categorised, particular cases are discouraged. On this reading, the rigidity of a large organisation is not an administrative failing: it is variety management under a cognitive-resource constraint — and it was, until yesterday, the only management possible.

What abundance changes is exactly this: **the variety of the regulator ceases to be expensive**. It becomes possible to handle the particular case without its costing more than the general one — and therefore, for the first time, to stop flattening reality in order to make it manageable. That is very precisely where the promise lodges — and one sees at the same time why it is not a promise of speed: it is a promise of fineness.

### 4.4 — Organic does not mean disorder

The misreading must be forestalled, because every leader will make it, and they will be right to be wary of it. "Less structure" is not the conclusion of this reasoning — it is its most predictable failure.

An organisation that withdraws its procedures and puts nothing in their place does not become organic: it becomes a mechanistic organisation that has lost its instruments. The organic form **demand**s more discipline than the mechanistic form, not less. It requires knowing at all times who decides what, where hierarchy said it once and for all. It requires fast and honest feedback loops, where the cascade made do with a quarterly report. It requires active integration, failing which differentiation dissolves it. And it is infinitely easier to throw out of tune: Jay Forrester and then Peter Senge established this for feedback systems — transmission lags and over-strong corrections produce oscillations there that amplify on their own, that familiar phenomenon in which

each correction calls forth the next, more violent one. That steering by loops is preferable to steering by cascade, they did not say, and we advance it on our own account; what they did show, on the other hand, is that it tunes more finely and goes out of tune faster.

That is what makes the subject difficult, and what makes it a work of design rather than a movement of mood. One does not become organic by removing rules. One becomes organic by designing the apparatus that replaces what the rules were doing.

**The old form was not irrational. It was rational under a constraint that is lifting — and nothing guarantees that what replaces it will be better.**

## 5. Abundance plugged into the old form

---

Here is what the intermediate period produces — the period of the dynamo on the steam engine's base. These three mechanisms are given as **argued hypotheses**, deduced from the reasoning above and from what we observe among the people we talk to; they are not study results, and the reader will test them against their own organisation better than we can.

**The acceleration of the plausible.** A chain of intelligence produces in hours what the organisation took weeks to produce. But the control mechanisms have not changed cadence: the second reading, the challenge, the committee that meets once a fortnight. Volume therefore passes the control without being controlled — not through negligence, but because the control was sized for a throughput that no longer exists. And what passes is the plausible: well written, well structured, coherent, and unverified. The organisation gains in production what it loses in reliability, exactly the reverse of what bureaucracy guaranteed it.

**The silent displacement of responsibility.** The note comes from the machine, therefore from no one. Whoever passes it on did not write it; whoever receives it assumes it has been verified; whoever signs it ratifies work whose path they have no time to retrace. The signature formally remains, but it has ceased to be an act of judgment and become an act of procedure. This is the gravest of the three, because it is painless: nothing breaks, until the day something breaks and it is discovered that no one, at any point, had really concluded.

**Defensive hardening.** The organisation that senses the first two dangers and does not yet have the means to govern them takes the only decision left to it: it prohibits. An internal memo, a usage restriction, an approval committee. The reflex is sound — it really does protect, and it is better than carelessness. Its flaw is not that it is too cautious, it is that it is provisional without saying so: while the door is shut, use goes through the window, with neither traceability nor control, which produces exactly the risk it was meant to avoid. Prohibition is a good holding measure; it becomes costly the day it stands in for a policy. That, in our view, is where most serious organisations are today — and it is a good place from which to start designing.

These three pathologies share one feature, and it is the interesting one: **none of them is cured by more technology**. They do not come from a flaw in the models, they come from the absence of *design*. Which is what makes them interesting to a leader: they are problems that fall to them.

**Giving a frontier model to an organisation that is not expecting it does not make it more intelligent: it produces the plausible, faster.**

## 6. Designing the interface of intelligence

---

### 6.1 — What one does not design

This *design* is work, and that work is becoming a profession. What it bears on still has to be stated, for that is where almost the whole market goes wrong.

What has to be designed is not intelligence. The most powerful models are available to whoever pays for them, and our conviction — it is one, not a fact — is that no one will take a durable advantage there: anyone who had built their position on access to one of them would watch it evaporate at the next generation. Choosing a supplier is a purchasing decision, not a design decision.

What has to be designed is **the interface of intelligence** — and "interface" is to be heard here in the architects' sense, not the computer scientists': not a screen, but the contact surface between a capability and the organisation that has to live with it. Where intelligence enters your organisation. What it opens that you were not doing. What it is allowed to touch. What it must report, to whom, in what form. What it never concludes. And who answers for what it produces.

This profession has a name. In French it is called *l'AI design* — the design of the interface of intelligence; in English there is one right rendering of it and one only: **the design of intelligence**. *Design* in the strict sense, the architect's and not the decorator's: the overall plan before the ornament. It demands two halves that almost no one holds together: knowing what a chain of intelligence can do, and knowing what a human collective can carry. The first half goes out of date every six months; the second is as old as organisations.

### 6.2 — The five questions of the design of intelligence

Here is the grid. It is not proprietary, it requires no tool, and the reader can apply it tomorrow to any artificial intelligence system in service in their organisation — their own, or a supplier's. **A design that leaves one of these boxes empty is not finished.**

**The perimeter.** Where does intelligence enter, and what is it given to touch? The question looks technical; it is strategic, because it contains the sorting we spoke of above. The right perimeter question is not "where can I go faster?" — that is a supplier's question. It is: *which questions are we not investigating, and which of them deserve to be, now that they are affordable?*

**What comes back.** What must the chain report, in what form, at what moment, to whom? An investigation that arrives after the decision is worth nothing; an investigation no one can read in ten minutes will not be read. The return is what turns a capability into illumination, and it is almost always the most neglected box.

**The red line.** What does the chain never conclude? Every serious design has a negative part, and it is that part which defines it: an interface is judged by what it forbids. In decisions that bear on people, the prohibition is

not caution added on — it is the source of the value. A chain that scores dries up its own source: no one opens up to a system that assesses them. The restriction does not diminish the quality of the material, it produces it.

**Traceability.** Where does each assertion come from, and how is it verified? This is what separates an investigation from a well-written opinion. A chain that cannot say where what it asserts comes from is not an instrument of decision: it is a generator of groundless confidence, and that is the most dangerous form of the plausible.

**The signature.** Who answers, by name, for what is produced? Not who approved — who **answers**. The difference is that of the second pathological mechanism described above, and it is not settled by a tick-box: it is settled by an order of work in which ratifying costs more than judging.

### **6.3 — Why the design of intelligence is a profession, and not one more skill**

Put these five questions around your table, and watch what happens. Each will answer for their part, and each will be right for their part. The information systems function will know how to answer about tools, security and data, and much less on what deserves to be investigated — that is not its mandate. The legal function will know how to state the prohibition and will state it well, without having to design what the prohibition makes possible. The business functions will know what they would like to obtain, and rarely what the chain can actually produce. And the supplier, whoever they are, will have a structural reason to answer all five with a single "yes, we can do that" — ourselves included, which is precisely why these questions must be put to us too.

Each of these answers is right and partial; what is missing is their articulation, and it belongs to none of the existing posts. Now, holding together a capability that moves every six months and an organisation that takes years to change is the very definition of a design profession. Our conviction — stated as such — is that it will establish itself within five years as the function that runs information systems established itself in its own time: first a piece of improvisation entrusted to whoever took an interest, then a function, finally a discipline.

In the meantime, the design of intelligence is practised — and it is already practised, by those who have understood it, without carrying a name.

**What we design is not intelligence — it is its interface.**

## 7. Three ways of getting it wrong

---

The world that is coming will offer three of them, and they are worth naming, because they resemble one another very little and are cured differently.

**The craftsman without the machine.** Sincere, conscientious, often excellent — and exposed to being outdone in depth of investigation by an equipped competitor. It is not their judgment that weakens, it is their footing: they triangulate what they remember while the other triangulates everything. The error is comfortable because it has the shape of a virtue.

**The software without judgment.** The score no one signs, the recommendation no one takes on, the probability sold as a conclusion. It fails at the exact point where the other fails, by the opposite route: there is no one at the end of it. And one does not buy a probability; one buys someone who answers for a reading.

**Deployment without design** — the most widespread, because it looks like modernity. Abundance plugged into the old form and ordered to redo yesterday faster: the dynamo set on the steam engine's base. It is the error that produces the three pathologies of section 07, and the only one of the three committed in good faith and at great expense.

The pincer closes on both errors at once — *frontier intelligence widens what a practitioner can investigate; governed judgment is what turns it into a decision*. Between the two, *design* decides. And *design* is never finished: the frontier moves every month; what it allows, what it costs, what has to be refused it move with it.

**The state of the art is not a state. It is a conversation.**

## 8. What this document demonstrates, and what it does not

---

We owe the reader the exact measure of what they have just read, because the times produce a great deal of self-assurance on this subject and we do not intend to add to it.

**What is established** and does not belong to us: bounded rationality as a structuring constraint on organisations; the reading of the organisation as an information-processing apparatus; the distinction between mechanistic and organic forms, and the difficulty — political, status-related, never technical — of the passage from one to the other; the law of requisite variety; the historical delay of four decades between the first power station and manufacturing productivity gains, together with the part played in it by the reorganisation of factories. This work is solid, dated, cited at the end of the document, and the reader can go and verify it — we encourage them to.

**What is observable** and anyone can see on their own desk: the rapid fall in the cost of cognitive work produced by machine, and the widening of the range of tasks a chain carries out end to end without intervention.

**What is assumed**, and which we give as such — the list is deliberately long:

The central articulation of this document — *scarcity justified the form, scarcity is lifting, the form is no longer justified* — is a practitioner's hypothesis resting on those lineages, not a result. The cognitivist reading of organisational costs we draw from Coase is ours, and it is partial: those costs also rest on drivers that have nothing cognitive about them. The application of the law of requisite variety to the rigidity of large organisations is a reading, not a theorem — the theorem bears on regulating systems in general. The inversion of the hierarchy of value, which is our commercial thesis, is a deduction and deserves to be examined as such by anyone with an interest in contradicting it. The extension of the cost curve to five years is a forecast. The three pathologies of section 07 are deduced mechanisms, consistent with what we observe among the people we talk to — conversational material, not a survey, and we present it as nothing else. That the design of intelligence will become an established function within five years is a conviction. And the 2031 scene that opens this document is a working fiction, written to make thinkable what the figures do not yet show.

**What this document does not contain**, and should not have contained: no engagement result, no client case even anonymised, no proprietary data, no performance figure. This is not an omission of convenience — it is the regime we impose on ourselves for as long as the series does not exist. A practice that sells judgment must be beyond reproach on the distinction between what it has demonstrated, what it has observed and what it advances. The day we have cases, we shall publish them anonymised, and we shall say so.

What the reader can therefore legitimately expect of these pages: a line of reasoning they can contradict, and a grid they can apply. Nothing more. We would rather they closed this document knowing exactly what they have just received.

## 9. The Tuesday morning question

---

Return to the room at eight in the morning. What makes it possible is not the machine: five years from now it will be everywhere, like the current. What makes it possible is a *design* — choices made months earlier, in cold blood: what the night is allowed to investigate, what the morning alone can settle, what the chain reports and to whom, what it will never conclude, and the name of whoever signs. That *design* is improvised neither on the day of the closing nor on the evening of the crisis. It is designed beforehand — like the factory around the dynamo.

For what is most precious about that room is not what gets done faster in it: it is what gets done *at last* — the file that went uninvestigated, the question no one dared pay for, the contradiction no one could ever afford to look for. The world that is coming will not be ours on fast-forward. It will be wider — and the gap between two organisations will be measured there less by their power than by their discernment: by what they will have chosen to open, and by the rigour with which they will have governed that opening.

Five years from now, your organisation will be steeped in the same intelligence as every other, and the possible will have changed size for everyone. One question alone will set it apart, and it is the question of the design of intelligence: **who is its designer?**

# Colophon

*This document was born of the world it describes: investigated by a chain of frontier artificial intelligence, under human direction — and governed by the very rules it defends: the machine never assesses a person · no decision is automated · the process is consented to and its findings returned · every conclusion is signed — this text was weighed, concluded and signed by someone who answers for it. That is the only proof worth anything in this matter: not to declare that one works at the frontier, but to show it on the documents, and to accept being looked at.*

**Where to go next.** This thesis has two applications, and they read separately. The Blind Spot (</en/papers/the-blind-spot>) (white paper no. 2) shows what scarcity had left out of the field in decisions of capital, and how that variable is read — for those who invest. Intelligence, by design (</en/papers/intelligence-by-design>) (white paper no. 3) goes down into the practice: how a chain that investigates everything and assesses no one is concretely designed, with its parts and its limits — for those who lead.

**References.** Ronald Coase, "The Nature of the Firm", *Economica*, 1937 · Herbert Simon, *Administrative Behavior* (1947) and *Models of Man* (1957), on bounded rationality · Max Weber, *Economy and Society* (posthumous, 1922), on rational-legal domination, and *The Protestant Ethic and the Spirit of Capitalism* (1904-1905) for the iron cage · W. Ross Ashby, *An Introduction to Cybernetics* (1956), the law of requisite variety · Stafford Beer, *Brain of the Firm* (1972), on the engineering of variety · Tom Burns and George Stalker, *The Management of Innovation* (1961) · Paul Lawrence and Jay Lorsch, *Organization and Environment* (1967) · Jay Galbraith, *Designing Complex Organizations* (1973), on the organisation as an information-processing apparatus · Jay Forrester, *Industrial Dynamics* (1961), and Peter Senge, *The Fifth Discipline* (1990), on feedback loops · Warren Devine, "From Shafts to Wires: Historical Perspective on Electrification", *The Journal of Economic History*, 1983 · Paul David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox", *American Economic Review*, 1990.

*Frontière — Frontier intelligence for leadership advisory — [frontiere.ai](http://frontiere.ai) — [alice@frontiere.ai](mailto:alice@frontiere.ai)*