

FRONTIÈRE

FRONTIER INTELLIGENCE FOR LEADERSHIP ADVISORY

WHITE PAPER N° 3

Intelligence, by design

FRONTIER INTELLIGENCE FOR LEADERSHIP ADVISORY · FOR THOSE WHO LEAD

What agentic systems make possible for those who decide — and what has to be designed for that power to be worth anything

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

Bringing objectivity to the human factor in the decisions that have no second chance.

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	Opening — The Machine That Has Read Everything
02	1. What Has Just Become Possible
03	2. The False Dilemma
04	3. What the Law Draws
05	4. The Architectural Principle — Investigation, Never Verdict
06	5. Inside the practice — the division of labour, station by station
07	6. Science at the Inputs, Judgment at the Integration
08	7. Provable Governance — Proof by Deed
09	8. Conclusion — Governed Judgment

Opening — The Machine That Has Read Everything

The status of this text, stated at the outset because it governs the reading: it is a theoretical demonstration. The scene that follows is a situation type, inspired by real situations, and its quantities are illustrative. This document invokes no engagement result, no client case, no proprietary data.

By the end of an assessment engagement on an executive committee, the machine has read everything. Fourteen hours of interviews, transcribed to the word. The strategic plan and its three successive revisions. Two years of executive committee minutes. The official organisation chart, and the one nobody has ever drawn but which the exchanges make visible. The instrument reports. The documentation of the project this team is supposed to carry. It holds all of it in memory, simultaneously, without fatigue and without forgetting.

It would take a few seconds to produce what is asked of it everywhere else: a ranking of the five executives, a score per person, a flight-risk index, a replacement recommendation.

It will not do so. Not because it is incapable — it is perfectly capable, and that is precisely the problem — but because the chain in which it works was built so that it never would, and because that prohibition can be verified line by line.

That machine did not exist three years ago, and the fact that it does is the subject of another of our documents: *The World That Is Coming* argues that intelligence has always been rationed, that the scarcity is lifting, and that what then opens is not, in the first place, an age of replacement but one of widening. What remained to be settled was what one does with it when what one investigates is people. That is the object of this one.

An artificial intelligence system is ordinarily judged by what it can do. In a profession that bears on people, whose conclusions decide careers and commit capital, it has to be judged by something else: by what it has been forbidden to do, by the reasons for that prohibition, and by the ability to prove that it holds when nobody is watching. That is the subject of this document. It is not a brochure: it is the description of an architecture, delivered with its limits and exposed to contradiction. A reader who wants to find it at fault will find here what is needed to do so — that is the condition of its being worth anything at all.

1. What Has Just Become Possible

Let us begin with the capability, because that is what has changed, and because no discussion of safeguards means anything until one has measured what there is to guard.

A leader who decides about people has always worked from a sample. He has met three people out of twelve, read the summary note rather than the documents, retained whatever his memory chose to keep of the March meeting. This is not a want of seriousness: it is the normal regime of every human decision, and it has been the normal regime because thinking attention was expensive. What follows describes what ceases to be true.

Reading everything, not a sample. A well-designed chain takes on the whole of a corpus: interviews transcribed to the word, the strategic plan and its three revisions, two years of committee minutes, the documents of the project the team must carry. Not a summary of those documents — the documents. The difference is not one of speed, it is one of nature: one stops reasoning on what one had time to read and begins reasoning on what exists.

Forgetting nothing, and holding it all at once. The most rigorous human mind loses the thread beyond a few dozen elements; it compensates with notes, grids and summaries — that is, with deliberate loss. A chain holds the entire file simultaneously, which permits a move until now impossible: bringing together two sentences spoken by two different people, six months apart, in two unrelated contexts.

Seeking the contradiction instead of waiting for it. In ordinary practice, a contradiction is discovered by accident — someone remembers, someone is surprised. It can now be sought: systematically, across every document, cross-checking what is said with what is written and with what has been done. It is work nobody ever funded because it cost more than it returned on average — whereas in the rare cases where it finds something, it is worth infinitely more than its price.

Triangulating on a new scale. A practitioner cross-checks what he remembers; a chain cross-checks everything. The same signal — a leader describing how he decides — is found in his interview, in the minutes where he is seen deciding, and in the words of those who work with him. The convergence or the gap between those three sources is information of the first quality, and it was out of reach.

Investigating the particular case at the price of the general one. This is perhaps the deepest shift. An organisation that treats a case bespoke pays each time the price of fresh thinking; that is why it standardises. When investigation becomes abundant, the bespoke question ceases to be a luxury. Applied to an executive team: one no longer applies a generic grid to a particular situation, one investigates *that* situation, against *that* thesis.

Never stopping. Continuous watch on transactions, executive movements and the environment of a portfolio company no longer requires a subscription and two people: it runs. The file is no longer a photograph taken on the date of the engagement; it can become live again.

That is the state of the capability. It is not hypothetical, it is not five years away: it is what a well-designed chain does today, and the corpus of the practice — its method, its studies, this document itself — is produced by one, under human direction and human judgment. *Proof by deed, never by declaration.*

And now the bad news. None of the above is to be picked up off the ground. Each of these capabilities has a twin pathology, and it is triggered by the absence of design. Reading everything produces a volume nobody re-reads. Forgetting nothing produces connections without relevance. Seeking the contradiction manufactures some — a language model will always find something if asked to. Triangulating without discipline mixes sources of unequal quality. And continuous watch, with no recipient and no threshold, becomes a background noise one ends up switching off.

This symmetry is the reason for everything that follows. **Power is not deployed: it is designed.** What separates a chain that lights up a decision from a chain that clutters it is neither the model used nor the quantity of computation — it is a series of design choices someone made, or failed to make.

A capability without design is not an advantage: it is an additional workload, arrived faster.

2. The False Dilemma

Every leader who has sought help on a difficult decision about people — recruiting a member of the executive committee, arbitrating a succession, knowing whether the team in place will carry the phase now opening — has met two families of answer, and has found both unsatisfying without always being able to say why.

On one side, practice by hand: an expert judgment expressed with assurance. Four interviews conducted, what can be read is read, and a view delivered whose quality rests on personal acuity. The approach is sincere and often right. But it stands alone in front of a volume no human attention can cover: the strategic plan skimmed, retained from Tuesday's interview whatever chimed with Thursday's, triangulated whatever happens to be remembered. Confidence in an expert judgment, as research has long established, correlates very weakly with its accuracy. This is not a failure of competence: it is the normal condition of a human brain faced with too much material.

On the other side, the software. A platform that administers tests, aggregates, and delivers a dashboard: scores, a fit index, sometimes a traffic light. Volume, this time, is no longer the problem. But the object produced is signed by nobody. It cannot be explained to the person concerned except by reading her own score back to her. It cannot be defended before a board, because no challenger can interrogate the chain that produced it. And it silently shifts a legal risk onto whoever uses it — we shall return to this, it is the least understood point in the whole subject.

The error lies in believing that one must choose. Both answers fail on the same point, by opposite routes: they confuse the volume of the investigation with the quality of the judgment. The craftsman has the judgment and denies himself the volume that would make it solid. The software has the volume and believes it produces judgment.

If these two offerings survive side by side, it is not through market inattention: each answers a real budget line. The craftsman is bought when a decision is too sensitive to be entrusted to a tool, and he is often bought almost confidentially, by a person rather than by an organisation. The software is bought when volume is too great to be handled any other way, and it is bought by a function that must account for a process. Both needs are legitimate; they are simply never served together. And it is at the levels where decisions are at once rare and heavy — the executive committee, the board, the team of a portfolio company — that the absence of a third answer costs most dearly, since both are needed at once: the depth no single person produces, and the commitment no software carries.

The times, moreover, have just reshuffled this false dilemma, and they reshuffle it on both sides at once. The cognitive work of analysis has been repriced downwards, definitively: what was sold yesterday in man-days is produced today in machine-hours — which takes from the artisan his excuse. And what becomes scarce in proportion is precisely what the machine will not carry: judgment that commits a responsibility, consented access to people, the confidence of a committee, the signature at the foot of a conclusion — which takes from the software its pretension. We have argued this movement for its own sake in *The World That Is Coming*; here it is enough to retain that neither of the two existing offers comes out of it intact.

Hence the frame of this entire document: **frontier intelligence widens what a practitioner can investigate; governed judgment is what turns it into a decision.** The gap can no longer be closed by hand. The software cannot produce what a decision-maker actually buys, which is not a probability but someone who answers for a reading. The differentiator is therefore neither the tool nor the expert: it is the design that articulates them — and a design, unlike a talent, can be described, verified and contested.

3. What the Law Draws

A leader who commissions an assessment of his team sooner or later receives a question from his own data protection officer, or from his general counsel. It comes in two parts: on what basis are these people being processed, and does an artificial intelligence decide anything about them? Most providers answer with reassurance. Reassurance is not an answer — the answer is an architecture, and European law has written its specification with a precision few professionals in the sector have taken the time to read.

The category. Assessing people in a work context by means of an artificial intelligence system falls, under the AI Act, within Annex III — the list of uses qualified as high-risk, whose point relating to employment and workforce management explicitly covers the evaluation of persons. The timetable of this regime has moved: the package known as the *Digital Omnibus*, adopted in 2026, defers its application to December 2027. That deferral changes a deadline; it does not change a specification. A system designed to be defensible in 2027 is a system that is defensible today; the converse — building fast on a bet on the delay — amounts to scheduling a rebuild for the end of the year in which the market will begin to ask for accounts.

The hard prohibition, already in force. One provision benefits from no deferral and has applied since February 2025: the prohibition of artificial intelligence systems intended to infer the emotions of a natural person in the workplace, save for medical or safety reasons. It sits among the prohibited practices, those carrying the highest penalty ceiling in the regulation. For our profession it is the sharpest line there is: any chain that records interviews and entrusts them to a machine is working at the exact edge of that frontier. The question is not whether one approaches it — one approaches it by construction — but whether one can demonstrate which side one stands on, and demonstrate it otherwise than by assertion.

Transparency. The same regulation has required, since the summer of 2026, that people be informed when they interact with an artificial intelligence system or when content is produced by a machine. Translated into the profession: the person being assessed must know what the machine does with her words. This obligation costs something only to those systems that had an interest in its being ignored.

And the point almost nobody anticipates. The General Data Protection Regulation — the GDPR — prohibits a decision producing legal effects, or similarly significant effects, from being based solely on automated processing. This rule is ordinarily read as a constraint bearing on the provider of the system. The Court of Justice of the European Union has decided otherwise in its case law on *scoring*: where a third party produces a score and the recipient uses it in a determining manner to decide, it is the score itself that constitutes the automated decision. The reasoning transposes without effort to the assessment of executives. A service provider who delivers a score on a person to a client who will use it to decide is not merely taking a risk: he is transferring one to his client, generally without telling him. Conversely, a service provider who delivers a human conclusion, reasoned, signed, and who stipulates by contract that the deliverable informs the decision without making it, protects both parties at the same time as the person assessed.

This symmetry has a practical consequence that management teams generally discover too late: in the economy of the regulation, whoever *deploys* an artificial intelligence system answers for it, and not only

whoever supplies it. A company that commissions an AI-equipped assessment is therefore not merely buying a deliverable — it enters, as far as it is concerned, a regime of obligations of its own: informing the persons, documenting the use, ensuring the competence of those who interpret the results. The good service provider is the one who arrives with those documents already written and hands them to his client; the bad one is the one who leaves his client to discover them six months later, through a question from his own lawyer. The test is simple, and it discriminates very quickly.

One question remains that cannot be settled in a general document, and that it would be dishonest to treat as closed: the legal basis for the processing. In a relationship of subordination, an employee's consent is a legally fragile basis, because the imbalance between the parties weakens the freedom to refuse. Two things must therefore be distinguished that the sector's marketing readily conflates. **Fair consent** — the person knows what is being done, accepts it, and receives her debrief — is a requirement of method and of professional ethics: without it, the material gathered is worth nothing, and we shall return to this. **The legal basis for the processing** is a distinct legal question, assessed deployment by deployment, with counsel, and documented in the record of processing activities and in the impact assessment. A serious service provider handles both; he does not mix them in order to make the second fit inside the slogan of the first.

The state of the law summarised here is that of the summer of 2026 and it will continue to move; dates and scopes must be verified at source before any decision. What will not move, however, is the direction: the European regulator has written, in substance, that a system must not conclude alone about a person, that it must be transparent about what it does, and that whoever employs it answers for it. That is, word for word, the professional ethics of the trade — with penalties added. **Documented compliance is therefore not a cost of getting into line: it is the form taken by a pre-existing requirement, and, for whoever has built it before the others, a barrier to entry.**

4. The Architectural Principle — Investigation, Never Verdict

Two practices can employ the same models, read the same corpus, mobilise the same computing power, and produce objects of opposite natures. One delivers a score on a person. The other delivers an investigated file that a human being concludes and signs. The difference does not lie in the technology: it lies in where the line has been drawn, and in the rigour with which it is held.

Our line is the following, and it is invariant: **the machine never assesses a person · no decision is automated · the process is consented to and its findings returned · every conclusion is signed.**

This is not moral caution added after the fact. Three necessities of different orders converge on this rule, and it is that convergence that makes it solid.

The condition of the material. A leader does not open up to a system that scores him. He opens up to a practitioner who listens to him, within a process whose rules he knows, and from which he knows he will receive what concerns him. This obvious point has an architectural consequence that is rarely measured: the prohibition is not laid *on top of* the collection, it is *what makes the collection possible*. A chain that scores dries up its own source — it obtains front-of-house interviews, calibrated answers, everyone's presentable version. Put differently, the quality of what the machine will have to investigate is a function of what it has been forbidden. This is the most counter-intuitive reversal in the whole subject: the restriction does not amputate the system's performance, it produces it.

The legal protection of everyone. This is the subject of the preceding section and it need only be recalled: the line keeps the service on the right side of two regimes — that of high-risk systems, and that of automated decision-making — and it protects the client as much as the service provider, since exposure travels with whoever relies on the result in a determining manner.

The definition of what is being bought. An executive committee, a board, an investment committee do not buy a score. They buy a conclusion that someone signs and for which that someone will answer — before them, before the person assessed, before a court if it comes to that. A machine cannot sign, not for technical reasons but because responsibility is not a feature: it is a legal and moral relation between persons. One does not buy a probability; one buys someone who answers for a reading. The red line therefore defines the product.

It remains to translate all of this into a real chain, for a principle that does not descend into the specifications is no more than an intention.

The grammar of verbs. A *prompt* is not a way of talking to a machine: it is a specification. The chain is therefore written with the verbs of investigation — monitor, prepare, transcribe, structure, cross-check, corroborate, extract, verify, format — and never with those of verdict: assess, score, rank, predict, qualify, recommend. This is not a precaution of vocabulary. An instruction asking a model to "assess the chief executive's robustness" has built an assessment system, whatever the sales brochure may say, and the legal

qualification will follow the reality of the instruction, not its dressing. An instruction asking it to "extract every passage in which the chief executive describes how he decides, with references" has built an investigation system. Same model, same corpus, two products, two legal regimes. This is why the commercial formulations themselves carry consequences: marketing that implies the AI assesses can requalify a service the practice had, for its part, built correctly.

The strictest rule in the practice: no emotional inference by the machine. Here we deliberately place ourselves above the legal minimum. What is permitted is the analysis of **content**: what the person says, the facts she reports, the structure and coherence of her account, the convergences and contradictions between the interviews, the instruments and the documents. What is forbidden to the machine is the inference of an **affective or psychological state** — anxiety, tension, enthusiasm, dejection — from the voice, the face, the rhythm, the hesitations, from any para-verbal signal, and, under our house rule, from any channel whatsoever, including the choice of words alone. In practice: the transcript is verbatim and nothing else, with no annotation of tone and no affective marker; no prosodic, facial or video analysis; no sentiment analysis module anywhere in the practice, and any third-party tool that would embed one has it contractually disabled or is ruled out.

And when a model spontaneously produces a qualification of that nature — "seems uncomfortable", "defensive tone" — the passage is neither retained, nor reworded, nor quietly recycled: it is discarded, and the incident is recorded in the engagement log, with its date, the station of the chain concerned and the remedy applied. This detail looks minor; it is the only one that proves anything at all. **A rule whose breaches are not recorded is not a rule, it is an intention.**

Where, then, does human observation pass? Any observation of an affective or psychological order appearing in a deliverable is an observation made in interview by the practitioner, written and signed by her, and the deliverable documents its provenance. The machine never proposes a first draft in that register — there is no first-draft exception, because the first draft is precisely where the machine's judgment would enter the human's without a sound, under cover of saving time.

The serious objection: what if the human merely signs? This is the only objection that truly counts, and it must be faced, because the entire sector now invokes the human in the loop. The law itself does not ask for human intervention, it asks for *meaningful* human intervention — the qualifier does all the work. Research on automation bias establishes that an operator tasked with validating a machine proposal tends to accept it, all the more so as that proposal is well written, confident, and costly to contradict. A well-formed synthesis is a remarkably effective anchor. Asserting that "it is the human who decides" therefore proves nothing: it is precisely what has to be demonstrated, through the construction of the chain.

Three provisions see to it. **The order of operations, first:** the practitioner forms and writes her own reading *before* receiving the machine's structuring, never after. A judgment formed first can be corrected by the material; a judgment formed second is almost always a ratification. **The nature of what reaches the concluding stations, next:** the machine does not hand up syntheses but passages, with their references. A synthesis invites assent; a set of sourced quotations forces one to read again. The design deliberately makes the convenient path a little more costly than the right one. **The exposure of contradictions, last:** where sources

diverge, the machine's task is to report the divergence intact, never to harmonise it. A chain that smooths contradictions produces a comfortable document and an unusable diagnosis.

The control that follows is verifiable: the sentences that conclude must attach to sources the signatory has read again. That is what meaningful human intervention concretely means — not a box ticked at the end of the process, but an order of work in which ratifying costs more than judging. One reservation, in honesty, to finish: no design eliminates automation bias, it can only make it costly and visible. Whoever claims to have removed it has not looked for it.

The red line **is** the architecture. It is read in the instructions, in the logs and in the checkpoints, where any serious challenger can go and verify it.

5. Inside the practice — the division of labour, station by station

Consider the real calendar of an assessment engagement conducted the old way. The practitioner spends the majority of her time on tasks that demand no presence: reading documentation, taking notes, listening again, copying out, formatting, checking consistency. The scarce hours — the interview, the arbitration, the debrief — are squeezed by the abundant ones. The right design question is therefore not "how do we go faster", which is a supplier's question, but "how do we give the scarce hours the place they deserve", which is a professional's question.

Hence the principle of division: **the human carries everything that creates trust and commits judgment; the machine carries everything that consumes time without requiring presence.** Five stations, each with its explicit frontier.

Preparation. The machine reads the whole of the documentation, the project the team is supposed to serve, the sector, the public track records, and produces briefing files as well as personalised interview guides derived from the standard structured protocol. The frontier: the guide is a proposal, which the practitioner amends; and the structured protocol itself — the scientifically validated part of the system — is never modified by the machine, on pain of breaking what makes it valid. The effect is the first real gain of the frontier: arriving at an interview having read everything, which no craftsman can do.

The interviews. Conducted by the practitioner, recorded after prior and specific written consent, transcribed automatically. This is the strictest frontier in the whole chain, and it is deliberately asymmetric: what the machine receives of the interview is a text; what the practitioner has seen is a person. That asymmetry is not a technical limitation awaiting removal by the next generation of models — it is a permanent choice, and it is where the prohibition on emotional inference runs.

The analysis. Assisted coding of the verbatim transcripts across the six dimensions of the fuseki "human capital assessment", cross-checking against instrument results and documents, detection of convergences and — more valuable still — of contradictions. The frontier: the machine proposes, it does not conclude; it reports passages with their references, so that every element remains traceable back to its source. Recognising the regime the organisation is in — virtuous, broken, inverted — and the state of the transmission between governance and vision is an act of judgment, never a model output. The machine can say: here are the fourteen passages that touch on the relationship between the industrial director and the commercial director, and here is how they diverge. It cannot say what that relationship does to the company, or where to press in order to change it.

Production of the deliverable. Generation within the standard template, then rewriting for judgment. The frontier is at its subtlest here, and it is the one most often crossed inadvertently: the machine writes what **reports** — what was said, what was measured, what the documents establish — and never what **concludes** — what it means, what should be done, at which fulcrum to act. A rewrite is not a polish: the sentences that conclude are written by the person who signs them, failing which the signature no longer bears on very much.

Capitalisation. A governed methodological corpus — anonymised verbatim material, observed regularities, calibration — improves the next engagement. The frontier: capitalisation works only on anonymised material, under retention rules, and never feeds a third-party model, exclusion from training being contractual with every processor.

What this division changes, in one sentence: it alters the ratio between the hours of investigation and the hours of judgment. On the investigation side, days become hours; the hours thus returned go back to the interview, the arbitration and the debrief. Hence the answer to the objection that always comes — *if the machine produces, why this price?* The frontier changes the cost of production; it does not change the value of the verdict: fees do not pay for production, they pay for judgment, responsibility and signature. And a deeper investigation does not make the verdict cheaper to carry; it makes it better.

A word, finally, on how this description can be verified. The practice's corpus — its method, its studies, this document itself — is produced by this chain, under human direction and human judgment. 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. **Proof by deed, never by declaration.**

What the practice does not solve. The frontier shifts the cost of investigation; it does not shift the limits of the profession, and a document that kept quiet about them would not deserve to be read. It does not replace presence: what a person agrees to say, and what she lets be seen, depends on who is sitting opposite her — a badly conducted interview will be instrumented to perfection, and it will remain poor material. The chain amplifies the quality of the collection, in both directions. Nor does it correct a missed scoping: if the engagement has been scoped on the wrong question — and the presenting subject is rarely the real subject — the whole apparatus will investigate the wrong question with exemplary rigour. Depth of investigation is a multiplier of scoping, never a substitute for it. It does not make prediction certain: any reading of a human being remains probabilistic, and will remain so. A structured assessment improves the quality of a bet; it does not turn it into a certainty, and the service provider who suggests otherwise is preparing a disappointment, sometimes a dispute. Nor, finally, does it dispense with consented access to people: where that access does not exist — and there are situations where it cannot exist — the honest answer is that the engagement does not happen, not that it would be done on documents alone.

These limits are not usage caveats relegated to a footnote. They belong to the design on the same footing as the prohibitions: a system that claims to have lifted them has simply stopped looking at them.

6. Science at the Inputs, Judgment at the Integration

There is, in this activity, a temptation that must be discussed frankly because it is powerful and it sells well: announcing a proprietary instrument. An in-house index, calibrated by learning on thousands of profiles, presented as the scientific contribution of a practice that would have managed what sixty years of psychometrics supposedly failed to do. It is the direction in which the sector's technical energy spontaneously runs. We do not take it, and the refusal is reasoned.

The science of instruments is not reinvented by whoever has computing power. Research has established a hierarchy of validity that cannot be circumvented: the structured interview at the top of the methods for predicting performance, far above the unstructured interview whose main output is false confidence; the five-factor family for the description of traits, of which conscientiousness remains the most robust predictor; validated collective-level measures for the properties of the team itself. And, right at the bottom, the typologies popularised by the corporate away-day circuit, which do not meet elementary psychometric standards and have no place in a decision that commits someone's career. Building an unvalidated proprietary instrument, however technically ingenious, would amount to joining that last category with a better interface.

Artificial intelligence therefore industrialises the investigation **around** validated instruments; it adds none. Where it brings something real is not in the science, it is in the volume of the triangulation: systematically confronting every instrument result with the entirety of the interview corpus and the documents, instead of doing so by sampling and from memory. A craftsman triangulates what he remembers; a chain that investigates triangulates everything. The gain is considerable — and it is a gain in investigation, not in assessment.

A limit must be added that the profession states rarely, because it is commercially uncomfortable: **no instrument reads a collective**. Psychometric tests profile persons, one by one. A few instruments measure a property of the team itself — the quality of its enabling conditions, the psychological safety that prevails within it. None produces the systems reading: the real complementarities, the fault lines, the state of relations between the official company and the real company, the point at which a limited action would produce the maximal effect. *A team is not the sum of its profiles — it is a system; the Goban is how it is read.* That reading is a human act, and it is that act one buys.

The most current temptation. One practice now spreading must be named, and we refuse it: having a language model produce, from an interview transcript, a personality profile — scores on the five factors, inferred from the text. Technically this works, in the sense that the output is stable and convincing. Three cumulative reasons exclude it from our practice. The scientific reason: the validity of an instrument is not a property of its output, it is a property of its construction and of its validation on populations. A model that imitates the result of a validated questionnaire has imitated its form, not its validity — the number resembles a score, it is not one. The legal reason: this is very precisely the inference, by a machine, of characteristics of a person, which our line excludes, and depending on what is inferred it comes close to the prohibition set out above. The practical reason, finally, which would suffice on its own: the operation makes traceability

disappear. A score inferred from a text traces back to nothing verifiable; it can neither be contested by the person, nor re-examined by a challenger. It is, literally, the object nobody can sign.

The question to put to any system that produces a number about a person is, incidentally, the same one put to an analyst: where does it come from, and what would it take to contradict it? A number that cannot be contradicted is not a measurement — it is an opinion dressed as data.

One last discipline, less visible but structuring: sobriety. Every instrument added is one more data flow, one more processor, one more surface of risk, and one more burden on the people being assessed. One therefore assembles the shortest toolkit that covers the reading, never the most abundant — the apparent richness of a system is often the first symptom that it has not been thought through.

Science is chosen; integration is judged. The layer that can be bought from nobody is the second.

7. Provable Governance — Proof by Deed

Suppose a person assessed last year asks what has been done with her words. Or that a data protection officer wishes to verify that no automated decision was taken about her. Or that a director, before resting a governance decision on a report, wishes to know how that report was made. What can be shown to them? A promise, or a file?

The file exists, or the system is worth nothing. It contains the specific written consent to the recording, distinct from the agreement given to the principle of the assessment itself, accompanied by information on recipients, purposes and durations. It contains the record of the individual debrief: every person assessed receives what concerns her. This point deserves emphasis, because it is the best test of seriousness for an assessment system — well ahead of the sophistication of the method. A chain incapable of giving back to a person what it has produced about her is a chain that has produced something indefensible; and one observes fairly regularly that what prevents the debrief is not the confidentiality invoked, but embarrassment.

It further contains the contractual exclusion of model training with every processor, documented hosting, retention periods and effective purging at term. It contains the logs of the chain: instructions and processing operations archived by engagement, which make it possible to show — and not merely to assert — which verbs were used at each station. It contains the named checkpoint that precedes every delivery, whose object is verified and dated: no emotional or psychological qualification of machine origin in the document, and the incident recorded where applicable. It contains, finally, the use clause on the contractual side: the deliverable informs a decision, it does not make it — a stipulation that protects the person assessed, the client and the service provider together.

None of this is spectacular, and that is the sign that it is the right work. What will be noticed above all is that this list was not written by a lawyer after the fact: it overlaps almost exactly with what a scrupulous practitioner was already doing before the law took an interest. The regulatory framework says roughly the same thing as the professional ethics of the trade, with penalties added. Whoever built his practice around the requirement finds compliance already done; whoever built it around performance will have to rebuild, and will rebuild badly, because one does not make up with clauses for an architecture that has taken the wrong set.

This produces a market effect that must be named without circumlocution, since it grounds part of our position. Documented compliance is a barrier to entry. It offers no protection against the large platforms, which will know how to document it. It separates, within the existing offering, those who can show their file from those who will have to rebuild — the craft offering that has nothing to show, and the fast software offering that showed too early that it had built the wrong object.

It will be objected that all this remains unverifiable from the outside: a file can be well kept and the practice badly governed. The objection is fair, and the only answer consists in opening up. We therefore accept three things that are not customary in the sector: that the protocol be examined before the engagement by the client's data protection officer or by his counsel; that the chain be shown — shown, not described — station by station, to whoever asks; and that a challenger designated by the client be able to interrogate the method at the time of

the debrief. None of these three openings costs anything to whoever has built correctly. All three are impracticable for whoever has built otherwise, which makes it the best test available — including at our own expense, and that is precisely why it is worth something.

The double register — what the frontier promises, what judgment guarantees — is not a slogan: it is an auditable system. The frontier promises a depth of investigation beyond craft reach, a memory without gaps, exhaustive triangulation, a cost of production that funds judgment rather than replacing it. Judgment guarantees that no person is assessed by a machine, that the process is consented and given back, and that a human conclusion, reasoned and open to challenge, is signed by someone who answers for it. The first half of that sentence without the second describes a dangerous product; the second without the first describes an obsolete craft. They are quoted only together.

8. Conclusion — Governed Judgment

Let us return to the machine of the opening, the one that has read everything and says nothing about anyone. It is thought to be shackled. It is in its place — and it is that place, precisely defined, that makes the rest possible: depth without arbitrariness, speed without haste, memory without filing people away.

What this design promises those who decide — leaders, boards or investors — comes down to two propositions that hold together. An investigation nobody could have produced three years ago, because no human attention covered that volume within that time. And, at the end of it, a conclusion that someone signs and for which that someone answers — before the committee that decides, before the person assessed, before the court if it must come to that.

What this says about the profession now coming is broader. As artificial intelligence commoditises the analytical work on which the advisory professions had built their value, those professions will refound themselves on the one thing that does not automate: the responsibility of concluding. But that refounding will not come from virtue alone. It will come from design — and that is why we hold that a profession is appearing before our eyes, that of the designer of intelligence: the one who decides where it enters, what it opens, what it will never conclude, and who answers for what it produces. Our white paper no. 1, *The World That Is Coming*, argues this thesis for all organisations; this one gives its practitioner's version, in the profession where the consequences are paid for most dearly.

The question that will soon separate serious providers from the rest is already no longer "what can your artificial intelligence do?" It has two halves, and both must be put together: **what does your design make possible — and what have you forbidden it? Show me both.** It is a question any decision-maker can put tomorrow, to us as to anyone, and whose answer — a file, or an awkward silence — says far more than any demonstration.

We have written this document so that it serves beyond us. A leader, a human resources director, a board member about to entrust the assessment of his team to a third party can take it as a grid and put four questions: show me your instructions; who signs the conclusion; what does the person assessed receive; what happens when the machine slips. If the answers arrive with documents, the provider is serious — whether it is us or another. If they arrive with adjectives, the decision-maker has just learned something useful.

What this document describes for one profession holds, in our view, well beyond it. Any organisation that lets intelligence in makes the same decisions — where it enters, what it opens, what it will never conclude, who answers for what it produces — whether those decisions are taken or endured. Taking them is work; we argue elsewhere that this work is becoming a profession.

There is no judgment without someone to answer for it. Everything else — speed, depth, memory, exhaustiveness — the machine gives, and will give better and better. That, it will not. Which is why it is governed.

Colophon

Frontière — [frontiere.ai](#) — alice@frontiere.ai

Where to go next. The general thesis of which this document is the technical application is set out in **The World That Is Coming** (</en/papers/the-world-that-is-coming>) (white paper no. 1) — rationed intelligence, what its lifting opens, and the discipline that follows from it. Its application to decisions of capital is set out in **The Blind Spot** (</en/papers/the-blind-spot>) (white paper no. 2), which also contains the full development of the scientific apparatus drawn on here.

References. The work drawn on beneath the surface of this document is that which grounds the hierarchy of validity of assessment methods (Schmidt and Hunter on the predictive validity of selection methods; Barrick and Mount on the five-factor model), the reading of executive teams as systems (Hambrick and Mason on upper echelons; Wageman and Hackman on the enabling conditions of top management teams; Edmondson on psychological safety) and the analysis of organisations as feedback loops (Burns and Stalker; Forrester; Senge). The full development of this apparatus appears in our white paper no. 1, *The World That Is Coming*.