

WORKING PAPER

When Observation Loses Authority

Applied research emerging from executive advisory, organisational transformation and investment practice

Author	Alessio Virgili — Virgili Studio
Status	Working Paper — First Release
Version	1.0 — July 2026
Audience	Advisory firms · Executive search · Private equity · Venture builders · Knowledge-intensive organisations
Contact	virgilistudio.com

Abstract

Generative AI has dramatically increased the speed at which organisations produce text, analysis, decisions, and strategic documents. This expansion does not automatically translate into greater knowledge. When no structure distinguishes what has been observed from what has been interpreted, what is demonstrated from what is merely plausible, and what has been decided from what has only been generated, production accumulates faster than understanding does.

This paper documents how three distinct operational problems — the representation of complex professional authority, the identification of real opportunities from well-formed ideas, and the governance of knowledge through a growing body of generative output — revealed a shared structural weakness. Each required not a better tool, but a discipline capable of keeping observation, knowledge, decision, and generation rigorously separated. On this basis, Virgili Studio developed a practical infrastructure: the CV Engine, the Opportunity Engine, and the Virgili Core Engine — a canonical knowledge kernel governing what may be treated as authoritative, in which conditions, and why. This

discipline is today used operationally within Virgili Studio across advisory, positioning, investment evaluation, and venture development work.

The contribution is not the invention of provenance, knowledge graphs, or executive assessment. It is their operational configuration around a principle existing frameworks address only partially: an organisation does not accumulate knowledge simply by storing information or producing documents. It accumulates knowledge when what it learns can survive the project that generated it, retain its epistemic status, and be interrogated again without losing its connection to the reality from which it came.

1 The Problem That Precedes the Solution

Virgili Studio was not trying to build a research programme. It was trying to solve recurring operational problems. The research emerged because those problems repeatedly revealed the same underlying structure.

The organisations that work with knowledge — advisory firms, executive search practices, private equity teams, venture builders, family offices — share a structural tension that generative AI has made more visible without resolving it. The tension is this: producing more does not mean knowing more.

When producing a document cost substantial time, each document received proportional attention. When production approaches zero friction, the cost moves from creation to verification, from writing to classification, from formulation to the governance of what status a formulation is allowed to acquire. Most organisations have built tools for the first problem. Very few have built the discipline required for the second.

The consequences are observable and recurring. A preliminary observation is repeated across several documents and progressively acquires the appearance of a verified conclusion. An estimate loses its label as estimate. An example becomes a pattern. The most recent formulation of a claim becomes more authoritative than the source it summarises — not because anyone decided this, but because the system rewards linguistic fluency over provenance.

The question is not how to produce more. It is how to ensure that what is produced does not erode the authority of what was known before it existed.

Three problems made this tension concrete. They appeared, initially, to be unrelated. What they revealed — progressively, through their failure modes — was that they shared the same structural cause.

2 Three Problems, One Structure

2.1 Representing authority that exceeds chronology

The first problem emerged from the work of positioning a senior executive whose professional value did not reside in a single function or a linear career path. The value derived from the combination of capabilities exercised across different contexts: product architecture, manufacturing governance, commercial scaling, brand construction, category management, operational transformation — applied across more than two decades in roles simultaneously entrepreneurial, industrial, technical, and commercial.

The market for senior mandates presents a structural characteristic that compounds this problem. At the highest levels of seniority — Chief Product Officer, Operating Partner, General Manager, Board Advisor — a significant proportion of opportunities is intermediated by a limited number of executive search firms, funds, investors, and professional networks that identify and assess candidates before any formal mandate is published. The window through which the market can read a profile is narrow, and it opens only when the profile has been made legible in advance.

Four failure modes were identified. A Translation Failure occurs when a capability exists but the market reads it incorrectly. A Depth Failure occurs when the capability exists but not at the depth the mandate requires. A Domain Failure occurs when the required domain is genuinely absent. A Visibility Failure — the most frequent — occurs when the capability exists at the correct depth in the correct domain, but the market does not see it.

OBSERVED PATTERN · EXECUTIVE REPRESENTATION

Problem	A senior profile with 23 years of documented operational authority across knitwear, RTW, manufacturing, commercial scaling, and P&L governance is reduced, in CV form, to a chronological list of company names and job titles. The market cannot read the capability because the representation obscures it.
Consequence	Translation Failure and Visibility Failure coexist. The profile reaches the wrong interlocutors or does not reach the right ones at all. The value exists; it is not seen.
Response	Separation of entity from capability. The same verified corpus produces different outputs — CPO CV, Operating Partner profile, PE snapshot, search firm brief — selected and translated for each buyer type, without reconstructing the person each time.
Transferable principle	Any organisation whose expertise does not map to a standard functional title faces the same representation problem. The discipline separates what exists from what is seen,

and governs the translation between them.

2.2 Distinguishing a real opportunity from a well-formulated idea

The second problem emerged in a different domain: the identification and evaluation of entrepreneurial and investment opportunities. Generative AI has made it relatively straightforward to produce concepts, market scenarios, and possible business models. This abundance does not solve the more important problem — how to distinguish an interesting idea from an opportunity with sufficient commercial, industrial, and strategic potential to justify investment, time, and organisational construction.

The problem encountered was located before existing evaluation processes begin. The object to be assessed was not yet a company or a business case. It could be a daily friction, a change in human behaviour, an industrial capability not fully used, a material with applications not yet built, or an intermediate space between existing categories. A normal scoring model was insufficient, because scoring requires the object to already be defined.

OBSERVED PATTERN · OPPORTUNITY IDENTIFICATION

Problem	A well-described concept acquires the appearance of a validated opportunity. The quality of the narrative substitutes for the quality of the evidence. Investment of time and attention follows the story, not the proof.
Consequence	Resources allocated to ideas that were never opportunities. The problem may not even exist in the form assumed. The buyer who would pay for the solution has not been identified. The capability required to deliver it has not been verified.
Response	A methodological chain in which the idea is not the starting point. Observation → Requirement → Hypothesis → Evidence → Evaluation → Decision. Each step requires explicit documentation before the next begins.
Transferable principle	Any organisation that evaluates opportunities — whether for investment, product development, or market entry — can apply this sequence to prevent narrative fluency from substituting for evidence.

2.3 Preventing generation from rewriting its own sources

The third problem was not initially visible as an autonomous challenge. It emerged as the first two systems grew. With the accumulation of analyses, documents, stress tests, and versions, each new output tended to reinterpret the previous ones. A definition could be rendered more incisive and, precisely for that reason, lose the qualifications that had made it accurate. A hypothesis could be taken up in a subsequent document and progressively become a fact. A number introduced to summarise a corpus could be repeated without it being clear whether it derived from a verified source, a count, or a formulation created to reinforce the narrative.

The problem was not the obvious hallucination. It was the gradual modification of the epistemic status of information — a drift that no single step causes and that becomes visible only in accumulation.

O B S E R V E D P A T T E R N · K N O W L E D G E D R I F T	
Problem	Across a growing body of documents, the most recent and most fluent formulation of a claim acquires authority over the source from which it was derived. Estimates lose their labels. Hypotheses become facts. The conversation history becomes more authoritative than the verified corpus.
Consequence	The system learns from its own outputs. Each generation cycle moves further from verified reality. The organisation cannot distinguish, from the document it holds, whether a claim is observed, inferred, or constructed.
Response	An architectural boundary between generation and canonisation. AI assists in drafting and proposing. Every output remains draft assistance until a human accepts it and the package passes a governed kernel. Generated documents do not become sources by virtue of being approved or published.
Transferable principle	Any organisation using generative systems in cumulative workflows — where today's output becomes tomorrow's input — requires a structural, not merely editorial, boundary between what is produced and what is authoritative.

3 Research Method

The research developed through an iterative process grounded in operational practice rather than academic design. Recurring patterns observed in the first projects generated interpretive hypotheses. These hypotheses were progressively tested across subsequent projects belonging to different domains — executive positioning, opportunity identification, investment evaluation, advisory mandates, venture

construction, brand development. Hypotheses that did not show continuity were abandoned. Those that continued to manifest independently of context were formalised as structural principles.

The corpus includes the complete professional patrimony developed across Lunaria Cashmere — founded 2003, international luxury knitwear, 100+ wholesale accounts across 15+ markets, eight Italian manufacturing districts, full P&L, exit 2024 — and Oratio, an investor-backed ultra-luxury menswear platform built from inception across six runway seasons. It includes the applied mandate work of Virgili Studio across advisory, product development, capital, and venture domains. It includes the applied research conducted to build the three engines documented in this paper.

The method is qualitative and applied. It does not claim to produce a general theory of organisational decision processes. What remains to be validated through independent cases is the comparative claim: whether this discipline produces better decisions than alternative approaches, and by what measure. This paper presents the implemented solution and the reasoning behind it.

4 What the Problems Revealed

Entities are not capabilities

The companies, brands, and platforms built over time are not capabilities. They are the contexts in which capabilities manifested. Confusing entity with capability is the source of most translation failures in senior executive representation. When a profile says 'Founded and led Lunaria Cashmere for 21 years,' it names an entity. The capability that entity evidences — product architecture, manufacturing governance, international commercial scaling, category building, P&L ownership — is a different level of description, and it is what the market actually purchases. The entity provides the proof. The capability is what the buyer acquires.

Preserved information and available knowledge are not the same thing

Preservation concerns the survival of information. A document can be archived and made accessible. Continuity concerns the persistence of the capacity for response that the knowledge contains — the ability to understand, judge, and apply it when the context changes. A procedure can be perfectly documented and prove insufficient when reality presents an unforeseen configuration. A decision can be preserved without preserving the judgment that had connected the circumstances to the choice made. Generative AI makes this tension more visible: it can produce a plausible response even when the distance between available content and specific reality has not been recognised.

Production is not the same as accumulation

Most strategic and entrepreneurial work produces episodic knowledge. A project is completed, a document is delivered, the team moves on. When a similar problem arises, the organisation rarely restarts from its actual previous patrimony — it restarts from a reconstruction of it. Knowledge becomes cumulative when it satisfies four conditions: it survives the occasion that generated it; it can be reused without being separated from the conditions that limit its validity; it preserves its provenance and epistemic status; and it increases the capacity to handle future cases, not only document past ones.

5 The Infrastructure Built

5.1 The CV Engine

The CV Engine is not a system that automates CV writing. It is a system that builds a verifiable representation of professional authority and transforms it into different outputs without reconstructing the person from scratch each time. The Capability Pattern separates entities from capabilities and maps recurring combinations across the full professional corpus. The Value Creation and Evidence Library catalogues, for each capability, the documented evidence — outcomes with sources, metrics with context, transformations with conditions. The Problem-to-Proof Matrix connects buyer problems with the proofs the corpus can provide.

The output layer produces documents calibrated for different interlocutors — an executive CV for a specific mandate, a CPO profile for search firms, an Operating Partner snapshot for private equity, a B2B document for advisory contexts — selecting, organising, and translating without rewriting. The Outcome Tracking Framework connects candidacies, responses, shortlists, and mandates to the way the profile was represented, producing over time the first real dataset: not a theory about how the market works, but evidence of how it actually responded to specific configurations.

5.2 The Opportunity Engine

The Opportunity Engine intervenes before the opportunity is already a definable object. The Observation Library collects observation protocols by domain — not queries or trends, but structured ways of recording a friction, a behaviour, a change of context, or an unused capacity without immediately producing a concept. The Requirement Engine transforms observations into requirements when a subject exists for whom the friction is real and conditions emerge that distinguish an episode from a pattern.

Only then does the system construct an opportunity hypothesis, connect it to evidence, compare it with available capabilities, potential buyers, channels, acquisition models, foreseeable objections, and development conditions. The engine also preserves rejected opportunities: the reasons for exclusion,

missing evidence, and conditions that could make them relevant again. A negative decision contributes to the system's knowledge rather than disappearing with the end of a conversation.

5.3 The Virgili Core Engine — a canonical knowledge kernel

The Virgili Core Engine (VCE) is the architectural response to the third problem and the common foundation beneath the first two. The decisive step occurred when the CV Engine, the Opportunity Engine, and the broader analytical work of Virgili Studio were no longer treated as autonomous systems. The question became: what functions remain necessary regardless of domain?

The kernel governs three primary object types: METRIC, a quantity defined through an identifiable criterion; PROOF, an evidence that supports, qualifies, or contradicts a claim, with a traceable connection to a source; and PROJECTION, a description of a future outcome or possible transferability, explicitly labelled as such. The relations between these objects are governed through a validation layer that prevents unsupported relationships from entering the canonical corpus. A critical principle governs the boundary: documents produced for external use do not become canonical sources simply because they were approved or published. To re-enter the system, they must pass through a new governed intake process. This prevents the system from learning from its own formulations.

The Kernel V1.0 is implemented, tested, and certified as a technical baseline, passing 174 tests. The architecture is not a logging system. It is a governance system with explicit lifecycle management for canonical objects, policy enforcement at the boundary between observation and canonisation, and structural separation between generation and authority.

6 What This Contribution Is and Is Not

The construction documented here does not claim to have invented provenance, knowledge graphs, executive assessment, or opportunity scanning. These domains exist and are developed by institutions, universities, technology companies, and international consulting practices. The contribution is their operational configuration around a starting point that most enterprise architectures do not address.

Enterprise knowledge architectures typically begin from organisations that already possess data, taxonomies, roles, processes, and decision criteria. Virgili Studio began from a prior condition: the experience existed but was not yet represented as governed knowledge. The capabilities had been exercised but not separated from the companies and roles in which they manifested. The evidence was distributed across documents, conversations, projects, and memory. The method existed in practice but had not been rendered independent of any particular domain.

The result is an implemented operational sequence: a career into a corpus of authority, capabilities, and evidence; the corpus into differentiated market outputs; observations and frictions into evaluable

opportunities; the repetition of the method into a domain-independent kernel; generation into a process subordinate to canonical knowledge. The distinction most important to state: existing provenance research seeks primarily to reconstruct from where an output came. The Virgili Core Engine seeks also to govern what status that output may acquire within the organisation. These questions are connected but do not coincide.

7 Transferability

The discipline is transferable not because it can be replicated exactly in different contexts, but because it addresses a structural problem — the distinction between observation, knowledge, decision, and generation — that appears in any organisation that must build reliable representations of complex realities.

Executive search and leadership advisory

The evaluation of a candidate requires the construction of a representation of professional authority, acquired capabilities, evolutionary potential, and coherence relative to a specific context. The same failure modes documented here — Translation, Depth, Domain, Visibility Failure — apply to every complex profile that does not correspond to a standard functional path. The CV Engine provides a model for keeping capability, evidence, and context of transferability structurally separate, producing different outputs for different buyers from the same verified corpus.

Private equity and operating partners

Investment decisions are made on the basis of representations built progressively. The Opportunity Engine addresses the phase before the object to be evaluated exists as a company or business case. The Kernel provides a framework for connecting an investment hypothesis or value creation thesis to evidence, metrics, capabilities required, and conditions that support its validity — while keeping projections explicitly labelled as projections.

Venture builders and innovation teams

The chain from observation to requirement to opportunity to evidence to decision provides structure for the phase that most innovation processes manage informally — the point at which a signal claims to be an opportunity. The discipline does not eliminate intuition; it prevents intuition from automatically becoming an investment thesis without passing through a verification sequence.

Advisory firms and consulting practices

Each mandate produces documents, analyses, and recommendations. The question the Kernel addresses is whether the knowledge produced in one mandate is available — in verified, provenance-carrying, epistemically labelled form — for a subsequent mandate. The governance discipline provides an architecture in which knowledge produced in one context can be reused in another without losing its connection to the conditions under which it was generated.

When the answer to these questions is no, the organisation produces episodically. When the infrastructure exists to answer yes, it accumulates. The difference, compounded across projects and years, is substantial.

8 Limits of This Research

The research corpus is qualitative and single-organisation. The patterns identified emerged from a specific professional history and were tested across a range of mandates and opportunity contexts within that history. They have not been validated through independent organisations, controlled experiments, or comparative studies against alternative methods.

The Kernel V1.0 demonstrates that the architecture exists and its primary contracts have been verified. It does not demonstrate comparative superiority over alternative knowledge governance approaches. Whether the capability-and-evidence representation increases the quality of interlocutors, shortlists, mandates, and compensation over time remains to be measured against a sufficient dataset of real candidacies.

The limits are stated here because the discipline documented in this paper is founded on the principle that claims must be traceable to their evidence and projections must remain visible as projections. Claiming comparative superiority without the evidence to support it would be an instance of exactly the failure mode the Kernel is designed to prevent.

9 Conclusion

Organisations are not entering a period characterised by scarcity of content. They are entering a period characterised by scarcity of distinction.

It will become increasingly easy to produce a synthesis, a strategy, a recommendation, a market thesis, or a professional representation. It will become more difficult to establish what, within that

production, merits trust; what derives from a source; what has been inferred; which decision modified the patrimony.

The three problems from which this work arose appeared different. The first concerned the possibility of making complex professional authority transferable without reducing it to a chronology. The second concerned the possibility of transforming observations into evaluable opportunities without confusing narrative plausibility with evidential quality. The third concerned the possibility of using generative systems in a cumulative process without allowing generation to rewrite the knowledge on which it depended.

The common response was to separate four functions: observe, know, decide, generate. Observing means recording reality without attributing a definitive explanation prematurely. Knowing means elevating some information through provenance, verification, relationship, and governed acceptance. Deciding means applying judgment to a specific configuration and preserving its reasons. Generating means translating the available patrimony into a useful form, without automatically acquiring authority over it.

The CV Engine, the Opportunity Engine, and the Virgili Core Engine are three applied responses to this same structural challenge — today in operational use at Virgili Studio, and transferable to any organisation encountering the same problem: the management of knowledge across a growing body of generative output, the representation of complex authority, the identification of real opportunities before they have been defined.

The primary limit of generative systems is not their capacity to formulate responses. It is the tendency of organisations to mistake production for knowledge. An organisation does not know more because it generates more documents. It knows more when what it observes can become reliable without becoming rigid; when what it decides preserves its connection to the reasons that made it possible; when what it produces does not cancel the difference between reality and representation. It is in this continuity — not in the volume of outputs — that knowledge becomes infrastructure.