Contextual Residue: What the Industrialisation of Knowledge Does Not Transfer

A minimal axiom, five layers of residue, and the embedded deployment of AI in regulated systems as the domain of demonstration.

This corpus has produced, over recent months, five distinct observations on the embedded deployment of artificial intelligence in the enterprise. They were established separately, on separate ground, none of them claiming to explain the others. This note advances a hypothesis contrary to their dispersion: they are not five results but five manifestations of one, produced by a single cause. That cause is neither AI, nor the cloud, nor the strategy of any one vendor. It is a property of every industrialisation of knowledge. The reasoning does not depend in any way on the actual performance of the models; it bears on the economic and organisational logic of the deployment apparatus, not on what the models can do.

Five anomalies, without a common cause

The first anomaly is temporal. The Dual Value Horizon Model establishes that the vendor of an embedded deployment recognises its return within a window that saturates near go-live, whereas the client's net return only reveals itself later, if it is to reveal itself at all, once the full cost of permanence has been observed and the benefits causally attributed. Two parties, one deployment, never the same window.

The second anomaly is institutional. An agent that decides within a process may have a perfectly identifiable operational engagement without any authority having subscribed ex ante to the consequences of its decisions.

The third anomaly is patrimonial. The asset actually deployed is not the model. It is the operating capability constituted at the client, distinct from documented competence and from the delivered tool.

The fourth anomaly is metrological. Adoption at go-live is not deployability, and billed consumption is not demand. Lock-in makes consumption persist independently of the value produced.

The fifth anomaly is strategic. When the logic of platform return on investment seizes up, an infrastructure vendor responds with a deployment apparatus rather than with a feature.

Five domains, five anomalies, one author. The temptation would be to write that the corpus proves them. It proves nothing: it observes. The question is therefore not to demonstrate each anomaly again, but to explain why they appear **together**. A theory does not ground the observations that precede it. It accounts for their co-occurrence.

Two definitions, before any demonstration

Two terms carry the entire argument. Leaving them to float would let the text say whatever the reader projects onto it.

To industrialise knowledge means here to produce a representation stable enough to be replicated at scale independently of its context of origin. This definition separates four operations that discourse conflates: to standardise, to deploy, to replicate, to scale. Only the first is at issue here, the other three being its operational consequences. What matters is the claimed independence from the context of origin. Without it, there is no industrialisation, only displaced craft.

Contextual residue designates the set of determinants of a decision that cannot be captured, transferred or represented by the abstraction used to industrialise a body of knowledge. Residue is therefore not missing information that could be collected later. It is what the representation, by virtue of the very thing that makes it replicable, cannot carry.

These two definitions hold for the whole text. They will not be restated.

The axiom, and nothing more

A single statement, as primitive as possible.

Every abstraction loses the context it does not represent.

The statement is information-theoretic before it is organisational. To abstract is to retain a representation and discard the rest. What is discarded is not destroyed; it ceases to be carried. The statement speaks neither of AI, nor of embedded deployment, nor of liability. It is credible to anyone who has never heard those terms. That is what one expects of an axiom: that it be short, nearly self-evident, and that it carry more than itself.

Its founding intuition is not mine. Herbert Simon showed that organisations do not handle complexity by representing it in full, which their bounded rationality forbids, but by constructing hierarchical and deliberately incomplete representations, carving the real into nearly independent subsystems. Near-decomposability is not a property of the world. It is what the organisation imposes on the world in order to process it. The axiom states the price of that imposition: what does not enter the carving is not absent from the real, it is absent from the representation. Simon does not come along to prove the axiom after the fact. He explains why an organisation that industrialises has no option but to produce residue.

An objection must be raised here, before a reader formulates it, because it is the only one that could kill the text. If the axiom were merely an abstract restatement of the anomalies it claims to generate, the theory would be circular. The safeguard is twofold. The axiom must be statable without any of its consequences; it has just been. And every link in the derivation must import a premise external to the axiom. Were a corollary to follow by mere renaming, the axiom would fall. The reader is invited to verify the entry of each external premise.

From residue to the local actor, and from the local actor to the client

To industrialise, in the sense defined above, is to abstract at scale. By the axiom, that abstraction loses context. The lost context is the residue, and it is local by nature since it **is** the unrepresented context, attached to the place and the moment that produced it.

From this follows a theorem, to be distinguished from the axiom on pain of confusing the cause with its manifestation. For an activity whose value is contextual, the residue cannot be reduced to zero without cancelling the abstraction, hence without cancelling the industrialisation itself. One cannot standardise an activity and eliminate its residue in the same gesture. This property, *contextual irreducibility*, is not an axiom but what the axiom produces when the value of the activity lies in context. The distinction separates *decomposing without loss* from *decomposing with remainder*, and everything that follows plays out in the second term.

Since the residue can be neither standardised nor eliminated, some actor must carry it, and that actor is necessarily local, because the residue is. First external premise: the theory of incomplete contracts, from Grossman and Hart to Hart and Moore, establishes that when future states cannot all be specified, residual control rights determine who decides in unforeseen situations. The actor who carries the residue holds those rights de facto, being alone in a position to arbitrate what the representation did not anticipate.

One must stop here, because the next step is where the demonstration could substitute words for proof. That the actor is local does not say that the actor is the client. A systems integrator, a managed operator, an industrial partner, an operating subsidiary or a regulator may occupy that position. Identifying the local actor with the client is not a consequence of the axiom: it is an additional hypothesis, true within the domain studied.

In the embedded deployment models observed, where the vendor takes neither operating responsibility nor residual control rights, the local actor is the client. Elsewhere it will have to be shown rather than assumed. This restriction does not weaken the theory; it delimits its field, which is not the same thing.

Five layers of residue

A concept that explains everything explains nothing. Residue, as just posited, would account for return on investment, governance, liability, contracts, value capture, deployability and break-even. That fecundity is suspect: it signals an omnibus concept, and an omnibus concept is hard to refute, therefore uninformative.

The remedy is not to reduce the ambition. It is to stratify. Residue is not a single substance manifesting under five disguises; it is a chain whose every link produces a remainder of a distinct type, with its own premise and its own falsification condition. The stratification is therefore not a typology appended after the fact to save the theory. It is the chain of derivation, read across.

Informational residue is what does not pass into the representation. It is residue in the strict sense of the axiom, and the only one that follows from it directly. It would be refuted by a representation of contextual knowledge that discards nothing, which would amount to abolishing abstraction.

Organisational residue is what remains to be coordinated locally. It arises from co-evolution: an architecture of knowledge evolves continuously with processes, data, teams, regulation and objectives, and each local evolution regenerates a remainder. The problem is not *knowing* the business, which a vendor's experts are often perfectly able to do, but *evolving with* the business, which does not replicate across thousands of simultaneous engagements. Domain depth does not run up against competence, it runs up against the scale of co-evolution. Falsification condition: an apparatus that would supply, per engagement and across thousands of simultaneous ones, a depth equivalent to that of the client's internal experts.

Decisional residue is what remains to be arbitrated. It is organisational residue seen from the side of action: unforeseen situations call for a decision that only the holder of residual control rights can take. Two movements must be distinguished here, which dominant discourse conflates. The vendor externalises an *operational* risk by placing its engineers in contact with the real. The client internalises a *decisional* risk, that of unspecified situations. The second is demonstrable, and it is probably what produces, seen from outside, the impression that the vendor is offloading. It is not the vendor abandoning a risk; it is the client inheriting another.

Legal residue is what remains to be assumed institutionally. To sign is to subscribe ex ante to consequences; yet the consequences that matter are those the representation did not specify. A signature bearing on the residue does not standardise, since the residue does not specify. It is therefore omitted. One term must be fixed here, as two neighbouring notions are commonly confused: an organisation may be *liable* without having signed, and may *sign* without being wholly liable. The term retained in this text is the *institutional signature*, understood as the act by which a designated authority converts a technical delegation into a commitment of the institution. It is the signature that is missing, not the liability, which always ends up being attributed, generally by a court and generally late. Falsification condition: a deployment whose contract designates, ex ante, the institutional signatory for the behaviour of the agents.

Economic residue is what remains to be captured as value. What is industrialised is measured, and what is measured becomes the proxy on which return is computed: consumption, billing, metered usage. These proxies are standardised, hence captured early. Net value resides in the residue, hence reveals itself late. External premise: Teece showed that creating value and appropriating it belong to two distinct regimes, governed by appropriability and complementary assets. Complementary assets are precisely the mechanisms by which an actor captures a share of the value produced by a residue it does not directly control. The vendor does not hold the residue; it holds the infrastructure without which the residue produces nothing. A reservation is in order: this does not establish that the vendor captures value at the client's expense. Both may gain. What the structure renders indeterminate is the *distribution* of the value created, not its existence. The inequality between the two windows of the break-even model is not a postulate, it is that indeterminacy seen over time. Falsification condition: a demonstration that billed consumption tracks symmetrically the net value appropriated by the client.

Five layers, five premises, five falsification conditions. None is obtained by renaming the axiom. The theory is therefore not circular, subject to the reservation that its external premises are themselves debatable, as every premise is.

The observable trace

A theory of residue must be able to show it somewhere. There is a trace, and it calls for a caution the subject does not spontaneously invite.

Proponents of embedded deployment speak abundantly of governance: trust, tooling, responsible AI, agent management. They never name, ex ante, the institutional signatory for the behaviour of the agent deployed in a regulated process. Microsoft Frontier Company, announced on 2 July 2026 with two and a half billion dollars and six thousand people embedded at clients, insists on measurable outcome and on trust without attributing to anyone the legal responsibility for the decisions produced. This occurrence is an illustration, not a demonstration: the theory would hold without it, and that is precisely how an illustrative example is recognised. It has the merit of being recent, documented and public.

The silence is consistent with the theory. It is what the theory predicted would be found: legal residue does not contractualise, therefore does not get stated. What that silence does not establish must be named. It establishes no intent. Three competing explanations account for it just as well: commercial simplification, since legal matter weighs down a pitch; contractual prudence, since one avoids committing in writing to a liability; immaturity of the framework, since no one yet knows, in law, who signs for an agent. None has been eliminated. The silence is therefore an empirical signature compatible with the theory, an indication, never a proof. A proof would presuppose the contrary, and the contrary is not established.

That modest status is what makes the indication useful, because it comes with its condition of refutation. Let a single proponent of embedded deployment publish a communication designating ex ante the institutional signatory for the agents it deploys, and the indication collapses. Its absence to date demonstrates nothing; it is merely consistent with a legal residue left to the local actor.

Domain of validity, conjecture, prescription

What this note demonstrates bears on one domain: the embedded deployment of AI in regulated systems, where the five anomalies were observed, where the chain holds, where the local actor is the client. What the axiom licenses one to conjecture goes further. Every industrialisation of a cognitive service abstracts contextual knowledge, therefore produces residue, therefore leaves something to a local actor. Consulting, software publishing, medicine, law and education ought, if the axiom is sound, to exhibit the same stratification. This is not a result but a research programme, testable and not yet tested. A general theory proved on a single case would be an essay. Stated general, proved on a single case, the remainder posited as conjecture, it is a programme. I claim it as such, and no more.

The limits must be set without evasion. The axiom does not forbid embedded deployment; it forbids embedded deployment *without residue*, which is not the same thing. An apparatus with strong capability transfer brings the client's break-even nearer and reduces organisational residue without ever cancelling it: residue is minimised, it is not suppressed. A managed operator remunerated on persistent performance does not refute the theory, it displaces the local actor. Identifying that actor with the client remains a domain hypothesis, not a result. Finally, the axiom rests on an information-theoretic intuition this note has not formalised; it treats it as a point of departure, not as a demonstrated theorem, and this is the most attackable point of the whole.

From this follows one prescription, architectural and not moral. Legal residue is not filled at the level of governance, where the tools operate, but at the level of signature, where no one operates. That is therefore where to act. A contract of agentic governance must designate the institutional signatory ex ante, as a precondition of deployment and not as documentation produced after the fact. Designating the signatory does not eliminate the residue. It names who carries it, which is the only operation available. That is the difference between a risk assumed and a risk inherited.

You can industrialise a representation. You cannot industrialise what made it right. That remainder has a name, a bearer and a price: the residue, the local actor, and the commitment no one has signed.