The Client's Deferred Break-even: a Structural Asymmetry of FDE Models

Why supplier and client do not observe their return in the same window

Introduction

In the summer of 2026, several major infrastructure and model providers announced or reinforced arrangements of engineers deployed at the client's point of contact. The sums committed run from hundreds of millions to a few billion dollars, the announced headcounts in the thousands. The *forward deployed engineer*, which a data-analytics software vendor turned some twenty years ago into a distinctive model in the service of public agencies, has become one of the most visible industrial models for deploying AI in the enterprise.

The ambient discourse treats this arrangement as a *delivery* model, a more intimate way of putting a system into production. That is a surface reading. The real subject is not the typology of FDEs. It is an asymmetry: supplier and client do not observe their return within the same temporal windows. The question that decides value is therefore not *how* one deploys, but who finances the deployed engineer, by what mechanism its cost is recovered, and at what moment the client's net return can be established.

A point of vocabulary must be settled at the outset. For convenience, we speak of ROI. In reality, a supplier first recognizes revenue, a margin, a learning, or a strategic option. ROI appears only once that value is set against the full cost of the intervention. The confusion between *return* and *ROI* is not neutral. It lets what is still only a receipt be presented as a yield.

We propose here a general framework, the **Dual Value Horizon Model (DVHM)**, of which the FDE is only the first field of application. The DVHM describes a property of sociotechnical systems in which the supplier's value capture and the client's value demonstration obey distinct, desynchronized temporal horizons. The FDE financing mechanisms are not the object of the model. They are its most legible illustration.

The thesis holds in two sentences. The recovery mechanism sets the window within which the supplier can recognize its return. The client's net return belongs to another window, independent and most often later, because it can be established only after observing the *cost of permanence*, that is, the cost of operation, governance, supervision, adaptation, and reversibility over the whole horizon, and after causal attribution of the benefits to the deployment.

The reason is structural, and rests on the incompleteness of contracts developed below. The client retains the control rights on which the persistence of the result depends, so that its return reveals itself only in use. The domain of validity is explicit. Return here designates the net financial return after full cost. The scope is that of the platform FDE, a technical and product team operating at the client's operational point of contact for a supplier whose return depends on the adoption of its product or ecosystem. This excludes independent consulting with no product interest.

The FDE is not merely a delivery model

Three questions precede any discussion of value. Who pays or employs the deployed team. On which actor the cost ultimately falls. Who is exposed to loss if the mission fails. They can be broken down into six reading variables, taken up in synthesis below: immediate payer, economic bearer, recovery mechanism, captured asset, return horizon, and exposure to failure.

Four mechanisms, and two trades

Four recovery mechanisms coexist. They engage neither the same payer nor the same captured asset. They differ in register, financing, governance, and legal structure. But they all read through a single variable, the way the supplier recovers the cost of the intervention, because it is that variable which sets the window where the return can be recognized.

The first finances the FDE as a **product cost**, close to quasi-R&D. The supplier bears the intervention on its books because the learnings flow back into the platform. The client's ground becomes a shared product-learning environment, whose value is principally capitalized by the vendor, even though the client receives adaptation, integration, and accelerated resolution. The counterpart is an integrated platform that lowers the cost of coordination while raising the cost of exit. One cannot say, for lack of comparable public data, which of the two prevails. An analytics vendor whose services share has declined over several fiscal years offers a compatible illustration, without on its own establishing the mechanism.

The second finances the FDE as an **accelerator of consumption or licensing**. A cloud consumption commitment decrements on effective usage and is lost if not consumed by the end of the term. The deployed engineer can then help convert an already contracted commitment into effective usage. Its operational function is not reducible to that financial mechanism. It can also accelerate *time-to-value* and reduce adoption failures. The supplier's economic mechanism and the function rendered to the client are two distinct readings of the same arrangement.

The third finances the FDE through a **dedicated structure**. An investment vehicle carries the deployment and distributes a yield to its investors. This is the most financialized case. It requires above all not to confuse four actors: the investor's yield, the vehicle's return, the technology supplier's benefit, and the value captured by the client do not coincide.

The fourth finances the FDE as a **billed services activity**. The client pays the team or the implementation package directly. This is not consulting rebranded, provided the supplier's return remains attached to the adoption of its product. The criterion is not the billing. It is the product interest, that is, here, stimulating consumption of the underlying product.

These families do not classify suppliers. They decompose mechanisms that a single mission may combine. Read through the six variables, they summarize as follows.

Mechanism	Immediate payer	Economic bearer	Captured asset	Horizon	Principal exposure
Quasi-R&D	Supplier	Supplier then future clients	Product learning	Medium term	Supplier
Consumption or licensing	Supplier or global contract	Client via consumption	Usage and commitment	Short to medium term	Shared
Dedicated structure	Investors and vehicle	Vehicle and clients	Distribution, participation, yield	Medium to long term	Investors and clients
Billed services	Client	Client	Service revenue and adoption	Short term	Client, per contract

A fifth actor traverses these families without constituting one: the integrator. It sometimes co-finances part of the deployment, captures its own margin, and above all displaces the bearer of the cost of permanence. Depending on the arrangement, the integrator carries the capacity durably through a managed-services contract, which makes it available without being internal. Or it withdraws it with the end of its mandate, leaving the client without the competence it believed acquired.

The supplier, moreover, is rarely a single actor. A vendor, its cloud platform, a code forge, a model provider, an integrator, and a services practice may share value creation. One must then reason in terms of a *capture chain* rather than a single supplier. Each link has its own recognition window, which does not reduce the asymmetry but multiplies it.

A second distinction, orthogonal to financing, decides the share of benefit actually attributable to the deployment, because not all FDEs do the same job. A first profile analyzes the client's data to reconstruct its ontology and discover risks within it; its ground is data. A second profile presupposes knowledge of the client's industry,

competitive landscape, digital maturity, processes, and flows, in addition to the technical skills of deployment, in order to propose a transformation of scale. This profile is rare.

A third case deserves to be isolated without being overinvested. A profile that would additionally master the regulatory and legal aspects of the deployment would render to the client a capacity it does not always hold internally. This case is structurally unlikely in the consumption mechanism: the supplier's economic incentives are not aligned with a maximization of legal constraints liable to slow adoption. This is not a trial of intent. It is a reading of incentive alignment: the recovery mechanism selects the profile, and a profile optimized for adoption is not optimized for regulatory friction.

Field observation confirms this reading without proving it. In several vendor services organizations, the dominant competence is technical or pre-sales, not deep business knowledge; sector advisors more often fall under *consultative selling* than under enterprise architecture applied to AI. This gap is not a detail of resource management. When a deployment meant to produce a sectoral transformation is entrusted to a pre-sales profile, the benefit actually attributable to the FDE diminishes, and the capacity transferred to the client with it.

Why do these distinctions matter. Because each mechanism sets a different window within which the supplier can claim to have created value. And because each profile decides the share of that value which will actually return to the client.

Four questions about the return

The word *return* covers four distinct properties, which it is better to name in order to reuse them: recognition, realization, persistence, attribution. Is the return *recognized* in the supplier's metrics, revenue, margin, estimated net present value. Is it *realized*, collected or cost avoided. Does it *persist*, after the team withdraws and governance costs are integrated. Is it, finally, *attributable*, actually caused by the FDE, established against a reference scenario.

Attribution is not a fourth step after persistence. It is a causal property that runs through the other three and can be interrogated at any moment. A gain can be recognized, realized, and durable without being caused by the FDE, simply because it would have occurred anyway. A collected return is not yet an attributable return. And without a reference scenario there is no attribution at all: any ROI demonstration without a counterfactual is a narrative, not a measurement. This requirement protects the analysis against before-and-after *business cases* that attribute every observed change to the deployment.

These four readings hold for each actor. That is what ruins the too-convenient opposition between a supplier that would gain early and a client that would gain late. The thesis does not say that the supplier always recognizes its return first. An analytics vendor may

recognize it very late; a dedicated vehicle as a function of future consumption; a hyperscaler several years after its commitment. It says that the recognition windows are *independent*, each commanded by its own recovery mechanism. The most frequent configuration remains the one in which the supplier recognizes part of its return before the client has been able to observe its full cost.

Two value functions, two temporal horizons

The asymmetry between supplier and client does not rest only on their economic interests. It rests on the fact that they do not observe their return over the same horizon. The supplier seeks to capture rapidly the value created by the deployment. The client can establish its net return only after having borne the full cost of operation.

The preceding analysis generalizes into a simple model describing the desynchronization between the supplier's value capture and the client's value demonstration. This is the DVHM announced in the introduction. It does not describe a particular case of FDE, but a property of sociotechnical systems in which creation, recognition, and demonstration of value obey distinct horizons.

The value captured by the supplier is represented by a function $CVF(t)$, which aggregates the forms of value it recognizes over time.

$$CVF(t) = REV(t) + LRN(t) + ECO(t)$$

where $REV(t)$denotes the revenue effectively recognized (licenses, subscriptions, services, contracted consumption), $LRN(t)$the economic value of learnings reusable to the supplier's benefit (platform improvement, acceleration of future deployments, reduction of development costs, competitive advantages), and $ECO(t)$the expected value of future usage induced by the deployment (cloud consumption, tokens, additional licenses, cross-sells, *churn* reduction). This function is not an accounting identity. It is an economic decomposition of the principal capture mechanisms, expressed in a common unit independently of the mode of financial recognition.

The dynamics of this function illuminate the whole thesis. $CVF(t)$grows sharply around the deployment, then its slope collapses. The reason is economic, not mathematical. Each of its components sees its marginal contribution attributable to the deployed engineer's presence decline once the system is in production. Contractual revenue is recognized or secured. The consumption commitment is already contracted. The exploitable learnings of the mission have already flowed back into the platform. Ecosystem value continues to run, but no longer depends on on-site presence. In other words, it is not the value that ceases; it is the share of value attributable to the FDE that saturates.

There is therefore a horizon T_F, generally close to production go-live, beyond which the marginal value directly imputable to the FDE becomes small.

$$T_F = \inf\left\{ t \;\middle|\; \frac{dCVF(t)}{dt} \approx 0 \right\}$$

T_F marks the saturation point: the moment when the derivative vanishes not because the supplier stops earning, but because the deployed team no longer brings marginal value. The essential of what the supplier draws from the deployment is already recognized or secured when the team withdraws. T_F is not a universal constant. It depends on the recovery mechanism chosen. The formalization does not seek to fix its value, but to distinguish its logic from that of the client break-even.

The client's trajectory is different. Its net value begins negative, since it immediately bears the cost of the project and of capacity constitution, then evolves under the effect of genuinely attributable benefits and of the cost of permanence.

$$NVC(t) = -(P + IC_0) + \int_0^t [\, BA(\tau, C(\tau)) - O(\tau) - K(\tau, C(\tau)) - X(\tau) \,]\, d\tau$$

where $BA(\tau, C)$ denotes the economic benefits causally attributable to the system, established against a credible counterfactual scenario; $O(\tau)$ the recurring technical operating costs; $K(\tau, C)$ the costs of governance, human supervision, compliance, and auditability; $X(\tau)$ the costs of adaptation, remediation, reversibility, and exit; P the initial project cost; IC_0 the cost of constituting the organizational capacity that lets the system be operated durably.

The nomenclature is kept stable throughout. In particular, reversibility is written $X(\tau)$ and not $R(\tau)$, in order to avoid any collision with supplier revenue $REV(t)$. One symbol, one meaning.

Reading the sign of the integrand explains the delay of the break-even. Early on, BA is small: adoption is ramping, benefits are neither fully realized nor yet attributable against a counterfactual. At the same moment, $O + K + X$ is high: governance, supervision, and adaptation weigh most in the early period. The integrand is therefore negative or weakly positive, and $NVC(t)$ stays below zero. Later, if the benefits materialize and recurring costs stabilize, the integrand turns positive and cumulative net value eventually crosses the axis. This crossing is late by construction, because it can occur only after observation of the full cost of permanence.

The term $C(t)$ is the mediating variable between the two horizons. It denotes the capacity effectively transferred and institutionalized at the client, $C(t) \in [0, 1]$. It links the DVHM

to prior work on capacity transferability. Its effect is twofold and of opposite sign on the two terms it governs.

$$\frac{\partial BA}{\partial C} > 0 \text{ and } \frac{\partial K}{\partial C} < 0$$

The more capacity is transferred, the more attributable benefits rise, because the client knows how to operate, correct, and extend the system. And the more governance costs fall, because supervision ceases to depend on an external operator. A deployment with low Ckeeps BAlow and Khigh, delays the crossing, and may render it unreachable. A deployment with high Cbrings the break-even nearer. Transferred capacity is therefore not a variable of comfort. It is the channel through which a deployment with a short supplier horizon can, or cannot, produce a client return within a finite horizon.

The net value trajectory then defines a characteristic time, the client break-even, the first instant at which cumulative net value becomes positive.

$$BET_C = \inf\{\, t \mid NVC(t) > 0 \,\}$$

The central property of the model is an inequality between two durations.

$$BET_C \gg T_F$$

$NVC(t)$does not compute an ROI but a characteristic time. In dimensional analysis, the difference is fundamental. ROI is a dimensionless quantity; the break-even is a duration. The question is therefore no longer "what is the project's ROI," but "from what moment can one demonstrate that cumulative net value becomes durably positive."

The supplier reaches its recognition horizon while the client has not yet observed the full cost of permanence. The two actors evaluate neither the same economic object nor the same temporal horizon. High adoption at production go-live principally validates the supplier's capture function. It does not yet demonstrate the existence of a durable net return for the client. This formulation no longer depends on the particular case of FDEs. It describes the structural desynchronization between the supplier's capture window and the client's demonstration window. It is this desynchronization, more than the delivery model itself, that constitutes the economic core of the thesis.

After the FDE withdraws

The cost of permanence is not only technical. It depends on the organization that retains, after the FDE withdraws, the capacity to decide, to operate, to finance, and to answer. These four functions, better named by their verbs than by the portmanteau *ownership*, do

not follow the same actor. To decide is to hold authority. To operate, to execute and maintain. To finance, to bear the initial and recurring costs as well as exposure to failure. To answer, to justify and correct.

The theory of incomplete contracts, from Grossman and Hart to Hart and Moore, supplies the mechanism. When the future states of a system cannot all be specified in the contract, it is residual control rights that determine who decides in unforeseen situations. It is because these rights remain with the client that the window of its return is pushed back to use. Control determines who can act; the contract organizes who must pay. The verb matters, because the final bearer of the cost also depends on the law, on civil liability, on insurance, and on the solvency of the actors.

A decision made by an automated system creates an accountability function that must remain assigned to an identifiable person or institution. It mobilizes a durable capacity of supervision, justification, and remediation. It is what feeds the governance cost K of the equation, and it is what transferability C comes to lighten. It would be tempting to conclude that a single owner of the result is missing. That would be a mistake. The performance of a sociotechnical system is legitimately collective, and designating a single owner would amount to making an actor bear a risk without sufficient control. The stake is to be able to reconstitute who controlled what, who was to intervene, and under which rule when the system drifted. In other words, the allocation of responsibilities must remain reconstructible.

The responsibility that remains with the client is sustainable only on one condition, and that condition is not internalization. Three notions follow one another here without merging. *Competence* is an individual knowledge. *Capacity* an organized ability to act. *Institutionalized capacity* a capacity endowed with durable roles, authority, and budget. The right criterion is not that capacity be internal, but that it be durably available, internal or contractually guaranteed by an identifiable managed operator. One must still not confuse the degrees of this transfer, the very ones that $C(t)$ aggregates. Means can be delivered without competence being transferred. Competence can be transferred without being institutionalized. Capacity can be institutionalized without being actually exercised. A documented competence is not yet an institutional capacity.

Counter-models, limits, and conclusion

The most serious counter-thesis is the outcome-based contract. One objects that a well-designed *gainshare* would align the supplier with the client's performance. The objection is real, but it runs into the structure of the system. Even well designed, the gainshare bears on a performance co-produced by distributed control. The supplier principally controls certain technical components, the conditions of access to the platform, and part of the remediations. The client principally controls the business data, the operational decisions, adoption, and thresholds. The problem does not rest first on the actors' bad faith, but on

the non-contractibility of a performance of which none alone controls or observes all the variables.

There remain the limits proper to the model, which it is better to state than to conceal.

The DVHM is **bilateral**. It opposes supplier value and client value, and that opposition is also its blind spot. Some deployments create a regulatory, ecological, or systemic value captured by neither actor and appearing in neither function. A deployment may destroy client return while producing a collective benefit, or the reverse. The model does not see it. To extend it would require a third function of social value, whose terms reduce to neither CVF nor NVC. This extension is not treated here.

The cost of permanence is treated as a **cost**, never as an **asset**. This is an assumed simplification. The governance, documentation, upskilling, and auditability produced during operation also create future options: reuse on other systems, compliance already acquired, increased bargaining capacity. There is therefore a *value of permanence* symmetric to the cost of permanence, which the model does not capture. To formalize it would amount to adding to the integrand a positive term also dependent on C. This is an identified extension, not an acquired result.

The function $C(t)$ itself remains underspecified. The model posits its existence, its two partial derivatives, and its mediating role. It proposes neither a functional form nor a measurement protocol. Yet it is probably the most determining and the most difficult variable to observe. Measuring it presupposes distinguishing delivered competence from exercised capacity, which no adoption indicator does.

Finally, a limit of demonstration must be named without evasion. This text establishes that the recognition and demonstration windows are *distinct* and governed by different mechanisms. It does not yet establish that this difference *causes* the observed gap in return between supplier and client. It shows a conceptual correlation, not a measured causality. The distinction is exactly the one the model otherwise imposes on its own objects: a recognized link is not an attributed link.

This limit sketches the research program rather than invalidating it. The testable hypothesis is the following: in a substantial proportion of missions, the inclusion of the cost of permanence and causal attribution materially modify the initial evaluation of the client return, and this return cannot be reliably established within the same window as the supplier's. The protocol holds in a few lines. On a population of some fifty FDE missions, compare the initial *business case* to the observed break-even, then to the net return observed at twenty-four months. Document for each the project cost, the cost of permanence by its five categories, the level of adoption, the capacity effectively transferred C, and a counterfactual scenario. The discriminating test bears on C: if high-C missions reach BET_C significantly earlier, at comparable recovery mechanism, the mediation ceases to be narrative and becomes measured. A philosophical hypothesis

then becomes a falsifiable one, and the DVHM a refutable framework rather than a reading grid.

The real counter-models do not confirm the thesis: they circumscribe its domain of validity. A managed operator remunerated on persistent performance, a contract allocating partial responsibilities reconstructibly, a genuinely institutionalized capacity transfer, show the conditions that avoid the problem. The thesis describes a frequent structure of the platform FDE, not a universal law. It does not claim that the client's net return generally becomes negative. A correctly analyzed deployment may show a positive client return after full cost, and that case does not refute it.

The debate on deployed engineers has focused on *delivery*, on the most efficient way to put a system into production. That was the wrong question. The FDE does not only modify the cost of deployment. It shifts in time the conditions under which each party can claim to have created value. Upstream, the recovery mechanism sets the window where the supplier recognizes its return. Downstream, the cost of permanence and causal attribution set, in another window, the one where the client can establish its own. The *business model* organizes the window of the supplier's return; the cost of permanence reveals the client's return. As long as these two windows remain conflated, adoption dashboards will keep taking early revenue for demonstrated ROI.

A final word on scope. The FDE is here only the first field of application of the DVHM. The chain that runs from the recovery mechanism to the recognition window, then to realization, persistence, attribution, and finally to demonstrable return, does not depend on the presence of deployed engineers. It holds for long-term staff-augmentation firms, large ERP programs, cloud platforms, software medical devices, digital twins, and organizational transformations. Wherever a supplier recognizes its return in a window the client does not share, the same asymmetry operates. The FDE merely has the merit of making it visible.

FDE models do not merely shift the distribution of value. They shift the temporal distribution of value creation.