



Firm Valuation When AI Shapes the Business Model

A Milestone-Based Real-Options Framework for the AI Valuation Uncertainty Problem

Walter Kurz¹, Wojtek Stricker¹, Stefan Marx², Frank Reinhardt², Florian Kollberg²

¹Swissi Institute for AI, kurz@swissi-ai.institute; stricker@swissi-ai.institute

²Hochschule für Wirtschaft und Umwelt Nürtingen-Geislingen

ISSN 3043–1921 | Volume 2026 | <https://journal.swissi-ai.institute> | journal@swissi-ai.institute | CC BY 4.0
Article SAIJ-f3jtignfyge3 | DOI: 10.5281/zenodo.21901247 | <https://journal.swissi-ai.institute/doi/f3jtignfyge3>
Submitted 03/26; Revised 04/26; Published 05/26

Abstract

Standard valuation methods, including discounted cash flow, the income approach standard IDW S 1 of the Institute of Public Auditors in Germany, and market multiples, compress milestone probabilities, continuation options, and risk shifts into opaque aggregate parameters; none provides a structured protocol for decomposing AI integration into auditable option-level assumptions. We propose an industry-agnostic taxonomy separating AI Integrators from AI Providers. AI Integrators are further classified by their Integration Depth Level, ranging from no integration to AI at the core of the product or process. A milestone-gated real-options overlay decomposes milestone state value into five components, and an Analytic Hierarchy Process-based Success Readiness Index derives per-option probabilities from structured pairwise comparisons for scenario analysis. Applied to an AI-native energy software-as-a-service firm, the framework yields a coherent valuation band traceable to identifiable option-level assumptions. Risk concentrates in later-stage continuation options, matching the structural prediction for AI Providers. The protocol applies across the firm lifecycle, including mergers and acquisitions due diligence. The case is a single-firm demonstration of protocol coherence, not empirical validation; multi-case testing against realised post-exit valuations is left to future research.

Keywords: firm valuation, AI integration, real options, milestone-based valuation, intangible assets, AHP, multi-criteria decision analysis

1 Introduction

Artificial intelligence has moved from pilot budgets to strategic capital allocation. In the United States, AI-focused firms have reached private-market valuations of hundreds of billions of dollars within years of founding, Anthropic, 2026; Databricks, 2025 while no European company created in the last fifty years has achieved a market capitalisation above EUR 100 billion. Draghi, 2024 The speed and magnitude of this repricing are not explained by technology alone. Value realisation remains highly heterogeneous across firms, sectors, and jurisdictions, Brynjolfsson et al., 2023; Eisfeldt et al., 2023; Singla, Alex, Sukharevsky, Alexander, Hall, Bryce, et al., 2025; Singla, Alex, Sukharevsky, Alexander, Yee, Lareina, et al., 2025 and the pattern is consistent: most deployments remain at limited integration depth, while a smaller group restructures workflows, decision rights, and data feedback loops. Singla, Alex, Sukharevsky, Alexander, Hall, Bryce, et al., 2025; Singla, Alex, Sukharevsky, Alexander, Yee, Lareina, et al., 2025 The difference between experimenting with AI and restructuring around it is a structural divide, and standard valuation practice has yet to absorb it. The standard discounted cash-flow framework values a firm as

$$V_0 = \sum_{t=1}^T \frac{CF_t}{(1+r)^t} + \frac{TV}{(1+r)^T} \quad (1)$$

where CF_t denotes expected free cash flow, r the cost of capital, and TV a terminal value capturing perpetual growth. Koller et al., 2020 This architecture works when the cash-flow engine is stable and projectable from historical patterns. AI integration strains this architecture beyond its design range: it restructures the engine itself through

staged technical, organisational, and regulatory decisions. Benaroch and Kauffman, 1999; Schwartz and Moon, 2000 Each milestone along the integration path (a workflow redesign, a proprietary model deployment, a regulatory clearance) changes what CF_t will look like downstream. The variable the analyst must forecast is itself being reshaped by decisions the firm has not yet made. An analyst who compresses this into a single growth-rate assumption is implicitly pricing every milestone probability, every continuation option, and every risk shift in one opaque number. The DCF number itself may remain coherent; the reasoning behind it becomes invisible, so two analysts who disagree cannot locate the source of their disagreement.

The accuracy of any DCF valuation depends on the data available to the valuator at the valuation date, and that data is a function of operating history. Figure 1 illustrates the valuator's problem. Before founding, there is no observable data: projection uncertainty σ is at its maximum. The act of establishing a firm marks a transition; from that point, costs, contracts, and revenue begin to accumulate, and σ decreases as the projection gains empirical grounding. The longer the firm operates, the lower σ becomes, but it never reaches zero because DCF remains a forecast at every point. When an operational firm introduces AI into its processes, σ spikes: the business model is being reconfigured, and the historical patterns the valuator relied on lose predictive power. The spike is not gradual; at sufficient integration depth, it is discontinuous. As AI milestones clear, σ decreases again, but the rate of recovery depends on the depth and scope of the integration, and σ may not return to its pre-AI level because the business model has structurally changed.




Figure 1: *Projection uncertainty σ over the firm lifecycle (a) and around AI introduction (b). Illustrative.*

What makes AI integration particularly resistant to ex ante valuation is the confluence of mechanisms it activates. Even within a given integration regime, deployment choices exhibit sensitive dependence on initial conditions: small differences in sequencing produce divergent firm-level trajectories because each decision reshapes the capability base available for subsequent ones. Lorenz, 1963 The resulting system exhibits non-linear feedback loops across functions, costs, and competitive position, characteristic of complexity-economic regimes where equilibrium assumptions break down and the distribution of outcomes is fat-tailed rather than normally distributed around a central projection. Arthur, 2021 Strategic interaction amplifies this further: platform economics, winner-take-most dynamics, and moat construction turn each firm's AI decisions into moves in an evolving competitive game whose payoff structure changes as other players act. Rochet and Tirole, 2003; Von Neumann and Morgenstern, 1944 For a valuator, these mechanisms operate simultaneously. The future configuration of an AI-integrated firm carries two stacked layers of uncertainty: magnitude uncertainty driven by the firm's own decisions, and kind uncertainty driven by what competitors do.

Within DCF, three mechanisms compound the problem. First, staged optionality collapses into the long-term growth rate: a contingent, milestone-gated payoff that management may abandon or expand is priced as if it were a deterministic cash-flow trajectory. Second, risk concentrates in a single discount rate that cannot represent the structural shift from high uncertainty at early integration stages to lower uncertainty after milestone clearance. Third, management flexibility (the right to defer, expand, or abandon at each decision node) carries positive value that DCF assigns implicitly at zero.

U.S. private-market practice already reveals this tension. Venture rounds price preferred claims with contractual protections, producing headline post-money valuations that average roughly 48% above common-share fair value. Gornall and Strebulaev, 2020 Tax and disclosure workflows value common shares separately under dedicated standards such as 409A and pre-IPO scenario methods. Maplebear Inc., 2023; U.S. Treasury Department and Internal Revenue

Service, 2026 so the same firm carries two structurally different valuations at the same date. The very existence of this dual-track architecture signals that practitioners already distinguish between option-loaded headline figures and fair-value estimates. A European practitioner reading a U.S. post-money figure as a common-equity equivalent makes a structural category error: the figure prices contractual optionality alongside operating forecasts.

In European and Swiss settings, the mismatch runs deeper. Dominant valuation standards such as IDW S 1 mandate discounted earnings as the primary method and require demonstrable historical patterns. Institut der Wirtschaftsprüfer in Deutschland (IDW), 2016 EU and Swiss AI governance frameworks formalise compliance gates that affect execution cost, admissibility risk, and time-to-scale. European Parliament and Council of the European Union, 2024; Swiss Federal Council, 2025; Swiss Financial Market Supervisory Authority (FINMA), 2024 These instruments were built for firms with established cash-flow histories, and for those firms they work well. The problem arises when the same instruments are applied to milestone-gated AI ventures: not as a deliberate analytical choice, but because no structured alternative exists within the European valuation ecosystem. The result is a measurement mismatch, not conservative valuation. European venture capital stands at less than one-third of U.S. levels as a share of GDP, and comparable startups are valued approximately seven times higher in the United States at the same development stage. Arnold et al., 2024; Draghi, 2024 Part of this gap reflects genuine differences in market structure and risk appetite. Part of it may reflect the absence of a valuation framework that can price what these firms actually are.

The core problem is not whether AI matters for firm value, but how to value integration pathways whose pay-offs depend on depth, sequence, and jurisdiction-conditioned feasibility. Existing literature provides the necessary components: real options for staged technology investment, multi-criteria methods for structured probability assessment, maturity models for depth classification, and market-pricing studies for empirical calibration. Benaroch and Kauffman, 1999; Eisfeldt et al., 2023; Hansen et al., 2024; Saaty, 1980; Schwartz and Moon, 2000 The unresolved task is to integrate these components into a single valuation logic that identifies the relevant milestones, specifies how each probability is estimated, and indicates where integration depth changes the parameter set. Within the present work, a structured, auditable protocol is developed that decomposes AI integration into option-level valuation inputs, comparable across analysts, applicable across integration depths, and consistent with existing valuation standards.

2 Related Work

Valuation research is still dominated by discounted cash-flow logic calibrated on relatively stable business models. Koller et al., 2020 The young-firm literature shows why that architecture weakens under negative earnings, high failure probability, and limited historical depth: conventional DCF and relative valuation break down precisely where the firm profile is dominated by intangible investment, survival uncertainty, and non-linear growth trajectories. Damodaran, 2009, 2018 Intangible-intensive firms compound this problem because their core value drivers resist balance-sheet measurement. Haskel and Westlake, 2017; Lev, 2001 The unresolved issue is not whether DCF remains useful, but where its assumptions become structurally incomplete for staged AI integration.

Market-based and early-stage methods face structurally different but equally limiting constraints. Revenue and ARR multiples, the dominant approach for private AI firm valuation in practice, compress option value into a comparable figure without identifying which milestone generates the premium or how sensitive the valuation is to individual probability assumptions; their application is also constrained by the scarcity of comparables for firms at Integration Depth Level (IDL) 2–3. Damodaran, 2009; Eisfeldt et al., 2023 Early-stage methods such as the VC method and First Chicago introduce scenario weighting but stop short of per-milestone probability estimation and depth-conditioned option architecture. Lerner and Nanda, 2020 The income approach under IDW S 1, the primary European statutory method, shares DCF's structural limitation: it projects a single income path and discounts it, compressing milestone probabilities and continuation-option value into the same opaque aggregate parameter. Institut der Wirtschaftsprüfer in Deutschland (IDW), 2016 The framework developed here does not replace these methods; it adds a decomposition layer in the form of option-level inputs with independent probability estimates linked to identifiable milestone assumptions.

Real-options research addresses staged commitment under uncertainty and supplies the formal language for sequential exercise, deferral, expansion, and abandonment. Dixit and Pindyck, 1994; Myers, 1977; Trigeorgis, 1996 Compound-option formulations extend this to multi-stage investment chains where each decision gate conditions the next. Cassimon et al., 2004; Geske, 1979 Applied work in internet and IT contexts confirms that technology investments are path-dependent and decision-contingent, with value concentrated in future optionality rather than static baseline projections. Benaroch and Kauffman, 1999, 2000; Schwartz and Moon, 2000, 2001 This research programme has also extended to natural-resource investment, where empirical evidence confirms that firms hold loss-making assets open when waiting option value exceeds closure value. Moel and Tufano, 2002 and to strategic investment contexts where option value interacts with

competitive pre-emption and deterrence. Smit and Trigeorgis, 2004 Across all applied domains, the consistent finding is that option value is substantial relative to static NPV and that standard discounted-earnings methods systematically underestimate flexibility value. These models treat technology generically: none addresses AI integration specifically or distinguishes integration depth as a structural variable that reshapes the option parameters themselves. Stage-level probabilities in applied implementations are typically derived from market-implied volatility or unstructured expert judgement, with no structured protocol for estimating each option's likelihood independently or varying individual probabilities in scenario analysis.

AI adoption and maturity studies document heterogeneous operational effects and depth-dependent capability formation. Dreyling et al., 2024; Hansen et al., 2024 Productivity evidence shows strong heterogeneity by worker skill and organisational complementarity, Brynjolfsson et al., 2023 delayed measured gains from costly intangible co-investments, Brynjolfsson et al., 2018 and market repricing that varies with AI exposure. Eisfeldt et al., 2023 Industry surveys confirm near-ubiquitous experimentation but limited deep workflow redesign. Singla, Alex, Sukharevsky, Alexander, Hall, Bryce, et al., 2025; Singla, Alex, Sukharevsky, Alexander, Yee, Lareina, et al., 2025 This stream identifies the empirical phenomenon (integration depth matters) but does not provide a valuation operator that transforms depth and milestone progression into auditable firm-value components.

Existing maturity scales classify depth within individual organisations; none distinguishes firms that embed AI into existing operations from firms whose core product is AI itself. The two cases generate structurally different risk profiles, option architectures, and competitive dynamics. No classification in this literature provides an industry-agnostic taxonomy that maps these distinctions to valuation mechanics.

Multi-criteria decision analysis provides structured tools for aggregating expert judgement under complexity. The Analytic Hierarchy Process decomposes assessment into pairwise comparisons with a built-in consistency diagnostic, Saaty, 1980 and has been combined with real options for telecom and IT infrastructure investment decisions. Angelou and Economides, 2008 Alternative methods (the Analytic Network Process, the Best-Worst Method, Rezaei, 2015 fuzzy extensions) relax specific AHP assumptions but share the core logic of structured criterion weighting. This stream has not been applied to AI-integration-depth-conditioned probability estimation: existing MCDM–real-options hybrids treat technology type generically without modelling how integration depth restructures the assessment criteria themselves.

Venture-finance and private-market evidence adds a further missing element: pricing mechanics. Staged financing is modelled as milestone-contingent commitment under uncertainty, Lerner and Nanda, 2020; Talmor and Cuny, 2005 and security design drives systematic divergence between headline round valuations and common-share fair values in unicorn markets. Gornall and Strebulaev, 2020 Regulatory and tax frameworks formalise this divergence through distinct valuation tracks for preferred and common equity. Maplebear Inc., 2023; U.S. Treasury Department and Internal Revenue Service, 2026 In AI markets, winner-take-most dynamics and complementary-asset moats concentrate option value in a small number of firms, Azoulay et al., 2024; Korinek and Vipra, 2025; Rochet and Tirole, 2003 compounding the disconnect between headline figures and operating-value fundamentals.

Cross-jurisdiction evidence indicates that valuation outcomes are shaped by financing ecosystems and regulatory institutions. European competitiveness diagnostics identify persistent scale-up constraints rooted in capital-market structure, exit availability, and risk-capital depth. Arnold et al., 2024; Draghi, 2024 EU and Swiss AI governance frameworks add a distinct layer: jurisdiction-specific compliance requirements that condition which milestones are feasible, at what cost, and on what timeline. European Parliament and Council of the European Union, 2024; Swiss Federal Council, 2025; Swiss Financial Market Supervisory Authority (FINMA), 2024

For the valuation framework, these requirements enter directly as $C_{r,s}$ in the V_s decomposition and as constraints on milestone feasibility that vary by jurisdiction. A firm pursuing IDL 3 integration under EU AI Act oversight faces compliance gates that alter both the timing and the probability of milestone clearance relative to the same firm operating in a lighter regulatory environment. Regulatory jurisdiction must therefore enter the framework as a parameter that varies across cases.

No existing framework provides an industry-agnostic taxonomy that distinguishes AI integration from AI production, links both categories to milestone-gated option structures, and estimates each option's probability independently so that scenario analysis can isolate individual sensitivities. The required synthesis spans depth-conditioned classification, real-options architecture, structured per-option probability estimation, security-design effects, and jurisdiction-conditioned feasibility. Each literature stream supplies components; none combines them into a single auditable valuation logic.

3 Contribution

DCF remains the necessary foundation for firm valuation; the issue is which protocol governs the assumptions embedded in it when AI integration is priced. The framework developed here decomposes the opaque growth-rate assumption into identifiable options and stated probabilities, thereby making the source of the valuator’s projection uncertainty (σ in Figure 1) explicit. The contribution is structural rather than predictive. The framework extends DCF with a decomposition protocol for AI-related inputs, without any claim to estimate option values precisely.

The scope is deliberately diagnostic and aims to make assumptions comparable across analysts, firms, and regulatory environments. The paper contributes:

1. An *industry-agnostic taxonomy* separating AI Integrators (IDL 0–3) from AI Providers (AI Wrapper and AI Native). IDL 0 and IDL 1 are valuation-neutral under this framework; IDL 2 and IDL 3 generate distinct option structures.
2. A *milestone-gated real-options overlay* where each milestone along the integration path gates a sequential investment decision with depth-specific and category-specific probabilities and payoffs, producing the additive structure $V_0 = V_{DCF} + V_{AI}$.
3. A V_s *decomposition* splitting each milestone’s opaque value into five auditable components: baseline value, operating uplift, continuation optionality, regulatory cost, and execution cost.
4. A *per-option probability architecture* where each option carries an independently estimated p_s derived from AHP-based intersubjective scoring (the Success Readiness Index), calibrated against depth-specific base rates and adjustable individually in scenario analysis. The SRI protocol accommodates mixed rater panels of internal and external experts; rater pool composition is an explicit, documented design choice governed by two credentialing dimensions (domain knowledge and conflict of interest), making rater governance itself auditable and contestable.
5. A *scenario simulation layer* producing structured valuation bands where each bound traces to identifiable option-level assumptions, enabling the valuator to locate where sensitivity concentrates.
6. *Testable predictions* on when markets react to AI milestones, how analyst dispersion changes with integration depth, and which milestone types move value most.

The framework applies at any lifecycle stage, from pre-investment planning to M&A due diligence on firms with fully completed AI integration, where option value has migrated into the DCF baseline.

4 Research

The valuation problem differs structurally depending on whether a firm uses AI to transform an existing business or sells AI as a product. A bank deploying machine-learning models for credit scoring faces integration risk, organisational adaptation cost, and regulatory compliance gates; a firm that develops and licenses the underlying model faces technology risk, platform economics, and competitive moat dynamics.

A single firm may contain elements of both, but the valuation logic for each component differs. Collapsing both into one classification obscures the mechanisms that drive option value. Figure 2 introduces a taxonomy that separates these categories and defines the sub-classifications used throughout this paper.




Figure 2: *Taxonomy of AI-integrated firms: AI Integrators by depth (IDL 0–3) and AI Providers by type, with progression arrows.*

AI Integrators are classified by Integration Depth Level (IDL 0–3). IDL 0 denotes no AI involvement and serves as

a boundary condition. IDL 1 (Baseline) captures firms using off-the-shelf AI tools such as general-purpose chatbots or embedded copilot features in standard software; this creates no switching costs, no proprietary capability, and no measurable cash-flow restructuring. Because neither IDL 0 nor IDL 1 generates a decision node with identifiable conditional value, the option overlay is zero by construction and the framework's analytical contribution begins at IDL 2. IDL 2 (Partial) marks the threshold where AI shapes specific business processes, producing measurable efficiency gains and the first real options. IDL 3 (Full) describes firms where AI is embedded across core operations, trained on proprietary data, and generates strategic feedback loops. At this depth, continuation-option value dominates operating gains, and the valuation problem shifts from measuring efficiency improvements to pricing future strategic optionality. The classification is cross-sectional; firms at any IDL level may remain there indefinitely.

AI Providers are subdivided into AI Wrapper and AI Native. A wrapper builds on third-party models, adding domain-specific interfaces or vertical integration without developing the underlying technology. A native develops proprietary models or infrastructure as its core product. The risk profiles diverge: wrappers face upstream dependency and compete on application-layer differentiation; natives face deep technology risk but can build platform-level moats. Progression from wrapper to native is rare but real, and the distinction matters for valuation. A permanent wrapper carries a narrower option set. A wrapper with a credible plan to develop proprietary models carries substantially more option value, because the transition unlocks moats, pricing power, and reduced dependency on upstream providers.

The unit of analysis for the valuation model is the *milestone state*: an observable, dateable point in the AI integration path at which the firm faces a decision node (continue, expand, pivot, or abandon). Each milestone state s carries a conditional value V_s and an independently estimated probability of success p_s ; the model prices the firm's AI component as a portfolio of such milestone-gated options. Table 1 summarises the valuation implications by integration depth.

Table 1: *Valuation implications by integration depth level.*

IDL	DCF impact	Option relevance	Risk	Valuator focus
0 (None)	Negligible	None	None	Ignore
1 (Baseline)	Negligible	None	None	Monitor
2 (Partial)	Margin improvement	Expand/abandon	Moderate	Core driver
3 (Full)	CF restructuring	Build/license/spin-off	High	Entire valuation

The framework proposed here builds on DCF and makes the assumptions an analyst already embeds in it explicit and decomposable. In standard practice, the AI-related growth premium is a single opaque parameter. The option overlay replaces that parameter with a structured set of milestone-level inputs: which milestones are priced, at what probability, with what conditional value, and under which IDL classification. A sufficiently rich multi-scenario DCF with probability-weighted stage gates approximates this structure; at that point, the analyst is performing real-options analysis with DCF mechanics. The contribution is a transparent architecture that makes the assumptions behind the number auditable and comparable across analysts.

4.1 Valuation model

The total firm value separates existing operations from AI-related option value:

$$V_0 = V_{\text{DCF}} + V_{\text{AI}} \quad (2)$$

where V_{DCF} is the discounted cash-flow value of the firm's existing operations and V_{AI} is the real-options overlay capturing the value of staged AI integration decisions. The AI component aggregates across all milestone-gated options:

$$V_{\text{AI}} = \sum_{s=1}^n \frac{p_s \cdot V_s}{(1 + r_s)^{T_s}} - I_0 \quad (3)$$

where n is the number of identified options, p_s the probability of reaching state s , V_s the value conditional on reaching it, T_s the expected time to reach s , r_s the risk-adjusted discount rate for option s , and I_0 the initial investment already committed.

Equations (2) and (3) decompose the opaque growth-rate assumption into explicit components. The parameterisation is more demanding, but it forces the valuator to state which milestones are priced, at what probability, and with what conditional value.

The value conditional on reaching milestone state s decomposes into five components:

$$V_s = V_{\text{b},s} + \Delta V_{\text{o},s} + V_{\text{c},s} - C_{\text{r},s} - C_{\text{e},s} \quad (4)$$

where $V_{b,s}$ is the baseline business value given sunk costs, $\Delta V_{o,s}$ the incremental operating value from measurable cash-flow changes, $V_{c,s}$ the continuation-option value from future decisions that only exist because state s was reached, $C_{r,s}$ the regulatory and compliance cost, and $C_{e,s}$ the execution cost including integration debt, retraining, and organisational friction.

The five-component decomposition makes the structural shift across integration depths visible. At IDL 2, measurable operating gains ($\Delta V_{o,s}$) dominate: a workflow redesign produces quantifiable cost savings or revenue gains. At IDL 3, continuation-option value ($V_{c,s}$) dominates: the firm's value depends on future decisions about platform scaling, new verticals, and data network effects that have not yet been made. Table 2 summarises this shift. The implication for valuation practice is direct: analyst dispersion increases with integration depth because $V_{c,s}$ is inherently harder to estimate than $\Delta V_{o,s}$.

Table 2: V_s composition by integration depth.

	IDL 2	IDL 3	AI Provider
$\Delta V_{o,s}$ weight	Dominant	Moderate	Varies
$V_{c,s}$ weight	Large	Dominant	Dominant
Estimation difficulty	Medium–high	High	High
Analyst dispersion	High	Very high	Very high

Estimation of V_s components differs by integration depth. At IDL 2, pilot data and operational metrics provide direct estimates of $\Delta V_{o,s}$: measured cost reductions in one department extrapolate to comparable departments through standard capital-budgeting methods. Typical milestones at this depth (proof of concept, pilot deployment, workflow integration, expansion decision) are shorter, less costly, and carry lower payoff variance. Continuation-option value remains small, valued as $p_{\text{expand}} \times \Delta V_{\text{next}}$. At IDL 3, direct measurement gives way to scenario modelling. Milestones (data infrastructure build-out, model training, production deployment, competitive differentiation) take longer, cost more, and carry higher payoff variance. $V_{c,s}$ involves strategic optionality (license the model, spin off the AI unit, enter adjacent markets), each priced as a separate mini-option anchored against comparable transactions. For AI providers, the biotech pipeline analogy applies: phase-gated investment with binary milestone outcomes at each gate, estimated from revenue traction, addressable market models, and funding-round repricing events.

A firm operating at different IDL levels across business functions (a bank at IDL 2 in credit scoring, IDL 0 in back-office correspondence) is treated as a portfolio: $V_{\text{AI}} = \sum_f V_{\text{AI},f}$, where f indexes business functions and each function is classified and valued independently. Functions at IDL 0 and IDL 1 contribute zero to the option overlay.

4.2 Probability architecture

The probability p_s in (3) is the most vulnerable parameter in any real-options application; in practice, it is often an unstructured expert guess. The protocol proposed here replaces that guess with an auditable, intersubjective estimate anchored in a multi-criteria framework. Its core instrument is the *Success Readiness Index* (SRI): a normalised score derived from Analytic Hierarchy Process (AHP) based expert assessment, reflecting how ready an option is to succeed. The SRI is a priority score on a ratio scale measuring relative readiness across criteria, not calibrated probability. Multiple experts score each option independently via the same AHP hierarchy; results are aggregated using the geometric mean, preserving the ratio scale per Saaty's group-aggregation protocol. The per-option probability p_s is then obtained by calibrating the SRI against IDL-specific base rates derived from comparable technology adoption outcomes. Each option carries its own p_s , enabling independent scenario variation across the portfolio: the SRI provides the ordinal ranking; the calibration step maps it onto a probability scale.

The rater panel may include both internal experts (founders, technical leads, domain owners) and external experts (independent advisers, sector specialists). Insiders contribute system-specific knowledge on data infrastructure, organisational constraints, and technical feasibility that external raters would struggle to replicate in a short engagement. The relevant distinction lies in two independent dimensions of each rater's contribution: domain knowledge $w_i^{\text{comp}} \in [0, 1]$ (expertise on this specific option type) and conflict of interest $b_i \in [0, 1]$ (financial stake, reputational commitment, or prior public positions on the firm). These yield an effective aggregation weight

$$w_i^{\text{eff}} \propto w_i^{\text{comp}} \cdot (1 - \lambda b_i), \quad (5)$$

where $\lambda \in [0, 1]$ governs how aggressively the protocol penalizes conflicts of interest. Both scores and the resulting weights are documented alongside the SRI output, so a reviewer can challenge any weight assignment, substitute

alternatives, and trace the effect on the aggregated score. Rater pool composition becomes an auditable design choice rather than an implicit assumption; Table 3 provides the scoring rubric for both dimensions.

Table 3: *Rater credentialing rubric: w_i^{comp} (domain knowledge) and b_i (conflict of interest).*

Band	w_i^{comp} ; domain knowledge	b_i ; conflict of interest
0.8–1.0	Deep operational expertise in this specific option type; multiple comparable prior engagements; demonstrable track record	Direct financial stake (equity, options, success fee); public prior endorsement of the valuation outcome; lead investor or co-founder
0.5–0.7	Solid sector knowledge; some comparable experience; relevant professional credentials	Indirect financial interest (adviser fee, ongoing contract with firm); prior public statements favouring a specific outcome
0.2–0.4	General industry knowledge; limited direct experience with this option type; familiarity from adjacent domains	Reputational interest only; no financial stake; known association with firm or management
0.0–0.1	No relevant domain expertise; generalist background only	No financial stake; no prior public positions on this firm or option; no ongoing relationship

The protocol proceeds in three stages. In the first, experts score each option against five criteria through AHP pairwise comparisons: technical feasibility, regulatory compliance, organisational readiness, market impact potential, and data access and governance. The output is a weighted priority vector constituting the SRI. Consistency ratios flag incoherent judgements, and multi-rater aggregation via geometric mean surfaces disagreements invisible in single-rater estimates.

In the second stage, the SRI is calibrated against IDL-specific base rates to produce p_s . A high SRI does not automatically imply a high p_s : a deep-integration option at IDL 3 may score well on readiness criteria but face a structurally lower base rate because the underlying milestone is more complex and less precedented. The calibration anchors are derived from comparable technology adoption rates, M&A stage conversion rates, or venture-capital milestone outcomes. These anchors are illustrative priors, not empirically derived base rates; they set the probability scale. In the third stage, scenario analysis shifts p_s values (selectively, across the board, or in targeted combinations) to produce a valuation band. The second stage sets the scale; the third stage stress-tests it.

The AHP criterion weights are not fixed across integration depths. At IDL 2, technical feasibility and organisational readiness carry the largest weight: the central question is whether the firm can build and absorb the AI system. At IDL 3 and for AI providers, market impact potential and regulatory risk dominate: the technology is assumed capable, and value depends on market acceptance and regulatory permission. This depth-dependent weight shift reflects the structural reality that the binding constraint on option exercise changes with integration depth.

The consistency ratio deserves explicit framing. The CR measures internal judgement coherence across pairwise comparisons; it does not measure empirical truth. A low CR (below the conventional 0.10 threshold) indicates that the rater’s comparative judgements are logically consistent; a high CR flags structural incoherence requiring re-elicitation. The hypothesis that higher-IDL options produce greater inter-rater dispersion in SRI scores, together with higher average CRs, is testable and follows from the structural prediction that $V_{c,s}$ -dominated options are harder to assess. This paper frames the hypothesis as exploratory; confirming or rejecting it requires the empirical data collection described in section 4.4.

Because each option carries its own p_s , the framework supports targeted scenario analysis. An analyst can shift the probability of a single option, a thematic group, or the entire portfolio and trace the effect through to V_{AI} . Selective shifts isolate sensitivity to specific assumptions (a regulatory tightening, a technology breakthrough); blanket shifts produce ceiling and floor estimates; mixed scenarios combine both. The output is a valuation band whose bounds trace to identifiable option-level assumptions. If shifting one option’s p_s from 0.3 to 0.6 moves V_{AI} by 40%, that option is where risk concentrates.

When a firm pursues multiple AI options, their effects may overlap. If Option A (AI customer scoring) and Option B (AI pricing) both claim credit for the same revenue uplift, summing their V_s values double-counts the gain. The framework must address this without collapsing into a combinatorial problem that defeats practical application.

Three approaches exist, each with a different precision–practicality trade-off (table 4). Attribution by assignment requires the valuator to allocate each value effect to one option, or to split it explicitly with documented percentages. The interaction matrix is an $n \times n$ diagnostic that flags option pairs as synergistic, overlapping, or independent, forcing the valuator to consider interactions before summing. Shapley value decomposition computes each option’s average marginal contribution across all possible coalitions; it is formally correct but computationally prohibitive for practical portfolios ($n = 10$ generates 1 024 coalition evaluations).

Table 4: *Approaches to option interdependency by mechanism, practicality, and precision.*

Method	Mechanism	Practicality	Precision
Attribution by assignment	Human allocation of value effects	High	Medium
Interaction matrix	$n \times n$ synergy/overlap flags	High	Low
Shapley values	Marginal contribution across coalitions	Low	High

The framework proposes attribution by assignment as the core rule, supplemented by the interaction matrix as a diagnostic check. Shapley values remain a direction for future work.

4.3 Lifecycle and empirical design

The framework applies at any point in the AI integration lifecycle, including post-milestone reassessment. As milestones are cleared, option value migrates into the DCF baseline:

$$V_{\text{firm},t} = V_{\text{DCF},t} + V_{\text{AI},t} \quad (6)$$

where $V_{\text{DCF},t}$ absorbs the cash-flow contributions of exercised options and $V_{\text{AI},t}$ contracts as uncertainty resolves. At any valuation date, each option falls into one of four status categories, summarised in table 5.

Table 5: *Option status classification at a given valuation date.*

Status	Value location	Treatment
Not started	Fully in V_{AI}	Full AHP $\rightarrow$ SRI $\rightarrow p_s \rightarrow$ scenario band
In progress	Split: cleared in DCF, remaining in options	Reduced V_{AI} ; p_s updated
Completed	Fully in V_{DCF}	Measurable cash-flow contribution
New continuation	In V_{AI} (new)	Fresh AHP scoring

Probabilities are not static. As milestones clear, p_s for the next milestone updates upward (demonstrated capability reduces uncertainty), r_s decreases (risk resolved at each gate), and new continuation options may emerge that were not identifiable at the earlier valuation date. A firm that has cleared four of five milestones carries a structurally higher p_5 than one that has not started. This Bayesian updating is the mechanism behind the endogenous risk shift: clearing milestones resolves risk and reprices the remaining option portfolio. The framework treats this updating as conceptual; specifying a functional form for $p_{s+1|s}$ cleared would risk overspecification for a framework paper and is deferred to future work.

In an acquisition context, buyer and seller see different option sets. The seller's V_{AI} reflects options exercisable with the seller's current resources, organization, and market position. The buyer's V_{AI} reflects options enabled by the combination: cross-selling into the buyer's customer base, scaling AI to the buyer's operations, or leveraging the buyer's data assets. The synergy premium becomes structurally transparent:

$$\Delta V_{\text{synergy}} = V_{\text{AI}}^{\text{buyer}} - V_{\text{AI}}^{\text{seller}} \quad (7)$$

This decomposition reduces M&A negotiation from a single headline disagreement to a structured comparison of specific options, probabilities, and continuation values that each party prices differently.

4.4 Illustrative application

The framework generates testable predictions: (a) milestone type and IDL depth predict the magnitude of market reactions to AI integration announcements, (b) higher-IDL milestones produce larger abnormal returns, (c) later-stage milestones within a given IDL level produce diminishing incremental effects as option value is progressively priced in, and (d) valuation dispersion increases monotonically with integration depth.

The framework is illustrated through a single evidence-anchored case reconstruction, parameterized from a real firm. The goal is to show that the model produces coherent, decomposable valuation bands under transparent, citable inputs, not to claim causal validation against market prices. All inputs are stated explicitly so readers can verify the arithmetic and substitute their own assumptions.

Case Firm E12 is a Swiss AI-native SaaS startup building an automated energy portfolio optimisation platform.¹ Its core product, E12-eAI, trades simultaneously across Day-Ahead, Intraday, Futures, power purchase agreements, flexibility markets, and CO₂ certificates, replacing the manual spreadsheets and fragmented legacy tools that most energy portfolio operators still rely on. Revenues combine platform licences, performance fees tied to realised customer margin uplift, data products, and integration services. Under the framework taxonomy, *Case Firm E12* is an *AI Provider (Native)*: E12-eAI is the product sold to customers, not a tool deployed internally; the energy portfolio operators who adopt the platform are the AI Integrators.

At the valuation date (Q2 2026), the AI core is at TRL 2–3: the optimisation algorithm has been validated on real market data, but the agent-orchestration layer is still in pre-development. The firm is targeting a CHF 15 million Seed round at a pre-money valuation of CHF 100 million. Four milestone gates over 24 months structure both the technical roadmap and investor decision rights (fig. 3): Gate 1 (month 3) confirms system stability for customer-facing pilots; Gate 2 (month 6) is the financing trigger, releasing the remaining CHF 13.5 million tranche only if two pilots have converted to paid agreements; Gate 3 (month 12) validates product-market fit through first recurring revenues (CHF 400k ARR); Gate 4 (month 24) establishes a reproducible go-to-market model with a proprietary TRL 7 model in production and 10+ active customers. Gate failure closes capital access and terminates all downstream options.




Figure 3: Milestone gate structure and TRL progression, *Case Firm E12* (Q2 2026–Q2 2028). Source: (*Case Firm E12*, 2026).

Each gate maps directly onto the option structure in eq. (3): milestone state s , conditional value V_s , and sequential dependency such that failure at any gate terminates downstream options. The V_s decomposition (eq. (4)) applies at each gate. In table 6, $w_s = p_s V_s / (1 + r_s)^{T_s}$ denotes the present-value-weighted contribution of each option to V_{AI} , and $|\delta w_s|$ captures the absolute change in w_s from a uniform ± 0.10 shift in p_s , holding other milestones at base. V_s combines all five components of eq. (4); $\Delta V_{o,s}$ and $V_{c,s}$ are shown separately as the dominant drivers. All values are in CHF millions.

Table 6: Milestone-gated valuation reconstruction, *Case Firm E12* (Q2 2026). Source: (*Case Firm E12*, 2026).

Milestone	T_s (yr)	p_s	$\Delta V_{o,s}$ (CHF m)	$V_{c,s}$ (CHF m)	V_s (CHF m)	r_s	w_s (CHF m)	$ \delta w_s $ (CHF m)
M1: Shadow mode stable; 2 API pilots in onboarding (TRL 4)	0.25	0.82	2.0	5.0	13.9	15%	11.0	1.3
M2: 2 active paid pilots; CHF 13.5m tranche unlocked (TRL 5)	0.50	0.67	6.0	15.0	34.9	20%	21.4	3.2
M3: First recurring revenues (CHF 400k ARR); TRL 6+	1.00	0.52	18.0	35.0	79.3	25%	33.0	6.3
M4: Reproducible GTM; 10+ customers; TRL 7 model	2.00	0.38	48.0	95.0	198.5	30%	44.6	11.7
$\sum w_s$							110.0	
$-I_0$ (seed capital committed)							-15.0	
V_{AI}							95.0	
V_{DCF} (IP/team baseline)							5.0	
$V_0 = V_{DCF} + V_{AI}$							100.0	

¹The firm’s identity is anonymised in accordance with data-protection requirements and standard research-ethics guidelines for confidential primary-source case studies. The underlying primary documents (investor masterplan, financial model, milestone register, TRL roadmap) are available from the authors upon reasonable request, subject to a non-disclosure agreement.




Figure 4: *PV-weighted option contributions w_s per milestone with ± 0.10 sensitivity bands ($|\delta w_s|$).*

Case Firm E12 has no revenue history at the valuation date. That means $V_{b,s}$, the baseline value given sunk costs, is carried by team credentials and IP rather than by cashflow extrapolation, which is why the $\Delta V_{o,s}$ values at M1 and M2 are small: no sustained customer margin uplift has been demonstrated yet. Value is almost entirely forward-looking, concentrated in $V_{c,s}$, the competitive-moat option that unlocks once proprietary TRL 7 models are in production and deliver continuously improving performance advantages over competitors still dependent on external-model wrappers. The option-killing gate structure has real teeth: the CHF 13.5 million tranche is contractually tied to Gate 2 pilot conversion, so two missed pilots terminate most available capital and effectively kill M3 and M4. The AHP–SRI probability priors reflect this structure directly: $p_{M1} = 0.82$ captures near-complete technical control over Gate 1 conditions; $p_{M4} = 0.38$ captures the compounded uncertainty of product-market fit, competitive dynamics, and 24 months of model maturation.

The base case yields CHF 100 million, matching the disclosed pre-money exactly. The inputs were derived independently from the masterplan’s financial model and gate definitions, so the convergence reflects genuine internal consistency. For context, the masterplan itself arrives at the same figure through a distinct methodology: a milestone-scoring overlay across 13 strategic components (readiness score 36.9% in Q2 2026; strategic value CHF 102.1m) is added to a negative DCF contribution (WACC 30%, pre-revenue; –CHF 1.3m), yielding CHF 100.8m, rounded to CHF 100m as the pre-money anchor ((Case Firm E12, 2026), Ch. 8). Two independent methods converging on the same value provides stronger evidence of internal consistency than either alone. The full valuation band runs from CHF 77 million (downside: all $p_s - 0.10$) to CHF 123 million (upside: all $p_s + 0.10$), a range of CHF 46 million. As fig. 4 shows, M4 drives more than 60% of that range ($|\delta w_s| = 11.7$): the proprietary model maturation and reproducible GTM option concentrate risk in the later stage, exactly as the structural prediction for AI Providers anticipates. Removing M4 entirely collapses V_{AI} from CHF 95 million to approximately CHF 50 million, meaning that roughly half of firm value rests on 24-month execution. Two analysts who disagree on the CHF 100 million pre-money can now identify whether the disagreement sits in p_{M4} , in $V_{c,M4}$, or in the 30% discount rate, each of which is individually challengeable against observable proxies.

5 Discussion

Standard valuation approaches do not ignore AI integration; they absorb it into an undifferentiated growth-rate assumption. The five-component V_s decomposition (eq. (4)) separates option value at each integration stage. $\Delta V_{o,s}$ captures the structural efficiency gain realised at the IDL 2 transition, $V_{c,s}$ captures the competitive-moat value that becomes accessible at IDL 3, and $C_{r,s}$ and $C_{e,s}$ record compliance and execution costs instead of netting them silently against projected revenues. The contribution does not lie in a different total valuation, but in the decomposition itself. Disagreement over V_{AI} can then be located in a specific component, probability estimate, input value, or discount rate.

This decomposition points to two systematic mispricing channels under conventional approaches. IDL 1–2 firms are likely to be undervalued when option value is ignored: the operational efficiency gains captured by $\Delta V_{o,s}$ are real and staged, but a single-stage DCF with a generic AI premium does not separate them from baseline revenues. IDL 3 firms face the opposite risk: markets may price network effects and platform control, represented here by $V_{c,s}$, as already realised although they remain contingent on milestone completion. The framework does not remove that contingency, but it does isolate the probability assumptions (p_s) that drive the valuation instead of embedding them in a single discount rate.

These mispricing channels are compounded by financial reporting standards. Under HGB (§248 II), development costs

carry an optional capitalisation right (Aktivierungswahlrecht) that conservative preparers typically do not exercise; AI development expenditure hits the income statement in full and the balance sheet carries nothing. Under IFRS (IAS 38.57), development costs must be capitalised once six criteria are met: (i) technical feasibility of completing the asset, (ii) intention to complete and use or sell it, (iii) ability to use or sell it, (iv) demonstration of probable future economic benefits, (v) availability of adequate technical and financial resources, and (vi) ability to measure reliably the attributable expenditure; research costs are always expensed. The capitalisation boundary is judgement-intensive and varies across firms and auditors. Under US GAAP (ASC 350-40 for internal-use software), capitalisation begins only at the application development stage and stops at post-implementation, with preliminary project work always expensed. The result: two firms committing CHF 10 million to economically identical AI development programmes may report dramatically different earnings, asset values, and return metrics depending solely on their reporting regime (fig. 5). A value driver that generates CHF 95 million in real-option value may simultaneously produce a multi-year operating loss under HGB and a partially capitalised intangible asset under IFRS. The economic reality is identical; the reported financial statements are not.




Figure 5: *Illustrative treatment of a CHF 10 million AI development investment under three reporting regimes.*

The framework addresses this directly. The cost term $C_{e,s}$ in the V_s decomposition (eq. (4)) records AI development expenditure as an economic cash outflow, regardless of whether the applicable standard classifies it as an asset or an expense. Two analysts applying the framework to the same firm under different reporting regimes will produce the same option-level inputs, because the decomposition operates on the economic layer beneath the reporting layer. This is precisely why it produces comparable valuations where unadjusted multiple-based methods produce incomparable ones.

The structural shift from $\Delta V_{o,s}$ -dominance at IDL 2 to $V_{c,s}$ -dominance at IDL 3 generates a directional empirical prediction: analyst forecast dispersion should increase monotonically with IDL depth. Competitive-moat value contingent on full integration success, represented here by $V_{c,s}$, is inherently harder to estimate than $\Delta V_{o,s}$, which has observable operational proxies. This prediction is testable from analyst forecast data without access to the internal valuations that produced it. The exploratory AHP hypothesis runs in parallel: if pairwise-comparison consistency ratios increase with IDL depth, that constitutes independent evidence that the estimation problem is structurally harder at higher integration levels, beyond what individual rater heterogeneity alone would generate.

Single-rater probability estimates are vulnerable to anchoring and overconfidence, both well-documented patterns in expert elicitation. The AHP-based SRI protocol constrains expert judgement rather than replacing it. Pairwise comparisons expose internal inconsistency before aggregation, and the credentialed weighting scheme (eq. (5)) discounts estimates from raters with domain gaps or conflict-of-interest exposure. The resulting gain over current practice is narrow and concrete: disagreement can be assigned to a specific p_s estimate for a specific option at a specific IDL stage, where a negotiated discount rate would otherwise absorb it.

$C_{r,s}$ enters the V_s decomposition (eq. (4)) as a cost term, but regulatory compliance may also carry option value that the present framework does not model fully. Firms that complete EU AI Act compliance milestones early may acquire a market-access advantage that late compliers cannot replicate quickly; in regulated industries that advantage may persist. In the Swiss context, FINMA's 2024 guidance on AI-driven financial services (Swiss Financial Market Supervisory Authority (FINMA), 2024) and the Federal Council's sector-specific approach (Swiss Federal Council, 2025) create distinct regulatory milestone paths whose option structure differs from the EU model. A fuller treatment of regulatory compliance as a staged option remains outside the present scope.

The lifecycle model (eq. (6) and table 5) and the synergy decomposition (eq. (7)) reframe M&A negotiation. Under current practice, acquirer and target disagree on a headline premium whose component drivers are opaque to both parties. The framework reduces this to a structured comparison of specific options, probabilities, and continuation values. An acquirer who assigns high probability to completing the IDL 3 transition can justify a higher premium transparently; a target that believes its p_s estimates are conservative can argue that point at the option level rather than relitigating the discount rate. This does not eliminate negotiation difficulty; it replaces narrative disagreement with structured disagreement that due diligence can address.

Three additional finance-side gaps are worth making explicit before turning to the literature.

First, valuation multiples. EV/Revenue and P/E are the standard plausibility checks on any DCF, but they presuppose a homogeneous peer group. An IDL 3 firm with a platform-level data moat trades at fundamentally different multiples than an IDL 1 firm using off-the-shelf AI tools, even within the same sector and the same GICS sub-industry classification. Assigning a technology-sector median multiple to both produces a spurious comparison. The IDL taxonomy makes peer selection explicit and IDL-conditioned: the right comparable set is determined by integration depth, not sector label alone.

Second, terminal value. In a standard DCF, 60–80% of total present value typically concentrates in the terminal-value term, a single growth rate applied to a normalised cash flow in perpetuity. For a firm at an IDL transition boundary, this growth rate implicitly bets on whether IDL 3 continuation value materialises. The lifecycle model (eq. (6)) addresses this structurally: as milestones clear, option value migrates into the DCF baseline and the terminal-value calculation is applied to an already-adjusted cash flow rather than one that embeds unresolved option payoffs. The framework does not eliminate terminal-value sensitivity, but it confines it to the post-exercise DCF rather than allowing it to absorb milestone uncertainty silently.

Third, preferred equity and dilution. The pre-money headline valuation of an AI-native firm diverges from its common-equity value. Liquidation preferences, ESOP pools, and anti-dilution provisions systematically decouple the two: Gornall and Strebulaev find that US unicorn post-money valuations exceed fair common-equity value by 48% on average due to the economic rights embedded in preferred shares (Gornall & Strebulaev, 2020). The Case Firm E12 illustrates this directly: the CHF 100 million pre-money is a Series seed headline. The V_s decomposition operates on economic cash flows attributed to each option, which means dilution and liquidation-preference effects must be modelled separately in the capital-structure layer rather than embedded in either p_s or r_s . Conflating them produces a valuation that is internally consistent at the option level but misleading at the equity level.

Table 7 maps each gap to its framework response; the common pattern is that the decomposition moves the problem to a layer where economic cash flows can be compared directly, rather than leaving it embedded in a multiple, a growth rate, or a headline number.

Table 7: *Three conventional finance-side gaps and the framework’s response.*

Gap	Standard limitation	Framework response
Valuation multiples	IDL heterogeneity makes peer groups non-comparable within a sector; median multiple mixes IDL 1 and IDL 3 firms	IDL classification conditions peer selection; multiples applied IDL-specifically rather than sector-wide
Terminal value	60–80% of DCF value in a single perpetuity growth rate that implicitly bets on IDL 3 continuation	Lifecycle model explicitly migrates option value into DCF baseline as milestones clear; terminal value applied to adjusted post-exercise cash flow only
Preferred equity / dilution	Pre-money headline $\neq$ common-equity value; liquidation preferences, ESOP pools, and anti-dilution decouple the two (average 48% overstatement for US unicorns (Gornall & Strebulaev, 2020))	V_s decomposition operates on economic cash flows; dilution and preference effects modelled separately in the capital-structure layer

The Case Firm E12 functions as a demonstration, not as an empirical validation. The case was chosen for disclosure quality and for a clean AI-native structure, conditions that favour the framework’s applicability and that bias case selection toward firms where the decomposition fits. One case establishes that the protocol can be applied coherently and yields a traceable band; it does not establish that the framework outperforms a conventional DCF on any observable metric, nor that the decomposition generalises to hybrid integrators, low-disclosure firms, or integration failures. A proper test requires a multi-case portfolio covering IDL levels, integration successes and failures, scored against realised post-exit valuations. That test is left to future research.

The framework connects three bodies of work. It introduces option structure into intangible-asset valuation (Lev, 2001),

where the disclosure problem is recognised but no mechanism is provided for staging value realisation. It introduces integration-depth architecture into real-options applications in technology valuation (Schwartz & Moon, 2000), where uncertainty is modelled at the firm level without conditioning option structure on the depth of technology embedding. It also links AI adoption research (Brynjolfsson et al., 2018; Hansen et al., 2024) to valuation by connecting milestone achievement to firm value. The AHP component draws on established multi-criteria methodology (Saaty, 1980) and uses it for IDL-conditioned probability estimation in staged AI integration options.

6 Conclusion

The central claim of this paper is modest in scope but consequential in implication: standard valuation methods are not wrong about AI-integrated firms, but they are structurally incomplete. DCF, IDW S 1's income approach, and market multiples all remain valid instruments for established cash-flow businesses. When applied to firms whose value is dominated by staged, contingent, and depth-dependent AI integration decisions, these methods compress option value, milestone risk, and continuation payoffs into parameters that obscure the very assumptions an analyst needs to challenge. The framework proposed here does not discard that machinery. It adds a decomposition layer, the IDL 0–3 taxonomy, the five-component V_s structure, and the AHP-based SRI protocol, that makes the embedded assumptions explicit and individually auditable.

The framework's contribution lies in the informativeness of the disagreement it surfaces, beyond any improvement in point estimates relative to DCF or multiples. A conventional DCF applied to Case Firm E12 would yield a single present value whose sensitivity concentrates in the terminal growth rate and the discount rate, two parameters that absorb everything the analyst believes about AI integration without revealing what those beliefs are. The framework decomposes that belief into four milestone-level probabilities, five value components per milestone, and a gate structure whose failure conditions are contractually defined. Two analysts who disagree on the CHF 100 million pre-money can now locate their disagreement in p_{M4} , in $V_{c,M4}$, or in the 30% discount rate. That precision is the contribution: a valuation whose internal structure can be interrogated.

Real-options theory has made similar promises before, and the track record warrants scepticism. Schwartz and Moon's option models for internet firms in 2000 were formally elegant but empirically uncalibrated; the biotech pipeline analogy that underpins much of staged-option pricing works well in pharmaceuticals, where phase-gate probabilities have decades of actuarial data, but lacks comparable base rates for AI integration milestones. The AHP-based probability architecture proposed here addresses part of this calibration gap. It replaces unstructured expert guesses with pairwise-comparison discipline, consistency checks, and credentialed aggregation; the resulting p_s estimates still sit in the domain of expert judgement, since observed frequencies for AI integration milestones do not yet exist at scale. The framework's honesty about this point is deliberate: it surfaces where judgement enters the valuation, exposing what a discount rate would otherwise absorb. Whether that exposure leads to better decisions is an empirical question the framework itself cannot answer.

The Case Firm E12 illustrates both the promise and the boundary. On one hand, the framework produces a coherent valuation band (CHF 77m to CHF 123m) from transparent, citable inputs, and the concentration of risk in M4 ($|\delta w_s| = 11.7$, more than 60% of total sensitivity) confirms the structural prediction that AI Provider valuations are dominated by late-stage continuation options. On the other hand, roughly half the firm's value rests on a single 24-month execution bet, a finding that no amount of decomposition makes less uncertain. The framework clarifies what the bet is and where it sits; it does not make the bet safer.

For European valuation practice, the implications extend beyond methodology. The measurement mismatch between how AI-integrated firms create value and how statutory standards capture it is not a temporary gap that market maturation will close. IDW S 1's income approach, built for firms with demonstrable earnings histories, systematically underweights optionality; HGB's treatment of development costs ensures that balance sheets remain silent about the very investments that drive AI value. The framework proposed here operates on the economic layer beneath these reporting conventions: $C_{e,s}$ records AI development expenditure as economic cost regardless of its accounting classification, and the V_s decomposition yields identical inputs under HGB, IFRS, and US GAAP. If European venture capital is to close even part of the sevenfold valuation gap with U.S. peers at comparable development stages, the valuation infrastructure will need to accommodate option-loaded, milestone-gated firm profiles. This paper proposes one architecture for doing so.

7 Limitations and Future Research

Milestone identification depends on disclosure quality, and firms may reveal neither their integration depth nor the timing of intermediate gates with sufficient precision for external reconstruction. Parameter estimation for p_s , V_s , and r_s remains sensitive to assumptions, especially at IDL 2–3 where comparables are scarce; the framework exposes that sensitivity rather than resolving it. IDL classification also requires judgement for hybrid firms operating at multiple levels across functions. The AHP component introduces further constraints: pairwise-comparison burden grows with criterion count, the IDL-specific calibration anchors remain illustrative priors rather than empirically derived base rates, and geometric-mean aggregation assumes a level of rater independence that may weaken when internal and external experts share the same organisational context. The credentialed weighting scheme partially addresses rater-quality and independence concerns, but assigning w_i^{comp} and b_i is itself a judgement that must be documented and remains open to challenge. Option interdependency handling through attribution by assignment likewise depends on human judgement, because the interaction matrix is qualitative and does not replace formal coalition analysis. The case design is illustrative rather than inferential: a single firm with detailed milestone disclosure can test framework coherence, but not predictive accuracy. Case selection was also purposive: the Case Firm E12 offered rich investor disclosure, an AI-native business model, and a clean milestone architecture, conditions under which the framework is expected to fit well. Cases with weaker disclosure, hybrid integration across IDL levels, or post-hoc failed transitions remain outside the present evidence base, and generalisation requires a larger and more varied sample. Regulatory conditions are also moving targets, because the EU AI Act is not yet fully in force and the Swiss sector-specific approach is still evolving.

Future research should move in five directions.

The first direction is empirical validation. Applying the framework to a multi-case portfolio across IDL levels and integration outcomes, with option-level estimates scored against realised post-exit valuations, would test whether the decomposition improves on conventional alternatives rather than merely describing them.

The second direction is longitudinal. Tracking dynamic IDL transitions would show how value migrates from the option overlay into the DCF baseline as milestones clear over time.

The third direction is empirical calibration. Panel data on AI project outcomes across industries could anchor the SRI-to- p_s mapping and refine milestone payoffs by sector.

The fourth direction is methodological. Bayesian updating for $p_{s+1|s \text{ cleared}}$, Shapley-value decomposition for option interdependencies, and robustness checks with alternative MCDM methods such as fuzzy AHP, ANP, or BWM would test how sensitive the framework is to its current simplifying choices.

The fifth direction is regulatory. As EU AI Act provisions take effect and Swiss guidance evolves, cross-jurisdiction evidence could show how compliance regimes alter milestone feasibility, risk, and continuation-option value.

References

- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press. Retrieved February 27, 2026, from https://openlibrary.org/books/OL26982584M/Prediction_Machines
- Angelou, G. N., & Economides, A. A. (2008). A decision analysis framework for adopting it innovation using real options and the AHP. *Technovation*, 28(7), 445–460. <https://doi.org/10.1016/j.technovation.2007.11.003>
- Anthropic. (2026, February). *Anthropic raises 30 billion in series g funding at 380 billion post-money valuation*. Retrieved February 27, 2026, from <https://www.anthropic.com/news/anthropic-raises-30-billion-series-g-funding-380-billion-post-money-valuation>
- Arnold, N. G., Claveres, G., & Frie, J. (2024). *Stepping up venture capital to finance innovation in europe* (IMF Working Paper No. WP/2024/146). International Monetary Fund. <https://doi.org/10.5089/9798400280771.001>
- Arthur, W. B. (2021). Foundations of complexity economics. *Nature Reviews Physics*, 3, 136–145. <https://doi.org/10.1038/s42254-020-00273-3>
- Azoulay, P., Krieger, J. L., & Nagaraj, A. (2024). *Old moats for new models: Openness, control, and competition in generative ai* (tech. rep. No. 32474). National Bureau of Economic Research. <https://doi.org/10.3386/w32474>
- Benaroch, M., & Kauffman, R. J. (1999). A case for using real options pricing analysis to evaluate information technology project investments. *Information Systems Research*, 10(1), 70–86. <https://doi.org/10.1287/isre.10.1.70>

- Benaroch, M., & Kauffman, R. J. (2000). Justifying electronic banking network expansion using real options analysis. *MIS Quarterly*, 24(2), 197–225. <https://doi.org/10.2307/3250936>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative ai at work* (tech. rep. No. 31161). National Bureau of Economic Research. <https://doi.org/10.3386/w31161>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2018). *The productivity j-curve: How intangibles complement general purpose technologies* (tech. rep. No. 25148). National Bureau of Economic Research. <https://doi.org/10.3386/w25148>
- Case Firm E12. (2026). Investor masterplan v2.0 [Unpublished investor document; firm anonymised per data-protection and research-ethics guidelines. Held confidentially by the authors and available on reasonable request, subject to a non-disclosure agreement.].
- Cassimon, D., Engelen, P. J., Thomassen, L., & Van Wouwe, M. (2004). The valuation of a nda using a 6-fold compound option. *Research Policy*, 33(1), 41–51. [https://doi.org/10.1016/S0048-7333\(03\)00089-1](https://doi.org/10.1016/S0048-7333(03)00089-1)
- China Law Translate. (2023, July). *Interim measures for the management of generative artificial intelligence services*. Retrieved February 27, 2026, from <https://www.chinalawtranslate.com/en/generative-ai-interim/>
- Damodaran, A. (2009). *Valuing young, start-up and growth companies: Estimation issues and valuation challenges* (SSRN Working Paper No. 1418687). Stern School of Business, New York University. <https://doi.org/10.2139/ssrn.1418687>
- Damodaran, A. (2018). *The dark side of valuation: Valuing young, distressed, and complex businesses* (3rd ed.). Pearson. Retrieved February 27, 2026, from <https://www.pearson.com/en-us/subject-catalog/p/dark-side-of-valuation-valuing-young-distressed-and-complex-businesses/P200000003482/9780134854106>
- Databricks. (2025, August). *Databricks announces series k financing at more than 100 billion valuation*. Retrieved February 27, 2026, from <https://www.prnewswire.com/news-releases/databricks-announces-series-k-financing-at-more-than-100-billion-valuation-302528687.html>
- Department for Science, Innovation and Technology & Office for Artificial Intelligence. (2023, March). *Ai regulation: A pro-innovation approach*. Retrieved February 27, 2026, from <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach>
- Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press. Retrieved February 27, 2026, from <https://mitpressbookstore.mit.edu/book/9780691034102>
- Draghi, M. (2024, September). *The future of european competitiveness: A competitiveness strategy for europe* (tech. rep.). European Commission. Retrieved February 27, 2026, from https://commission.europa.eu/topics/competitiveness/draghi-report_en
- Dreyling, R., Lemmik, J., Tammet, T., & Pappel, I. (2024). An artificial intelligence maturity model for public sector institutions. *TalTech Journal of European Studies*, 14(2), 217–239. <https://doi.org/10.2478/bjes-2024-0023>
- Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). *Generative ai and firm values* (tech. rep. No. 31222). National Bureau of Economic Research. <https://doi.org/10.3386/w31222>
- European Parliament and Council of the European Union. (2024). Regulation (eu) 2024/1689 laying down harmonised rules on artificial intelligence (artificial intelligence act). Retrieved February 27, 2026, from <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- Geske, R. (1979). The valuation of compound options. *Journal of Financial Economics*, 7(1), 63–81. [https://doi.org/10.1016/0304-405X\(79\)90022-9](https://doi.org/10.1016/0304-405X(79)90022-9)
- Gornall, W., & Strebulaev, I. A. (2017). *Squaring venture capital valuations with reality* (tech. rep. No. 23895). National Bureau of Economic Research. <https://doi.org/10.3386/w23895>
- Gornall, W., & Strebulaev, I. A. (2020). Squaring venture capital valuations with reality. *Journal of Financial Economics*, 135(1), 120–143. <https://doi.org/10.1016/j.jfineco.2018.04.015>
- Hansen, H. F., Lillesund, E., Mikalef, P., et al. (2024). An ai capability maturity model for ai-enabled organizations. *Information Systems Frontiers*, 26, 2147–2163. <https://doi.org/10.1007/s10796-024-10528-4>
- Haskel, J., & Westlake, S. (2017). *Capitalism without capital: The rise of the intangible economy*. Princeton University Press. Retrieved February 27, 2026, from <https://www.ingramacademic.com/9780691175034/capitalism-without-capital/>
- Iansiti, M., & Lakhani, K. R. (2020). *Competing in the age of ai: Strategy and leadership when algorithms and networks run the world*. Harvard Business Review Press. Retrieved February 27, 2026, from <https://store.hbr.org/product/competing-in-the-age-of-ai-strategy-and-leadership-when-algorithms-and-networks-run-the-world/10495>
- Institut der Wirtschaftsprüfer in Deutschland (IDW). (2016). Idw standard: Grundsätze zur durchführung von unternehmensbewertungen (idw s 1) [Stand: 04.07.2016; Veröffentlicht: 15.05.2017].
- Koller, T., Goedhart, M., & Wessels, D. (2020). *Valuation: Measuring and managing the value of companies* (7th ed.). Wiley. Retrieved February 27, 2026, from <https://www.wiley.com/en-us/Valuation%3A+Measuring+and+Managing+the+Value+of+Companies%2C+7th+Edition-p-9781119610885>

- Korinek, A., & Vipra, J. (2025). Concentrating intelligence: Scaling and market structure in artificial intelligence. *Economic Policy*, 40(121), 225–256. <https://doi.org/10.1093/epolic/eiae057>
- Lerner, J., & Nanda, R. (2020). Venture capital's role in financing innovation: What we know and how much we still need to learn. *Journal of Economic Perspectives*, 34(3), 237–261. <https://doi.org/10.1257/jep.34.3.237>
- Lev, B. (2001). *Intangibles: Management, measurement, and reporting*. Brookings Institution Press. Retrieved February 27, 2026, from <https://www.brookings.edu/books/intangibles/>
- Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2), 130–141.
- MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1), 13–39. Retrieved February 27, 2026, from https://econpapers.repec.org/article/aeaajeclit/v_3a35_3ay_3a1997_3ai_3a1_3ap_3a13-39.htm
- Maplebear Inc. (2023, August). Correspondence related to draft registration statement on form s-1. Retrieved February 27, 2026, from <https://www.sec.gov/Archives/edgar/data/1579091/000119312523221820/filename1.htm>
- McDonald, R. L., & Siegel, D. (1982). *The value of waiting to invest* (tech. rep. No. 1019). National Bureau of Economic Research. <https://doi.org/10.3386/w1019>
- Moel, A., & Tufano, P. (2002). When are real options exercised? An empirical study of mine closings. *Review of Financial Studies*, 15(1), 35–64. <https://doi.org/10.1093/rfs/15.1.35>
- Myers, S. C. (1977). Determinants of corporate borrowing. *Journal of Financial Economics*, 5(2), 147–175. [https://doi.org/10.1016/0304-405X\(77\)90015-0](https://doi.org/10.1016/0304-405X(77)90015-0)
- Organisation for Economic Co-operation and Development. (2026). *Artificial intelligence*. Retrieved February 27, 2026, from <https://www.oecd.org/en/topics/artificial-intelligence.html>
- Rezaei, J. (2015). Best-worst multi-criteria decision-making method. *Omega*, 53, 49–57. <https://doi.org/10.1016/j.omega.2014.11.009>
- Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990–1029. <https://doi.org/10.1162/154247603322493212>
- Saaty, T. L. (1980). *The analytic hierarchy process: Planning, priority setting, resource allocation*. McGraw-Hill.
- Schwartz, E. S., & Moon, M. (2000). Rational pricing of internet companies. *Financial Analysts Journal*, 56(3), 62–75. <https://doi.org/10.2469/faj.v56.n3.2361>
- Schwartz, E. S., & Moon, M. (2001). Rational pricing of internet companies revisited. *The Financial Review*, 36(4), 7–25. <https://doi.org/10.1111/j.1540-6288.2001.tb00027.x>
- Singla, Alex, Sukharevsky, Alexander, Hall, Bryce, Yee, Lareina, & Chui, Michael. (2025, November). *The state of ai in 2025: Agents, innovation, and transformation*. Retrieved February 27, 2026, from <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Singla, Alex, Sukharevsky, Alexander, Yee, Lareina, Chui, Michael, & Hall, Bryce. (2025, March). *The state of ai: How organizations are rewiring to capture value*. Retrieved February 27, 2026, from <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value>
- Smit, H. T. J., & Trigeorgis, L. (2004). *Strategic investment: Real options and games*. Princeton University Press.
- Swiss Federal Council. (2025, February). *Federal council opts for sector-specific regulation in the field of artificial intelligence*. Retrieved February 27, 2026, from <https://www.bakom.admin.ch/bakom/en/homepage/digital-transformation-and-ai/ai-regulation.html>
- Swiss Financial Market Supervisory Authority (FINMA). (2024, December). *Finma publishes new guidance on governance and risk management when using artificial intelligence*. Retrieved February 27, 2026, from <https://www.finma.ch/en/news/2024/12/20241218-mm-ki/>
- Talmor, E., & Cuny, C. J. (2005). *The staging of venture capital financing: Milestone vs. rounds* (EFA 2005 Moscow Meetings Paper No. SSRN 487414). Retrieved April 1, 2026, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=487414
- The President of the United States. (2023, November). Executive order 14110: Safe, secure, and trustworthy development and use of artificial intelligence. Retrieved February 27, 2026, from <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>
- The White House. (2025, January). *Initial rescissions of harmful executive orders and actions*. Retrieved February 27, 2026, from <https://www.whitehouse.gov/presidential-actions/2025/01/initial-rescissions-of-harmful-executive-orders-and-actions/>
- Thom, R. (1975). *Structural stability and morphogenesis: An outline of a general theory of models* [Translated by D. H. Fowler]. W. A. Benjamin.
- Trigeorgis, L. (1996). *Real options: Managerial flexibility and strategy in resource allocation*. MIT Press. Retrieved February 27, 2026, from <https://mitpress.mit.edu/9780262201025/real-options/>
- U.S. Treasury Department and Internal Revenue Service. (2026). 26 cfr 1.409a-1: Definition and coverage of deferred compensation. Retrieved February 27, 2026, from <https://www.law.cornell.edu/cfr/text/26/1.409A-1>

- Von Bertalanffy, L. (1968). *General system theory: Foundations, development, applications*. George Braziller.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.
- Yin, R. K. (2017). *Case study research and applications: Design and methods* (6th ed.). SAGE Publications. Retrieved February 27, 2026, from <https://us.sagepub.com/en-us/nam/case-study-research-and-applications/book250150>
- Zeeman, E. C. (1977). *Catastrophe theory: Selected papers, 1972–1977*. Addison-Wesley.

Declarations

CITE THIS ARTICLE (BIBTEX)

```
@article{saijf3jtignfyge3,
  author = {Kurz, W. and Stricker, W. and Marx, S. and Reinhardt, F. and Kollberg, F.},
  title = {Firm Valuation When AI Shapes the Business Model: A Milestone-Based Real-Options Framework for the AI Valuation Uncertainty Problem},
  journal = {Swissi AI Journal},
  year = {2026},
  volume = {2026},
  note = {Article SAIJ-f3jtignfyge3},
  issn = {3043-1921},
  doi = {10.5281/zenodo.21901247},
  url = {https://journal.swissi-ai.institute/doi/f3jtignfyge3}
}
```

PAPER INFORMATION

AFFILIATIONS	¹ Swissi Institute for AI, kurz@swissi-ai.institute ; stricker@swissi-ai.institute ² Hochschule für Wirtschaft und Umwelt Nürtingen-Geislingen
CORRESPONDING AUTHOR	Walter Kurz, kurz@swissi-ai.institute , ORCID 0009-0006-8045-4775
ARTICLE TYPE	Research Article
SWISSI ADDRESS	Sw2:academic01:obj:p1:f3jtignfyge3sorthchi2pzbjowklrynqvelcgfrfgfkozbktiwa:ba9cca65

PAPER DECLARATIONS

AUTHOR CONTRIBUTIONS	All authors contributed equally to this article.
FUNDING	This research received no external funding.
CONFLICTS OF INTEREST	The authors declare no conflicts of interest.
DATA AND CODE AVAILABILITY	The data and code supporting this article are available from the corresponding author on reasonable request.
USE OF AI TOOLS	Generative AI tools were used for editorial work only, such as language editing and formatting, in accordance with international scientific standards. The authors verified the content and remain fully responsible for the article.

JOURNAL INFORMATION

PUBLISHER	Swissi AI Journal
EDITORIAL DATES	Submitted 03/26; Revised 04/26; Published 05/26
LICENCE	CC BY 4.0
OPEN ACCESS	https://journal.swissi-ai.institute/authors#open-access
CREDIT TAXONOMY	https://journal.swissi-ai.institute/authors#credit-taxonomy
JOURNAL CONTACT	https://journal.swissi-ai.institute journal@swissi-ai.institute