Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Financial Innovation and Technological Change: The Timing of Compute Futures

Fuli Yang

dr@garch.ac

GARCH Institute

Yu Liu

Shenzhen University of Information Technology

Shiyu Lin

GARCH Institute

Research Article

Keywords: financial innovation, derivatives design, technological change, real options, compute futures, optimal entry

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11637775/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Financial Innovation and Technological Change: The Timing of Compute Futures

Abstract

Financial contracts must serve exposures that can change while the contracts are being prepared. We study how technological change affects reference choice and launch timing. Preparation raises readiness, while declining commercial relevance can close a fixed design's profitable launch window. A more durable reference can preserve entry after that window closes, including when it costs more and takes longer to prepare. We then derive the hedging values of old-vintage, composite and successor references. Balanced exposures can make the composite the best immediate design, yet waiting can eliminate its entire entry interval. Temporary operation followed by replacement restores entry only when interim value covers the initial launch cost after accounting for the successor opportunity. For a product family serving successive customer cohorts, the relevant exposure is the discounted future mix. With completed preparation and a single entry, sufficiently low launch costs imply immediate entry; higher costs can sustain bypass of a profitable composite. The compute case documents benchmark expansion and financial preparation proceeding on different schedules. The model links these developments to a common design problem: choosing a reference that can earn its entry cost from the exposures it will serve.

Keywords: financial innovation; derivatives design; technological change; real options; compute futures; optimal entry

JEL classification: G13; O33; D81

1 Introduction

A financial contract is built around a risk that users need to transfer. That risk can change during development, testing and regulatory review. Customers move toward new technologies, the mix of exposures shifts, and an initially attractive reference can lose its commercial role before the product begins trading. The sponsor must decide both when to launch and which reference to bring to market.

Compute futures make this problem concrete. A hardware generation can remain in use while hedging demand moves toward its successor. The reference that best serves current users may differ from the reference that will serve the customers of a continuing product family. Preparation improves the ability to launch, but the commercial value of the prepared design depends on how the underlying exposure evolves.

We develop a framework in which technological change affects both the reference selected for a financial contract and the decision to bring it to market. Preparation raises readiness while a fixed reference loses commercial relevance, creating a finite profitable launch window. We derive conditions under which retaining a more durable reference preserves entry after the original window closes, even when that reference costs more or takes longer to prepare. This connects the feasibility of financial innovation to the reference choices available during development.

Our central result identifies when an attractive transitional contract never enters. We derive the hedging values of old-vintage, composite and successor references from the exposures they serve. Within this menu, balanced exposures can make the composite strictly best for immediate entry. Proposition 9 gives a necessary and sufficient condition for that composite to be uniquely optimal

at some eligibility date; otherwise, waiting for the successor eliminates its entire strict-entry interval. Proposition 10 then derives the launch-cost threshold for temporary operation followed by replacement. A replacement option can preserve later flexibility while leaving the initial investment unjustified. Lowering the initial launch cost and lowering the additional redesign cost therefore have distinct effects on whether a transitional market opens.

We also establish how continuing customer service changes the valuation and timing decision. When launch fixes the exposure cohort, current exposures determine the design's capitalized value. When a listing serves successive cohorts, its value depends on the discounted future exposure mix. We derive this valuation and its entry condition, and use it for the exchange application. With completed preparation and one irreversible entry, sufficiently low costs imply immediate entry. At higher costs, a profitable composite can still be bypassed, including when subsequent replacement is available. The decisive margin is current operating income relative to the carrying cost of the launch expenditure. These results link financial product design to customer migration over the product's commercial life.

The compute case provides the institutional setting. Benchmark coverage expanded while the listing schedule was revised, and forward-curve dissemination proceeded before verified trading in the archived record. We organize these developments as linked product-design episodes, preserving the distinction between an available reference, a prepared product and an operating market. Sections 3–4 develop launch windows and exposure-based reference values. Section 5 studies renewal and continuing customer service. Sections 6–7 connect the model to measurement and the compute record.

2 Related literature

Our analysis draws on financial-contract design and dynamic investment. Duffie and Jackson (1989), Cuny (1993), and Tashjian (1995) connect futures specifications to hedging needs, participation and liquidity. Axelson (2013) studies dynamic contract choice under anticipated competition. Fernández-Fuertes and Grégoire (2026) examine compute benchmarks through heterogeneous exposures and hedging effectiveness. We use exposure-based values to study the timing and reference choice of a sponsor facing a changing customer base.

Technology-adoption research supplies the timing framework. Grenadier and Weiss (1997), Farzin, Huisman and Kort (1998), and Doraszelski (2001) analyze adoption across technological change; Cabral and Dezső (2008) study timing and alternative designs. Fernandes (2005) links technology adoption to financial innovation. The present model makes preparation, eligibility and reference relevance explicit, yielding tractable conditions for entry, substitution and abandonment.

Product-development models also treat design as adjustable. Bhattacharya, Krishnan and Mahajan (1998) study product definition in changing environments. Krishnan and Bhattacharya (2002) link technology commitment to uncertainty and design flexibility. Jain and Ramdas (2005) analyze repositioning during development, while Paulson Gjerde, Slotnick and Sobel (2002) and Krankel, Duenyas and Kapuscinski (2006) study successive product development and introduction. These approaches motivate treating the announced design as a choice that can be revised as the opportunity changes.

The renewal analysis builds on investment with switching. Décamps, Mariotti and Villeneuve (2004) compare alternative projects and subsequent conversion; Chronopoulos and Siddiqui (2015) compare sequential adoption with waiting and leapfrogging. Fusai, Kyriakou and

Castiglioni (2021) value component replacement. We apply this logic to reference-specific hedging values and separate the initial launch expenditure from the additional cost of redesign.

Two related unpublished manuscripts address benchmark architecture and successor listing (related manuscripts A and B; details supplied to the editor). Here the focus is selection among old-vintage, composite and successor references as exposures migrate, including the cash flows from serving future customer cohorts. For the empirical setting, Corkish, Holland and Vila (1997) motivate distinguishing early market indicators from subsequent contract success. Our episode structure records preparation, reference changes and entry separately.

3 Preparation and commercial relevance

Time t is measured from the start of preparation. An exchange or product sponsor chooses a launch time t at or after an externally specified eligibility date e ≥ 0. The sponsor can abandon the project for a payoff of zero. Preparation is costless in the baseline, and the launch cost K > 0 is paid only if the product enters. Preparation expenditure incurred regardless of entry is sunk and excluded from this decision.

Let A > 0 be the scale of the design's attainable commercial value. Readiness is q(t) = 1 − exp(−κt), with κ > 0. Commercial relevance for the fixed vintage is exp(−δt), with δ > 0. The gross value available at launch is

$$V(t) = A[1 - \exp(-\kappa t)]\exp(-\delta t). \tag{1}$$

V is the sponsor's launch-date present value of anticipated operating surplus before the fixed launch cost. A combines prospective demand, fee opportunities and operating costs; δ governs the decline in the reference's commercial relevance. Rental prices and trading activity are separate empirical quantities. This valuation integrates future operating flows before the entry decision is evaluated.

A cash-flow interpretation links relevance to the sponsor's operating opportunity. Suppose launching at t locks in service quality q(t), and net operating surplus at elapsed time u is M q(t) exp[−δ(t + u)], where M > 0. With r ≥ 0, integrating discounted surplus over u ≥ 0 gives V(t) = M q(t) exp(−δt)/(r + δ). Thus A = M/(r + δ). Readiness determines the service quality achieved at entry, while relevance determines the remaining commercial exposure served by that design. This interpretation holds post-entry quality fixed and assumes no later renewal; Section 5 studies a separate renewal choice. When comparative statics hold A fixed, they hold capitalized commercial scale fixed; holding the cash-flow primitive M fixed instead adds the corresponding capitalization effect.

With discount rate r ≥ 0, the decision is

$$J(e) = \max\left\{0, \sup_{t \geq e} \exp(-rt)[V(t) - K]\right\}. \tag{2}$$

Preparation proceeds before eligibility. The sponsor chooses whether and when to enter once the eligibility constraint permits it. This timing separates the accumulation of operational readiness from the legal availability of the trading venue.

4 Launch windows and optimal timing

Proposition 1. V(t) is strictly increasing until $t_m = \log(1 + \kappa/\delta)/\kappa$ and strictly decreasing afterward. Its maximum is $V_m = \frac{A\kappa}{\kappa+\delta}\left(\frac{\delta}{\kappa+\delta}\right)^{\delta/\kappa}$. A strictly profitable launch interval exists if and only if $K < V_m$. In that case there are exactly two roots $t_L < t_m < t_U$ of V(t) = K, and positive launch surplus occurs exactly on (t_L, t_U).

Proof. Differentiation gives $V'(t) = A\exp(-\delta t)[(\kappa + \delta)\exp(-\kappa t) - \delta]$. The bracket is strictly decreasing from κ to −δ and vanishes only at t_m. Substitution yields V_m. Since V(0) = 0 and V(t) tends to zero, continuity and strict monotonicity on each side establish the two roots and the interval. If $K = V_m$, entry only breaks even at a single point; if $K > V_m$, entry never covers its cost.

Proposition 2. When $K < V_m$, the unconstrained optimal launch time t^* is unique. For r > 0 it lies strictly between t_L and t_m and solves $V'(t^*) = r[V(t^*) - K]$. For r = 0 it equals t_m. With eligibility e, a strictly profitable optimal launch is t^* if $e \le t^*$, and e if $t^* < e < t_U$. If $e \ge t_U$, the optimal value is zero.

Proof. On (t_L, t_m), define $H(t) = V'(t) - r[V(t) - K]$. The second derivative is $A\exp(-\delta t)[\delta^2 - (\kappa + \delta)^2\exp(-\kappa t)]$, which is negative through t_m. Thus $H' = V'' - rV' < 0$ on this interval. $H(t_L) > 0$ and $H(t_m) < 0$ when r > 0, giving a unique root. After t_m and before t_U, $V' < 0$ and $V - K > 0$, so discounted surplus is strictly decreasing. Outside the profitable interval it cannot dominate abandonment. The stated constraint cases follow.

Corollary. Increasing e cannot raise J(e). It has no effect while $e \le t^*$. On the binding profitable segment, $J'(e) = \exp(-re)\{V'(e) - r[V(e) - K]\} < 0$. At and beyond t_U, launch value is zero. This is the private opportunity cost of a tighter eligibility constraint holding the design and preparation path fixed; it is not a welfare estimate of regulation.

Proposition 3. Faster relevance decay weakly shrinks the profitable launch set, while faster preparation weakly expands it. Both statements follow pointwise: for every t > 0, V falls with δ and rises with κ. Consequently, any previously infeasible date remains infeasible after an increase in δ. The changes need not identify a common causal effect across actual products, whose A and K may also change.

4.1 Joint choice of reference and entry

Consider two reference designs, H and D. H offers greater initial commercial value; D is more durable. Let $A_H > A_D > 0$ and $\delta_H > \delta_D > 0$. Both share $q(t) = 1 - \exp(-\kappa t)$, eligibility e, launch cost K and discount rate r. The sponsor may launch at most one design. The two gross values are $V_j(t) = A_j\, q(t)\exp(-\delta_j\, t)$, $j \in \{H,D\}$, and the joint decision is

$$J_F(e) = \max\left\{0, \sup_{\substack{j\in\{H,D\}\\ t\ge e}} \exp(-rt)\left[V_j(t) - K\right]\right\}. \tag{3}$$

The common-readiness assumption isolates reference durability. D can represent a broader reference or a design whose commercial relevance declines more slowly. Greater durability is a maintained property, not an automatic consequence of greater benchmark breadth. The cost of reduced initial fit is represented by $A_D < A_H$. The sponsor has both designs available without a new switching payment in this comparison; Section 5 introduces costly renewal.

Proposition 4 (Design crossover). For every t > 0, H has greater launch surplus before τ and D has greater launch surplus after τ, where

$$\tau = \frac{\log(A_H/A_D)}{\delta_H - \delta_D}. \tag{4}$$

If e ≥ τ and profitable entry is possible, D is an optimal reference. If e > τ, every strictly profitable optimal launch uses D. Thus a sufficiently late eligibility date selects the durable design from the feasible menu.

Proof. $V_D(t)/V_H(t) = (A_D/A_H)\exp[(\delta_H - \delta_D)t]$. This ratio is strictly increasing and equals one at τ. Common K and the positive discount factor preserve the ranking at each date. For e ≥ τ, D weakly dominates H at every feasible date. For e > τ, dominance is strict at every finite feasible date. A strictly positive maximum exists at a finite date because discounted launch surplus eventually becomes negative or approaches zero from below. Therefore no strictly profitable maximizer can select H.

The crossover is independent of κ, K and r because these terms are common in the pointwise comparison. They remain decisive for whether either design is profitable and when entry occurs. A crossover can occur after both entry windows have closed; ranking alone does not imply entry. Earlier eligibility also need not select H: the sponsor can voluntarily wait for D. The result identifies a sufficient late-eligibility condition for selection, rather than a claim that every delay causes a switch.

4.2 Reference flexibility and the survival of entry

Proposition 5 (Surviving entry). Let H have a profitable interval ending at $t_{U,H}$. If there exists s ≥ e ≥ $t_{U,H}$ such that $V_D(s) > K$, then $J_H(e) = 0$ and $J_F(e) > 0$. More generally, $J_F(e) = \max\{J_H(e), J_D(e)\}$ is weakly decreasing in e, and the profitable dates available under reference choice are exactly the union of the two designs' profitable sets.

Proof. Proposition 1 gives $V_H(t) \le K$ for all t ≥ e. Hence H cannot deliver a positive value. Launching D at s delivers $\exp(-rs)[V_D(s) - K] > 0$. Maximizing over a finite menu commutes with taking the supremum over dates, which gives the value identity. Each individual feasible set shrinks as e increases, so both individual values, and their maximum, are weakly decreasing. At any date at least one design is profitable if and only if that date belongs to the union.

Reference flexibility preserves an opportunity that the original design can lose. The gain is the value of the additional design option, $J_F(e) - J_H(e)$. A tighter eligibility constraint reduces the available opportunities, while retaining an alternative lets the sponsor adapt within the remaining launch window.

4.3 A numerical illustration

Normalize $A_H = 1$ and set $A_D = 0.45$, $\delta_H = 0.50$, $\delta_D = 0.12$, κ = 1, K = 0.10 and r = 0.05. These are illustrative inputs; time and value are normalized units. The crossover is τ = 2.101. With e = 0, H optimally enters at 1.052 with value 0.270, exceeding D's optimized value 0.186. With e = 3, D enters immediately with value 0.171, exceeding H's 0.096. With e = 5, H has no profitable entry, while D retains value 0.113. The same fixed menu therefore produces initial selection of H, later selection of D, and finally entry sustained only by D.

The illustration separates the timing of entry from the reference selected. Its empirical counterparts are design-specific relevance, readiness and commercial value. Rental prices

describe the market environment; observations on the references adopted and the exposures served identify the substitution margin.

4.4 Unequal preparation speeds and launch costs

A durable reference can require more work to prepare. Let $q_j(t) = 1 - \exp(-\kappa_j t)$, with $0 < \kappa_D \leq \kappa_H$, while $A_D < A_H$ and $\delta_D < \delta_H$ continue to hold. The ratio of gross launch values is $R(t) = (A_D/A_H)[q_D(t)/q_H(t)] \exp[(\delta_H - \delta_D)t]$. Its limiting value at zero is $A_D\ \kappa_D/(A_H\ \kappa_H) < 1$, and its limit at infinity is unbounded.

Proposition 6 (Durability with slower preparation). $R(t)$ is strictly increasing for $t > 0$. It crosses one exactly once, at $\tau_q \geq \tau$. The inequality is strict if $\kappa_D < \kappa_H$. More generally, for $K_D \geq K_H > 0$, let τ_K be the unique date at which $R(\tau_K) = K_D/K_H$. For $e > \tau_K$, every strictly profitable optimal launch selects D when the sponsor may choose either design or abandon. Slower preparation and a cost premium therefore delay the sufficient selection threshold without removing it.

Proof. The logarithmic derivative of R is $\delta_H - \delta_D + g(\kappa_D,t) - g(\kappa_H,t)$, where $g(\kappa,t) = \kappa/[\exp(\kappa t) - 1]$. For fixed $t > 0$, g decreases in κ: its derivative has numerator $\exp(x)(1 - x) - 1 < 0$ for $x = \kappa t > 0$. Therefore the logarithmic derivative of R is strictly positive. Continuity and the two limits give each unique crossing. Since $q_D/q_H \leq 1$, the gross-value crossing cannot precede τ. For $t > \tau_K$, $R(t) > K_D/K_H$. If H has positive surplus, then $V_H(t) > K_H$ and $[V_D(t) - K_D] - [V_H(t) - K_H] = [R(t) - 1]V_H(t) - [K_D - K_H] > 0$. If H has nonpositive surplus, abandonment weakly dominates it. Hence no strictly profitable optimum after τ_K uses H.

The cost-adjusted threshold is sufficient, not necessarily the first net-surplus crossing. With a positive cost premium, D need not dominate H at very late dates when both are unprofitable. Including abandonment in the decision is essential: selection concerns profitable opportunities. Existence of a finite τ_K establishes eventual relative advantage; positive entry still requires an open window for D. The survival condition in Proposition 5 applies with the design-specific costs.

A durable reference can overcome slower preparation and a higher launch cost while the remaining opportunity is profitable. The comparison therefore depends on preparation burden, cost and exposure served, as well as the rate at which each reference loses commercial relevance.

4.5 The value of keeping a reference option

Define $S(e) = J_F(e) - J_H(e)$, the value of adding the durable design to a menu that otherwise contains only H. This comparison holds the eligibility date fixed. It measures the value of design flexibility to the sponsor, whereas a change in $J_F(e)$ measures the effect of tightening the eligibility constraint on that same menu.

Proposition 7 (Flexibility and delayed eligibility). $S(e) = \max\{J_D(e) - J_H(e), 0\}$. It is continuous, nonnegative and eventually zero. Wherever $J_D(e) > J_H(e)$ and both value functions are differentiable, $S'(e) = J_D'(e) - J_H'(e)$. Hence flexibility becomes more valuable over any interval on which the original design loses optimized value faster than the durable design, even though $J_F(e)$ remains weakly decreasing.

Proof. Proposition 5 gives the identity. Proposition 2 represents each J_j as a constant until its unconstrained optimum, discounted launch surplus until its upper profitable endpoint, and zero thereafter. These pieces meet continuously; a design without a profitable window has value identically zero. Their difference and its positive part are continuous, and both individual values vanish beyond their finite upper endpoints. Differentiating the difference on a strict D-selection

interval gives the derivative claim. Monotonicity of J_F follows separately from the nesting of feasible launch dates.

The value of flexibility need not increase throughout a delay. It can rise while the original opportunity erodes and then fall as the alternative also approaches its last profitable date. Figure 1 displays this pattern using the parameters in Section 4.3. The original design is selected at early eligibility dates, the durable design later supports entry, and both opportunities ultimately disappear. Preserving a reference option mitigates the private cost of delay without converting delay itself into a benefit.

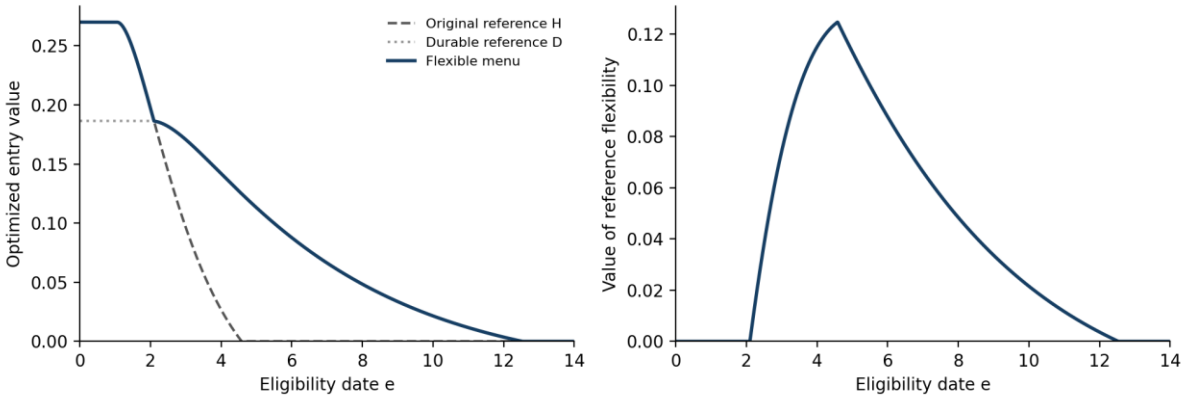


Figure 1. Eligibility, entry value and reference flexibility. Left: optimized values for H, D and the flexible menu. Right: the value of adding D to the original H-only menu. Inputs are the illustrative parameters in Section 4.3; values are computed by constrained optimization, not estimated from GPU prices.

4.6 Exposure migration and benchmark breadth

A broader reference need not be the most durable choice. Consider a separate extension in which value comes from hedging two exposures. At each potential launch date, let X_H and X_N be centered old- and new-vintage cash-flow shocks with unit variances and correlation $0 \le \rho < 1$. A hedger with exposure aX_i chooses h to minimize $\mathrm{Var}(aX_i - hY)$. The optimum removes $\frac{a^2\mathrm{Cov}(X_i,Y)^2}{\mathrm{Var}(Y)}$ from variance. Under mean–variance preferences this reduction has monetary value proportional to risk aversion. Let s(t) be the old-vintage share of aggregate risk-aversion-weighted squared exposure, rather than a share of physical GPU units. This is the minimum-variance hedging framework associated with Ederington (1979).

The sponsor can launch one of three pre-specified references: H settles on X_H, N on X_N, and C on $X_H + X_N$. The normalization of C does not affect an optimally sized hedge. At date t, let $L(t) > 0$ be the sponsor's capitalized share of aggregate hedging value, including the common readiness factor q(t). This assumes a common proportional revenue capture rule and sufficiently elastic outside risk-bearing capacity. It isolates hedge matching and does not solve participation, risk-premium or market-clearing equilibrium. For the chosen contracting horizon, s(t) describes the exposures served at launch. The gross values are $L(t)b_j(s(t))$, with

$$\begin{aligned} b_H(s) &= \rho^2 + (1-\rho^2)s, \\ b_N(s) &= 1 - (1-\rho^2)s, \\ b_C &= \tfrac{1+\rho}{2}. \end{aligned} \tag{5}$$

The two dedicated designs cost K; the composite costs K + ΔK, where ΔK ≥ 0. All three have the same eligibility and common discount factor. Let η(t) = ΔK/L(t) and define

$$w(t) = \frac{\rho}{2(1+\rho)} - \frac{\eta(t)}{1-\rho^2}. \tag{6}$$

Proposition 8 (A transitional composite). At any date with L(t) > 0, C strictly dominates both dedicated designs in net launch value if and only if |s(t) − 1/2| < w(t). This interval is nonempty if and only if ΔK < L(t)ρ(1 − ρ)/2. C is strictly preferred to both designs and abandonment if, in addition, L(t)(1 + ρ)/2 > K + ΔK. Above the interval H weakly dominates C; below it N weakly dominates C, with ties at its boundaries.

Proof. For H users, variance reduction is 1 under H, ρ^2 under N and (1 + ρ)/2 under C; the roles reverse for N users. Weighting by s gives (5). The better dedicated design has gross benefit $\max\{b_H, b_N\}$ = (1 + ρ^2)/2 + (1 − ρ^2)|s − 1/2|. Subtracting it from b_C yields ρ(1 − ρ)/2 − (1 − ρ^2)|s − 1/2|. Multiplication by L and subtraction of ΔK give (6) and the stated inequalities. Positive net value supplies the abandonment comparison.

At balanced exposure, the largest composite premium justified by hedge matching is Lρ(1 − ρ)/2. Holding L fixed, this premium is greatest at ρ = 1/2 and tends to zero as ρ tends to either endpoint. With independent shocks, an equal-weight index cannot strictly beat the best dedicated design in aggregate; with almost identical shocks, a dedicated design already provides nearly the same hedge. A composite is most valuable when the two risks share some, but not all, of their variation.

For illustration, let L = 1, ρ = 0.5, K = 0.10 and ΔK = 0.05. The composite wins for 0.4 < s < 0.6 and has net value 0.60. At s = 0.8, the old-vintage reference has value 0.75 and is preferred. At s = 0.5, each dedicated reference has value 0.525, so the composite is preferred. At s = 0.2, the new-vintage reference has value 0.75 and is preferred. These are assumed-input, same-date comparisons, not estimated market values or a claim that all three products launch sequentially.

If L and the cost premium are constant and s(t) decreases, the same-date ranking can move H → C → N. Once s(t) remains below 1/2 − w(t) at every feasible future date, N dominates C and excludes a strictly optimal composite launch. The next subsection determines whether a composite is ever selected before that point when the sponsor can wait.

The terminal entry window also depends on total exposure. If L(t) tends to L_∞ > K and s(t) tends to zero, N has strictly positive undiscounted net value at all sufficiently late dates. Discounting can make distant entry unattractive without producing a finite last profitable date. In the baseline, every fixed reference has value converging to zero, which produces finite windows. Exposure migration into a viable successor can instead preserve entry indefinitely. The general implication is to evaluate the remaining exposure and full feasible reference menu whenever eligibility changes.

This extension restricts the sponsor to one listed reference from three templates. Allowing continuously adjustable weights or simultaneous H and N contracts changes the design problem. The conditional ranking is useful precisely because it identifies what a claim of greater durability must establish: the exposure path, cross-hedging covariance, implementation cost and reference menu.

4.7 When waiting eliminates a transitional design

To distinguish date-specific design from entry choice, specialize the hedging extension to completed preparation and constant L. Reset time zero to this stage. A sponsor can launch one reference once, at any $t \geq e$, or abandon. Let $r > 0$ and let $s(t)$ be continuous and strictly decreasing from $s(0) = 1$ to zero. Migration can accelerate, slow down or follow an S-shaped path. Take $0 < \rho < 1$, $0 \leq \Delta K < L\rho(1-\rho)/2$, and $c = L(1+\rho)/2 - K - \Delta K > 0$. Define $B = L - K$, $D = L(1-\rho^2)$, and w from (6). Let a and b be the unique dates satisfying $s(a) = 1/2 + w$ and $s(b) = 1/2 - w$. The composite is the best immediate design on (a,b).

The old reference has declining net value $h(t) = L\rho^2 - K + Ds(t)$, so any profitable H launch occurs at e. The composite has constant net value c and also enters immediately if selected. The successor has increasing net value $n(t) = B - Ds(t)$. Express its optimized value in units of the eligibility date:

$$\begin{aligned} R(e) &= \max_{u \geq 0} \exp(-ru)n(e+u), \\ J_N(e) &= \exp(-re)R(e). \end{aligned} \tag{7}$$

The maximum exists and is positive: $n(t)$ tends to $B > c > 0$, while discounting sends the payoff to zero at infinity. The other values are $J_H(e) = \exp(-re)\max\{h(e),0\}$ and $J_C(e) = \exp(-re)c$. No differentiability or constant migration rate is required.

Proposition 9 (Bypassing the composite). Under these assumptions, C is the unique optimal entry for some eligibility date if and only if $c > R(a)$. If $c \leq R(a)$, C is never uniquely optimal, although it is the best immediate design throughout (a,b). If $c > R(a)$, there is one b^* in (a,b] such that C is uniquely optimal exactly on (a,b^*). Throughout this interval it enters immediately; b^* is characterized by $R(b^*) = c$.

Proof. R is continuous because the discounted objective is continuous and its tail vanishes uniformly on compact sets of e. It is strictly increasing: for $e_2 > e_1$, evaluate the objective at an optimizer u_1 for e_1. Since $n(e_2 + u_1) > n(e_1 + u_1)$, $R(e_2) > R(e_1)$. Before a, immediate H beats C; after b, immediate N beats C. At a, H ties C. On (a,b), H is inferior to C, and C beats optimized N exactly when $c > R(e)$. Finally, $R(b) \geq n(b) = c$. Continuity and strict monotonicity establish the condition and the unique endpoint.

The condition compares the full future exposure path with current hedge value. In eligibility-date units, faster discounting weakly reduces R and expands the set of dates selecting C. At a given exposure share, if migration is written $s(t) = s_0(\lambda t)$, the substitution $v = \lambda u$ makes the comparison depend on r/λ. Faster migration relative to discounting raises the attraction of waiting. This timing channel does not depend on exponential decay.

For the exponential case $s(t) = \exp(-\lambda t)$, $a = -\log(1/2 + w)/\lambda$ and $b = -\log(1/2 - w)/\lambda$. Differentiating $\exp(-rt)[B - D\exp(-\lambda t)]$ gives the unconstrained successor date $t_N = \max\{0, \log[D(r+\lambda)/(rB)]/\lambda\}$. With $T(e) = \max\{e, t_N\}$,

$$J_N(e) = \exp[-rT(e)]\{B - D\exp[-\lambda T(e)]\}. \tag{8}$$

The general condition becomes $c\exp(-ra) > J_N(a)$. If $t_N \leq b$, the full interval (a,b) survives. If $t_N > b$ and an interval survives, its upper endpoint is $b^* = \log[c/J_N(a)]/r < b$. The closed form supplies a convenient illustration of the general selection rule.

Use the preceding inputs $L = 1$, $\rho = 0.5$, $K = 0.10$ and $\Delta K = 0.05$, and set $\lambda = 1$. Then $a = 0.5108$ and $b = 0.9163$. At $r = 0.05$, $t_N = 2.8622$ and $J_N(a) = 0.7429$, whereas $J_C(a) = 0.5849$. The

composite has no strict optimal-entry interval. At e = log 2, it is the best immediate design, but its optimized value is 0.5796, below 0.7429 from waiting for N. At r = 0.50, the entire interval (a,b) survives; at e = log 2, C is worth 0.4243 and N is worth 0.3795. These are counterfactual calculations under assumed inputs, not market estimates.

The economic implication concerns a missing product, not merely a delayed one. Same-date hedging superiority does not establish that an exchange should launch a transitional benchmark. Its value must exceed the option to wait for the next exposure. Section 5 relaxes the single-entry restriction by allowing temporary operation followed by replacement, with revenues counted only while each design operates.

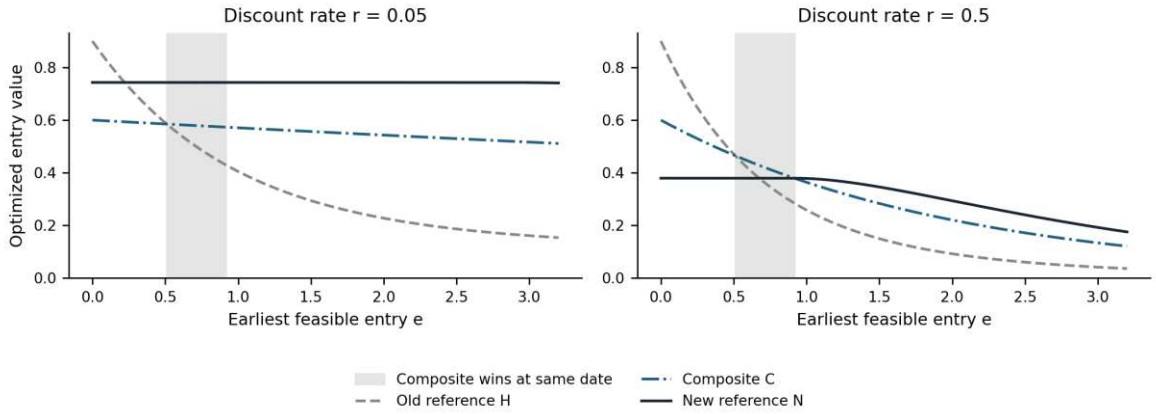


Figure 2. A current design advantage need not survive optimal timing. Shading marks the interval where C beats immediate H and N. Curves allow each reference to choose its launch date. The left panel bypasses C; the right selects C within the shaded interval. Inputs are the illustrative values in Section 4.7.

5 Reference renewal

5.1 Revising a reference before launch

Suppose that at date e the sponsor can retain the old reference or switch once to a new reference. Retaining produces the optimized value O(e), expressed in date-e units. Switching incurs an immediate redesign cost C ≥ 0 and imposes an additional preparation and eligibility delay h ≥ 0. Let B > 0 be the new design's gross commercial value measured at e, $\delta_n \geq 0$ its subsequent relevance decay, and $K_n > 0$ its launch cost. This stylized extension assumes no further switching and fixes h.

The value of switching and then deciding whether to launch is N(e) = −C + exp(−rh) max{B exp(−δ_n h) − K_n, 0}. The sponsor chooses max{0, O(e), N(e)}. When O(e) ≥ 0, switching strictly dominates retaining and abandonment if and only if

$$B > \exp(\delta_n h)\{K_n + \exp(rh)[O(e) + C]\}. \tag{9}$$

Proof. Strict dominance requires positive launch surplus for the new design. Rearranging N(e) > O(e) then gives (9); conversely, (9) implies positive surplus and strict dominance. The condition makes renewal costly in two ways: the direct redesign payment and the additional interval during which the new reference itself ages. Merely labeling a contract with newer hardware is therefore insufficient to establish an improvement.

5.2 Temporary operation followed by replacement

Now allow the sponsor to operate a temporary contract before introducing N. Retain r > 0 and the capitalized successor net value n(t) = B − Ds(t) from Section 4.7; this value already deducts the successor's launch cost K. Let A = L(1 + ρ)/2 be the composite's gross capitalized value and F_C = K + ΔK its nonrefundable launch cost, with A > F_C. Interpret A as a constant operating flow rA capitalized at rate r. A composite launched at τ can be replaced once, at T ≥ τ, with an additional redesign payment S ≥ 0. Replacement ends its operating flow. The successor then delivers its launch-date value n(T), with no further revision. This implements the frozen launch-cohort valuation used in Section 3 while adding one operating transition.

The date-zero payoff from this sequence is

$$\begin{aligned}\int_{\tau}^{T} \exp(-rt) rA\, dt - \exp(-r\tau)F_C + \exp(-rT)[n(T) - S] \\ = \exp(-r\tau)(A - F_C) + \exp(-rT)[n(T) - S - A].\end{aligned} \tag{10}$$

The subtraction of A removes the remaining composite annuity at replacement. Adding the two full launch values would count that revenue twice. Permanent composite operation is the case T = ∞. For any fixed finite T, the first term in (10) is maximized at τ = e because A − F_C > 0. The best sequence that starts with C therefore starts immediately.

Define the replacement option in date-e units by

$$Q_S(e) = \max\left\{0, \sup_{u \geq 0} \exp(-ru)[n(e + u) - S - A]\right\}. \tag{11}$$

Proposition 10 (The cost of temporary entry). The optimized value of entering C and optionally replacing it is exp(−re)[A − F_C + Q_S(e)]. It is strictly preferred to waiting for N if and only if

$$F_C < F_C^*(e, S) \equiv A + Q_S(e) - R(e). \tag{12}$$

The threshold weakly decreases with the additional redesign cost S. When S = 0 and n(e) < A, it lies strictly between zero and A. Consequently, there is a nonempty range of positive-net-value composite launch costs for which waiting is better even though redesign carries no additional charge.

Proof. Equation (10), immediate optimal initial entry and maximization over T give (11) and the sequence value. Subtraction of the direct-entry value exp(−re)R(e) gives (12). Increasing S lowers Q_S. For S = 0, R(e) > Q_0(e): if Q_0 > 0, evaluate direct entry at an optimizer of Q_0; otherwise use R(e) > 0. Thus F_C^* < A. Let u_R attain R(e). If u_R > 0, A + Q_0(e) ≥ R(e) + A[1 − exp(−ru_R)] > R(e). If u_R = 0, R(e) = n(e) < A ≤ A + Q_0(e). Thus F_C^* > 0.

For exponential migration, let z = B − S − A. If z ≤ 0, Q_S = 0 and the composite is never replaced. If z > 0, its optimal replacement date is T_C(e) = max{e, log[D(r + λ)/(rz)]/λ}, and Q_S(e) = exp{−r[T_C(e) − e]}[z − Dexp(−λT_C(e))]. The successor's launch cost remains in B = L − K. Setting S = 0 removes only the additional redesign charge.

An H-first sequence can also be admitted. Its gross annuity at launch is A_H(τ) = Lρ^2 + Ds(τ), with launch cost K. For fixed replacement date T, its payoff has the form (10) with A replaced by A_H(τ). Whenever A_H(τ) > K, earlier initial entry raises the payoff: it increases both the duration of profitable operation and its gross annuity. If A_H(τ) ≤ K, entering H is dominated by waiting for the same successor date. Hence any globally optimal H-first sequence starts at e and is valued by the same option formula with A_H(e) and K.

Use the parameters in Section 4.7, with r = 0.05 and S = 0. At e = log 2, waiting for N yields 0.7429. Entering H and later replacing it yields 0.7210; entering C and later replacing it yields 0.6928. The composite's replacement option is positive and its optimal replacement date is 4.6540, compared with direct N entry at 2.8622. Thus the calculation admits and values an actual replacement opportunity. Its critical initial cost is 0.0981, below the specified F_C = 0.15. Temporary operation raises the composite's value from 0.5796 to 0.6928 but does not change the chosen entry.

In this example C is bypassed at every eligibility date, as Figure 3 shows. Before a, H-first dominates C-first even after allowing replacement: its immediate net value is at least as large, and the replacement option is decreasing and one-Lipschitz in the forgone annuity. For e ≥ a, the exponential threshold F_C^*(e,0) decreases up to t_N, is convex between t_N and the unconstrained composite replacement date, and is zero thereafter. Its maximum over e ≥ a is therefore $\max\{F_C^*(a,0),0\}$, which is below 0.15. Direct N entry dominates throughout this region. The strict comparisons at e = log 2 also persist under small parameter changes.

The economic margin is the cost of opening a temporary market. Reversibility preserves a later choice, but it does not refund the first launch. A viable transition requires sufficient interim operating value after allowing for the successor opportunity. Lower initial costs can restore a transitional launch; the existence of a replacement option alone cannot.

The threshold also separates two kinds of cost reduction. Holding the successor's launch cost fixed, lowering the initial expenditure increases the composite-first value one for one in eligibility-date units. Lowering the additional redesign charge acts only through the discounted replacement option and has no effect while replacement remains unattractive. For a finite optimal replacement delay u_S, a unique optimizer gives $\frac{\partial F_C^*(e,S)}{\partial S} = -\exp(-ru_S)$; when the no-replacement option is strictly optimal, this derivative is zero. Thus cheaper redesign cannot be treated as an equal-sized reduction in the cost of opening the temporary market.

Corollary (A common launch-cost reduction). Hold L, ρ, the exposure path, r, ΔK and S fixed. At differentiability points, let $\beta_R = \exp(-ru_R)$ be the discount factor at the optimal direct-entry delay and let $\beta_Q = \exp(-ru_Q)$ at an optimal profitable replacement; set $\beta_Q = 0$ when replacement is strictly unprofitable. The composite-first advantage over waiting, in eligibility-date units, is $G = A - K - \Delta K + Q_S - R$. The envelope theorem gives $\partial G / \partial K = -1 - \beta_Q + \beta_R \leq 0$. Thus reducing the launch cost common to both products weakly favors composite-first entry relative to waiting. The gain reflects the cost reductions at both entry dates and need not be one for one.

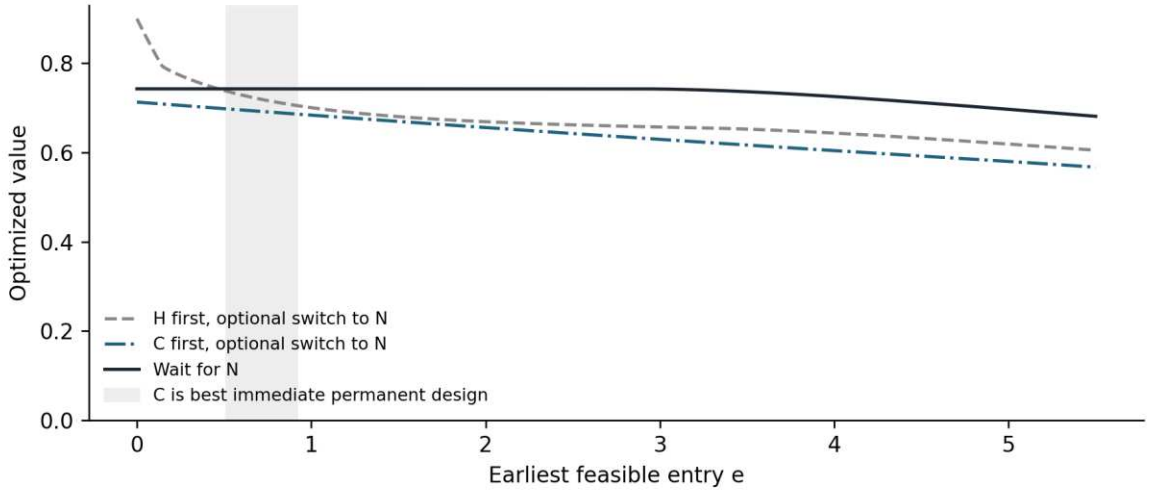


Figure 3. Allowing replacement does not necessarily restore the transitional design. All values include an optional later switch from H or C to N; direct N entry remains available. The shaded interval is where C is the best immediate permanent design. Inputs are those in Section 5.2, with zero additional redesign charge and positive launch costs.

5.3 Serving evolving exposures after entry

The launch-cohort specification fixes the exposures served by an installed contract. Consider instead a listing that continues to serve the changing market. Assume r > 0. Hold L constant and let the operating flow at calendar time t be rL b_j(s(t)), regardless of its launch date. With no subsequent replacement, its launch-date lifetime net value is

$$V_j^{\text{ongoing}}(t) = \int_0^\infty \exp(-ru)\, rL\, b_j\big(s(t+u)\big)\, du - K_j. \tag{13}$$

Because each b_j is affine in s, equation (13) equals L $b_j(\overline{s}(t)) - K_j$, where $\overline{s}(t) = \int_0^\infty r \exp(-ru) s(t+u)\, du$. Thus the relevant exposure is the discounted future mix. Current hedge fit and lifetime launch value are distinct objects.

Proposition 11 (Continuing service and immediate entry). For continuous s, differentiating the calendar-date discounted payoff gives $\frac{d}{dt}\left[\exp(-rt)V_j^{\text{ongoing}}(t)\right] = \exp(-rt)r\left[K_j - Lb_j\big(s(t)\big)\right]$. Hence a design whose operating benefit always exceeds its annuitized launch cost has no incentive to delay. In particular, $K < L\rho^2$ and $K + \Delta K < L(1+\rho)/2$ make all three discounted launch values strictly decreasing. With one irreversible entry, the design with the largest positive lifetime value enters at eligibility. A design that strictly wins this comparison cannot be eliminated by waiting.

For exponential migration, $\overline{s}(t) = [r/(r+\lambda)]\exp(-\lambda t)$. At e = log 2 with the parameters of Section 4.7 and r = 0.05, the immediate lifetime net values are 0.1679 for H, 0.6000 for C and 0.8821 for N, in date-e units. N enters immediately. The same-date values 0.5250, 0.6000 and 0.5250 in Section 4.6 describe the fixed launch cohort; they are not lifetime values for an ongoing service. These calculations preserve the bypassed composite in this example but change its explanation: N already has the larger lifetime value.

For a continuing product family, entry buys access to successive customer cohorts. The sponsor can capture future migration through an existing listing, so the relevant comparison includes future exposures from the outset. The fixed-cohort model instead represents an arrangement

whose returns are tied to the customers served at launch. Customer turnover, contracting horizons and revenue capture determine which specification applies.

Proposition 12 (Optimal successor entry). The low-cost condition is sufficient, not universal. For exponential migration let $\alpha = r/(r + \lambda)$, $B = L - K > 0$ and $D = L(1 - \rho^2)$. The ongoing-service launch values are $h_o(t) = L\rho^2 + \alpha D \exp(-\lambda t) - K$, $c = L(1 + \rho)/2 - K - \Delta K$, and $n_o(t) = B - \alpha D \exp(-\lambda t)$. A positive H payoff is optimally taken at eligibility; a positive C payoff is also taken immediately. The optimal date for N and its value at eligibility are

$$\begin{aligned} T_o(e) &= \max\left\{e, 0, \frac{\log(D/B)}{\lambda}\right\}, \\ R_o(e) &= \exp[-r(T_o(e) - e)]n_o\big(T_o(e)\big). \end{aligned} \tag{14}$$

Proof. To verify equation (14), the derivative of $\exp(-rt)n_o(t)$ is $r \exp(-rt)[D \exp(-\lambda t) - B]$. It changes sign at most once, from positive to negative. For H, h_o decreases with time, so delaying cannot improve any positive launch payoff. The single-entry choice is therefore $\max\{0, h_o(e), c, R_o(e)\}$. A composite that strictly wins the immediate lifetime-value comparison is bypassed whenever $R_o(e) > c$. Waiting requires $K > L[1 - (1 - \rho^2)s(e)]$: the successor's current operating income must be below the carrying cost of its launch expenditure. Future revenue capture reduces the value of waiting but need not eliminate it.

For example, set $L = 1$, $\rho = 0.5$, $r = \lambda = 1$, $K = 0.65$, $\Delta K = 0.03$ and $e = 0$. The immediate lifetime net values are -0.025 for H, 0.070 for C and -0.025 for N. The composite is profitable and strictly best now. Waiting until $T_o = \log(15/7) = 0.7621$ delivers 0.0817 from N, so C is bypassed. All values in this example are measured at date zero. The parameters illustrate the cost threshold; they are not estimates for compute futures.

Allowing C to operate and later switch to N does not overturn this example. With $A = L(1 + \rho)/2$ and incremental redesign charge $S \geq 0$, the replacement gain at T is $n_o(T) - A - S$. Since $n_o(T) < B = 0.35 < A = 0.75$, replacement never pays, even at $S = 0$. C-first remains worth 0.070. Thus bypass can survive ongoing customer service and an optional successor replacement, using the same launch-cost structure. Low-cost immediate entry and higher-cost bypass are different regions of the model.

6 Evidence and measurement

For an exchange, the launch decision concerns a product family capable of serving later trades and contract months. CME's 2025 Form 10-K describes clearing and transaction fees as dependent on traded volume and reports continuing market-data services (CME Group, 2026c, pp. 33–34). Its compute product page specifies financially settled monthly contracts extending 36 months (CME Group, 2026d). These institutional features favor ongoing service as the starting point for the exchange application. They do not identify a constant capture rate L, the launch cost K or the exposure mix. The fixed-cohort specification requires a separate revenue arrangement tying returns to the customers served at launch.

The empirical unit is a product-design episode. A redesign begins a new linked episode rather than overwriting the original plan. Each record retains its announcement date, named reference, proposed launch date, subsequent revisions, regulatory status, testing availability, production listing and first verified trade. Event time and retrieval time are separate fields. Unknown status remains unknown; the absence of a public settlement record is not coded as zero volume.

The existing CME, CFTC and Silicon Data collection provides the primary case record. Section 7 reconstructs the price environment during the preparation interval using matched-date H100 tier spreads, relative B200/H100 rental-index prices, and the sequence of disclosed institutional changes. Such quantities describe changes in published prices. Without performance, configuration and exposure data, they do not identify technological depreciation δ, aggregate hedging demand A or the sponsor's cost K.

A comparison panel should include all products meeting a prespecified inclusion rule, including successful and discontinued launches. Relevant comparisons require official announcement and production dates, an identifiable reference and evidence on design revisions. A single pair of related GPU contracts does not supply independent variation for a launch-hazard regression. The descriptive case and analytical propositions remain distinct until an adequate comparison sample exists.

A test of reference bypass follows the candidate design through four linked observations: an identified exposure and feasible reference; completed preparation and legal eligibility; a subsequent change in the exposure served; and entry of the successor while the earlier candidate remains unlaunched. Exposure shares and covariance locate the candidate's immediate hedging advantage, while launch and redesign costs locate the entry threshold. These observations distinguish the predicted sequence from a project that simply fails to become ready. A broad empirical test compares such sequences across product families and retains both launched and unlaunched candidates in the risk set. The valuation must also establish whether an existing listing captures future cohorts; equations (13)–(14) supply the ongoing-service comparison. The test must establish that a profitable immediate design loses to a later alternative after future customer revenues have already been included in both valuations.

6.1 Observable predictions

The baseline yields three empirical predictions. First, for otherwise comparable product episodes, reference revisions should become more attractive when the original exposure loses commercial importance relative to an available alternative. The relevant outcome is a change in the contract's reference or scope, and the relevant regressor is the relative exposure served by the two designs. Hardware age and rental prices alone do not measure that comparison.

Second, preparation and eligibility have different timing implications. Under the maintained baseline, a product whose eligibility date precedes its unconstrained optimum waits voluntarily, while a still-profitable product made eligible after that optimum enters immediately. A gap beyond both eligibility and the predicted optimum identifies an additional margin to investigate: incomplete readiness, revised commercial value, changing costs or a new design choice. Testing availability is a preparation milestone and cannot substitute for the legal eligibility date or the first production trade.

Third, reference flexibility predicts survival after the original design's opportunity expires. An empirical comparison must retain the original design as a counterfactual rather than counting a redesigned launch as fulfillment of an unchanged plan. Linked design episodes allow researchers to distinguish delayed execution, substitution and abandonment. A successful test requires evidence on the original exposure and on the adopted replacement; expansion of a public index menu establishes the availability of references, not their adoption by a futures contract.

The hedging extension adds a fourth implication: a feasible composite benchmark can remain unlisted while the sponsor waits for a successor reference. Evaluating this mechanism requires comparing the current composite payoff with the successor's optimized future payoff, using

exposure shares, covariance, implementation costs and discounting. A missing launch alone does not identify deliberate waiting; incomplete preparation, competition and uncertain demand can produce the same observation. With replacement available, equation (12) identifies the additional quantities needed: the initial launch cost, interim operating revenue and incremental redesign charge. These costs must be separated from the successor's own launch cost.

6.2 Identification and economic scope

A comparative study would estimate entry and redesign as separate outcomes and preserve discontinued projects in the risk set. Eligibility extensions are not randomly assigned: novelty or complexity can affect review length, preparation and demand together. Identification therefore requires a source of timing variation or a design that models these joint determinants. Treating the review extension as an exogenous experiment would obscure the very product characteristics that drive the mechanism.

The sponsor's private entry value also does not determine the socially preferred review period. Market integrity, participant protection and information generated during review lie outside the payoff in equation (2). A welfare analysis would add those benefits and costs explicitly. The model's institutional implication is operational: assess the remaining exposure and the feasible reference menu when eligibility changes, and record adaptation separately from delay.

7 The compute sequence: reference expansion before entry

7.1 Institutional chronology

The institutional sequence separates three activities that can proceed at different speeds: extending the set of published benchmarks, preparing market-data infrastructure, and obtaining permission to trade. Silicon Data announced a B300 Neocloud Index on September 11, 2026, effective September 14. Its register separately gives April 23 as the index start date. We use September 11 as the public announcement date and September 14 as the effective date; backdated index history does not establish earlier public availability (Silicon Data, 2026).

CFTC's September 21 letter extended review of the H100 and B200 futures submissions through November 9, citing novel or complex issues. CME's September 25 Special Executive Report 9785R replaced the initial listing schedule with a date to be announced. This revision is the operative evidence on the planned listing date in the case record. It is distinct from a realized launch outcome (CFTC, 2026a; CME Group, 2026a).

The October 12 Globex notice then specified October 15 dissemination of non-tradable Silicon Data rental-index forward curves through CME market-data channels. The notice identifies GPU1 and GPU2 and also describes customer testing. The event shows that price dissemination can advance while tradability remains a separate question. On October 26, NYMEX filed a GPU Compute Liquidity Provider Program; the official rules-filing page reports a ten-day review. This adds a dated incentive-design milestone, not evidence of realized liquidity (CME Group, 2026b; CFTC, 2026b).

The sequence places reference expansion, product scheduling and market-data preparation on separate timelines. Hardware benchmark coverage widened before the revised listing schedule, and market-data dissemination followed that revision. The case locates the sponsor's design choice within an evolving set of references.

7.2 Prices during the preparation interval

We extract the three named daily series from tracker version 89 and retain the observation dates and source classifications. The resulting sample contains 74 H100 Neo-cloud observations, 73 H100 Hyperscaler observations and 74 B200 observations from August 20 through November 4, 2026. There are 73 dates observed for all three series. Two separate August 27 product-page FAQ snapshots are excluded from these series. No missing dates are filled. These counts are an audit of the archived tracker, rather than a claim that every historical value remains visible on the current website.

Across the common endpoints, H100 Neo-cloud rises from $2.68 to $2.96 per GPU-hour (10.45%), H100 Hyperscaler falls from $7.27 to $7.12 (2.06%), and B200 rises from $5.67 to $6.53 (15.17%). The Hyperscaler-minus-Neo-cloud spread narrows from $4.59 to $4.16 per GPU-hour, a 9.37% contraction. The B200/Neo-cloud H100 price ratio increases from 2.116 to 2.206, or 4.27%. These are endpoint changes computed from matched dates, not sample-average changes or performance-adjusted price comparisons.

The price evidence rejects a simple description of this interval as a broad decline in GPU rental prices. Commercial relevance in the model must therefore be measured separately from the rental-price level. The economically interesting observation is that the benchmark menu and relative prices changed during preparation. A hardware-specific contract can face a changed set of competing references even while the rental price of its named hardware rises.

7.3 Source discipline and testable implications

The tracker retains four historical governance assertions that its October 17 audit quarantined or marked unverifiable. In particular, announcement 260911-1 currently describes the B300 launch, rather than a September 25 H100 baseline change; announcement 260630-1 describes the H200 launch, rather than a B200 provider expansion. We exclude all four disputed event mappings from event classification and do not attribute the September price movements to them. Their presence in historical notes is not independent corroboration. Daily observations remain usable for descriptive reporting subject to their own provenance, but they cannot establish a structural break from a disputed event label.

The archived CME records contain 94 product–maturity–date entries for October 29 and 30 with published zero activity fields and blank transaction-price fields. Their source notes classify the values as web reference settlements not based on market activity. We preserve those records as an information-product observation and exclude them from trade-based price discovery, trading-success and participant-hedging tests. The first verified production trade remains a separate outcome to be collected.

The next test compares linked product episodes, preserving each original reference and planned date. Redesign creates a new episode. Announcement-to-eligibility, eligibility-to-entry and redesign timing are separate outcomes. Public index launches identify reference expansion; exposure and adoption data identify substitution. The case establishes that benchmarks and forward curves can develop before verified futures trading, while the model determines when the interval changes optimal entry or renewal.

8 Conclusion

Financial entry depends on the exposures a reference can serve over its commercial life. In the model, preparation improves readiness while reference aging reduces the remaining opportunity.

Their interaction creates a finite launch window. Retaining a more durable design can preserve entry after the original reference becomes unviable, even when the alternative is slower or more expensive to prepare.

Exposure migration also determines whether a transitional market opens. A composite can offer the greatest immediate hedging value yet lose to waiting for a successor. Temporary operation can restore its entry when interim earnings cover the initial launch expenditure after allowing for the replacement opportunity. Lowering the launch cost and lowering the additional redesign charge have different effects: the first directly reduces the cost of opening the market, while the second matters through a later replacement decision.

A continuing listing must be valued against the customers it will serve in the future. With completed preparation and a single entry, sufficiently low launch costs make the best lifetime design enter immediately. At higher costs, waiting can dominate a profitable composite even when subsequent replacement is available. These results identify the operating income and launch-cost conditions that separate immediate entry from bypass.

The compute record shows benchmark coverage expanding while financial preparation and the listing schedule changed. It supplies a setting in which product design must be assessed against an evolving exposure. The model’s practical implication is to evaluate the remaining exposure and the available reference menu together at the entry decision. An announced specification is a starting point; the economic choice is which design can support an operating market.

References

Axelson, U. (2013). A theory of the evolution of derivatives markets. Financial Markets Group Discussion Paper 723. Manuscript dated 5 October; series publication November. https://www.fmg.ac.uk/sites/default/files/publications/DP723.pdf

Bhattacharya, S., Krishnan, V., and Mahajan, V. (1998). Managing new product definition in highly dynamic environments. Management Science, 44(11-part-2), S50–S64. https://doi.org/10.1287/mnsc.44.11.S50

Cabral, L., and Dezső, C. (2008). Technology adoption with multiple alternative designs and the option to wait. Journal of Economics & Management Strategy, 17(2), 413–441. https://doi.org/10.1111/j.1530-9134.2008.00183.x

CFTC (2026a). Extension of Initial Review Period of NYMEX Submissions 26-370 and 26-370S. September 21. https://www.cftc.gov/filings/documents/2026/orgdcmnymexcompcontr260921.pdf

CFTC (2026b). Designated Contract Market Rules Filing 65759: Implementation of the NYMEX GPU Compute Liquidity Provider Program. Received October 26. https://www.cftc.gov/IndustryOversight/IndustryFilings/TradingOrganizationRules/65759

CME Group (2026a). Special Executive Report 9785R. September 25. https://www.cmegroup.com/content/dam/cmegroup/notices/ser/2026/09/ser-9785r.pdf

CME Group (2026b). CME Globex Notice, October 12: New Silicon Data Rental Indices Forward Curve Market Data. https://www.cmegroup.com/notices/electronic-trading/2026/10/20261012.html

CME Group (2026c). Annual report on Form 10-K for the year ended December 31, 2025. Filed February 26, 2026, pp. 33–34. https://www.sec.gov/Archives/edgar/data/1156375/000115637526000009/cme-20251231.htm

CME Group (2026d). Compute futures. Product overview and contract specifications. Accessed November 5. https://www.cmegroup.com/markets/energy/power/compute-futures.html

Chronopoulos, M., and Siddiqui, A. S. (2015). When is it better to wait for a new version? Optimal replacement of an emerging technology under uncertainty. Annals of Operations Research, 235(1), 177–201. https://doi.org/10.1007/s10479-015-2010-6

Corkish, J., Holland, A., and Vila, A. F. (1997). The determinants of successful financial innovation: An empirical analysis of futures innovation on LIFFE. Bank of England Working Paper 70. https://www.bankofengland.co.uk/-/media/boe/files/working-paper/1997/the-determinants-of-successful-financial-innovation.pdf

Cuny, C. J. (1993). The role of liquidity in futures market innovations. Review of Financial Studies, 6(1), 57–78. https://doi.org/10.1093/rfs/6.1.57

Data archive (2026). Compute Futures Data Collection Tracker, version 89, November 5, 2026; worksheets 04_CME_DAILY_RAW, 06_BENCHMARK_RAW, 08_EVENTS_RAW and 18_COLLECTION_LOG. Daily series sources: https://portal.silicondata.com/gpu-index-chart?gpu=h100&mainTab=neo-cloud&standalone=true ; https://portal.silicondata.com/gpu-index-chart?gpu=h100&mainTab=hyperscaler&standalone=true ; https://portal.silicondata.com/gpu-index-chart?gpu=b200&standalone=true .

Décamps, J.-P., Mariotti, T., and Villeneuve, S. (2004). Irreversible investment in alternative projects. Working paper, revised July. https://www.tse-fr.eu/sites/default/files/medias/doc/wp/2004/irreversible_investment.pdf. Published in Economic Theory, 28(2), 425–448 (2006), https://doi.org/10.1007/s00199-005-0629-2.

Doraszelski, U. (2001). The net present value method versus the option value of waiting: A note on Farzin, Huisman and Kort (1998). Journal of Economic Dynamics and Control, 25(8), 1109–1115. https://doi.org/10.1016/S0165-1889(00)00023-3

Duffie, D., and Jackson, M. O. (1989). Optimal innovation of futures contracts. Review of Financial Studies, 2(3), 275–296. https://web.stanford.edu/~duffie/jstorlinks/d_jackson1.pdf

Ederington, L. H. (1979). The hedging performance of the new futures markets. Journal of Finance, 34(1), 157–170. https://doi.org/10.1111/j.1540-6261.1979.tb02077.x

Farzin, Y. H., Huisman, K. J. M., and Kort, P. M. (1998). Optimal timing of technology adoption. Journal of Economic Dynamics and Control, 22(5), 779–799. https://doi.org/10.1016/S0165-1889(97)00097-3

Fernandes, A. (2005). Knowledge, technology adoption and financial innovation. Working paper, June. https://anafernandes.name/finnov16bernwpc.pdf

Fernández-Fuertes, R., and Grégoire, V. (2026). Compute futures: Who is left exposed? Working paper, August. Author-hosted HTML edition, accessed 5 November 2026. https://rubenfernandezfuertes.com/papers/2026/compute-futures

Fusai, G., Kyriakou, I., and Castiglioni, M. (2021). Component replacement under uncertainty – a switching option perspective. Working paper, City Research Online. https://openaccess.city.ac.uk/id/eprint/25547/

Grenadier, S. R., and Weiss, A. M. (1997). Investment in technological innovations: An option pricing approach. Journal of Financial Economics, 44(3), 397–416. https://doi.org/10.1016/S0304-405X(97)00009-3

Jain, S., and Ramdas, K. (2005). Up or out—or stay put? Product positioning in an evolving technology environment. Production and Operations Management, 14(3), 362–376. https://doi.org/10.1111/j.1937-5956.2005.tb00030.x

Krankel, R. M., Duenyas, I., and Kapuscinski, R. (2006). Timing successive product introductions with demand diffusion and stochastic technology improvement. Manufacturing & Service Operations Management, 8(2), 119–135. https://doi.org/10.1287/msom.1060.0102

Krishnan, V., and Bhattacharya, S. (2002). Technology selection and commitment in new product development: The role of uncertainty and design flexibility. Management Science, 48(3), 313–327. https://doi.org/10.1287/mnsc.48.3.313.7728

Paulson Gjerde, K. A., Slotnick, S. A., and Sobel, M. J. (2002). New product innovation with multiple features and technology constraints. Management Science, 48(10), 1268–1284. https://doi.org/10.1287/mnsc.48.10.1268.270

Silicon Data (2026). GPU Index Announcements, items 260911-1 and 260630-1. Accessed November 5. https://docs.silicondata.com/products/gpu-index-announcements

Tashjian, E. (1995). Optimal futures contract design. Quarterly Review of Economics and Finance, 35(2), 153–162. https://doi.org/10.1016/1062-9769(95)90012-8

AI assistance

AI assistance. Generative AI tools assisted manuscript preparation and checks of mathematical derivations and code. The authors take responsibility for the manuscript.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ResearchSupplementAnonymous.zip