Why SpaceX Didn't Create a Space Economy

Architecture Lag, Premature Markets, and the Structural Conditions for Market Formation in Frontier Sectors

Sinéad O'Sullivan
s@sinead.co

May 2025

Abstract

SpaceX achieved a 20 times reduction in launch cost per kilogram between 2012 and 2024 while conducting 165 orbital launches in 2025—more than twice China's entire program. Yet no downstream space market segment exhibits functioning price discovery, repeat customers, or a learning curve. This paper explains why, despite incredible technological progress, there is still no viable space economy.

The architecture gap model decomposes market formation into three interdependent system layers, each measured on a scale from 0 (absent) to 1 (fully mature): *technical architecture*, which captures the performance and integration readiness of the technology itself; *market architecture* which captures the structural conditions for scalable economic activity—demand elasticity, cost stability, modularity, and specialization; and *institutional architecture* which captures the regulatory, standards, risk-sharing, and infrastructure frameworks that make coordinated economic activity possible at scale. *System readiness* is the combined maturity of market and institutional architectures. The *architecture gap* measures how far the technology has outpaced the system around it.

Markets form only when this architecture gap remains below a critical viability threshold. In the space economy, the gap is large and widening, meaning that the sector is deeply premature by any reasonable threshold. SpaceX's response—vertical integration via Starlink to manufacture its own demand—is an extraordinary private workaround, not a model for the creation of a space economy, as it does not close the architecture gap for the broader sector.

The paper derives the formal conditions under which cost curves can form, demonstrates that they require architectural stability rather than merely technological progress, and identifies four specific architecture-building interventions that would move the sector toward viability, and a case study of space-based data centers illustrates the framework.

Until governments and institutions build the missing market and institutional architectures of standards, resource rights, pricing mechanisms, and logistics infrastructure, the space economy will remain what it is today: a technologically extraordinary, but commercially premature sector.

JEL Classification: L52, L93, O25, O31, O38

Keywords: space economy, market formation, architecture lag, premature markets, cost curves, vertical integration, industrial policy, institutional architecture

Contents

1. Introduction

The space industry has a story it tells about itself, which is the following: SpaceX made launch cheap, the cost curves are coming down, and that competition will follow, allowing access to space to be fully democratized. Investors have therefore, on this narrative, bet heavily that the downstream space economy—of satellite constellations, in-orbit servicing, space-based manufacturing, and lunar resource extraction—will emerge as launch costs continue to fall.

The first part of this story is true, in that SpaceX *has* achieved a genuine and extraordinary reduction in the cost of reaching orbit. Consider that Falcon 9 launches routinely at a fraction of what legacy providers charge. In 2025, SpaceX conducted 165 orbital launches—more than twice the total of China's entire space program and roughly half of all launches worldwide.

But the rest of the story is wrong.

By 2025, the empirical record is unambiguous. Launch costs fell by a factor of twenty while launch cadence increased by an order of magnitude. Yet no downstream space market segment has emerged with functioning price discovery, elastic demand, repeat customers, or a learning curve. There are zero functioning lunar economy segments, zero in-orbit servicing markets with repeat customers, and zero priced resource-extraction transactions.

This paper argues that the failure of cheap launch to produce a space economy is not a timing problem—a sector that is "early" and will mature with patience—but a structural one. That the binding constraint was never launch cost, but that it was, and remains to be, the absence of the market and institutional architectures without which economic activity cannot form—regardless of how cheap access to orbit becomes.

The paper makes four contributions. First, it formalizes the architecture gap framework as applied to the space economy, providing rigorous definitions for the three architecture layers and the conditions under which they enable or preclude market formation (Section 3). Second, it demonstrates empirically that SpaceX's launch price is not a market price but a *transfer price* set by a vertically integrated company whose largest customer is... itself; and that this fact changes the economic interpretation of the cost reduction (Section 2). Third, it derives the formal conditions under which cost curves can form and shows that they require *architectural co-evolution*, not merely technological progress (Section 5). Fourth, it identifies four specific interventions in the market and institutional layers that would move the sector toward viability, and shows through a case study of space-based data centers why technology-first investment cannot substitute for these (Sections 6 and 7).

2. SpaceX Didn't Find Demand, So It Manufactured It.

In 2024, 89 of SpaceX's 139 Falcon missions carried Starlink payloads. Only 18 were standard commercial launches for external customers. SpaceX's launch price is not a market price, it is a *transfer price* set by a vertically integrated company whose largest customer is itself.

Definition 1 (Transfer Price vs. Market Price)**.** *A* market price p^m *is determined by the interaction of multiple independent buyers and sellers under conditions of price discovery and competitive substitution. A* transfer price p^t *is set by a vertically integrated firm for internal transactions between its own divisions. A transfer price does not require, and does not produce, market formation.*

In plain terms: when SpaceX launches a Starlink satellite, it is not selling a launch service on a market. Instead, it is moving a payload from one division of its own company to another. The price it assigns to that transaction tells you about SpaceX's internal accounting, not about the cost of launch as an economic commodity. Benchmarking competitors against this number is like benchmarking a restaurant's food costs against the price of a home-cooked meal; the comparison is structurally misleading because the two prices emerge from different economic relationships.

Recent estimates suggest that SpaceX's internal cost per Falcon 9 launch may be as low as $28 million, whereas the published price for external customers is $67 million—the price upon which competitors benchmark against . The real structural gap is against the *internal* cost.

Proposition 1 (Cadence as Architecture)**.** *SpaceX's cost advantage is not primarily a function of reusable technology as popularly discussed, but of* demand architecture*: the captive constellation that generates launch cadence no competitor can replicate independently of vehicle technology.*

In plain terms: even if Ariane, H3, or any emerging provider achieves full reusability tomorrow, they will still face a fundamental asymmetry vis-à-vis SpaceX in that they do not have a captive constellation generating 60–100 launches per year. Given that launch economics are driven by cadence, a reusable rocket that flies three times a year does not achieve the same unit economics as one that flies 165 times. This is because fixed costs such as a factory, workforce, launch infrastructure, and mission control, are amortized across each year's missions. At five to ten launches per year, fixed costs dominate regardless of vehicle type. At 165 launches, they nearly disappear.

Argument. Let $C(n)$ denote the effective cost per launch as a function of annual launch cadence n. Decomposing the total annual cost into fixed and variable components:

$$C_{\text{total}}(n) = F + v \cdot n$$

where F denotes fixed costs (factory, workforce, launch infrastructure) and v denotes marginal cost per additional launch. The effective cost per launch is:

$$C(n) = \frac{F}{n} + v \tag{1}$$

As $n \to \infty$, $C(n) \to v$. SpaceX operates at $n \approx 165$; while European providers are at $n \approx$ 7–12. For any realistic F, the per-unit cost differential is structural and cannot be closed by reusability alone. Ariane 6 with its design cadence of 10–12 launches per year will spread its fixed costs across an order of magnitude fewer missions than SpaceX, resulting in a per-launch cost that cannot converge with Falcon 9's pricing.

Even if you strip Starlink out of the equation, SpaceX's external launch cadence falls to perhaps 45–55 launches per year which is still substantial, but the cost base does not disappear with the Starlink missions. The pricing would therefore have to change for external customers.

> *Ariane and other launch providers are competing against an architecture, not a price.*

3. The Architecture Gap: A Formal Model of Market Formation

Markets do not "emerge" from good technology. Rather, they form when three distinct architectures *co-evolve*. In the space economy, only one architecture has been created so far.

3.1 Definitions

Definition 2 (Technical Architecture)**.** *Technical architecture* $T(t) \in [0, 1]$ *denotes the maturity of the hardware, software, and integration capabilities that define what a sector can do at time* t*.* $T(t)$ *is driven by R&D competition and engineering progress, and typically advances fastest of the three architectures.*

In plain terms: can the technology do "the thing"? In space, the answer is increasingly yes, considering that reusable launch vehicles work and autonomous docking has been demonstrated. So the technical architecture is mature and advancing rapidly.

Definition 3 (Market Architecture)**.** *Market architecture $M(t) \in [0,1]$ denotes the maturity of the conditions that allow repeated economic exchange at time t: elastic demand, stable pricing, modular products, standardized interfaces, specialized suppliers, and functioning supply chains.*

In plain terms: can you buy and sell the thing repeatedly, at predictable prices, from specialized suppliers? In space downstream services, the answer is currently no. There are no standard product categories, and no repeat customers outside Starlink. This leads to a missing pricing mechanisms for in-orbit services and a lack of cost curves, driven largely by every deployment being bespoke and the fact that nothing is comparable enough to learn from.

Definition 4 (Institutional Architecture)**.** *Institutional architecture $I(t) \in [0,1]$ denotes the maturity of the standards, certification regimes, liability frameworks, resource rights, logistics infrastructure, and governance arrangements that make economic activity legitimate and investable at time t.*

In plain terms: are there rules, standards, and infrastructure that allow economic activity to be insured, and scaled? In space, the answer is still largely no. There are no on-orbit docking or power interface standards, no liability regime for in-orbit servicing, or no jurisdiction framework for data hosted in LEO. The Outer Space Treaty, which is often erroneously cited here is a legal framework, not a market framework.

Definition 5 (System Readiness)**.** *System readiness is a weighted aggregate of market and institutional maturity:*

$$A(t) = w_M \cdot M(t) + w_I \cdot I(t) \tag{2}$$

where $w_M, w_I \geq 0$ and $w_M + w_I = 1$.

In plain terms: system readiness measures how prepared the non-technical parts of the system are, including the market conditions and the institutional infrastructure. Even if technology is brilliant, if nobody can buy it at a predictable price (market) and nobody knows whose law governs it (institutional), the brilliance is economically useless.

Definition 6 (Architecture Gap)**.** *The* architecture gap *is:*

$$\Delta(t) = T(t) - A(t) \tag{3}$$

When $\Delta(t) \approx 0$, the system is synchronized: markets can form, scale, and stabilize. When $\Delta(t) \gg 0$, the system is premature: technological capabilities outpace the supporting architectures and markets cannot form.

In plain terms: the architecture gap measures how far ahead the technology is relative to *everything else. A small gap means the ecosystem is ready to absorb what the technology can do; a large gap means the technology is pacing very far ahead while the market and institutional infrastructure needed to make it economically useful lags far behind. The gap is the core diagnostic of market or industry viability, as it tells you whether a sector is genuinely "early" (a small gap, that is closing) or structurally premature (a large gap, which is widening).*

3.2 Viability Conditions

Definition 7 (Viability Threshold)**.** *Let $\theta > 0$ denote the maximum tolerable misalignment beyond which markets cannot form. Market formation requires:*

$$\Delta(t) \leq \theta \qquad \textit{(viable)} \tag{4}$$

$$\Delta(t) > \theta \qquad \textit{(premature)} \tag{5}$$

In plain terms: θ is the maximum gap the system can tolerate and still form a functioning market. In sectors that have successfully industrialized, such as semiconductors, wind energy, and commercial aviation, the gap was kept within this threshold as the three architectures co-evolved. In sectors that remain premature, including nuclear fusion, quantum computing, hydrogen energy, and the space downstream economy—the gap exceeds it.

Differentiating equation (3) yields the *architecture-lag rate*:

$$L(t) = \frac{d\Delta}{dt} = \frac{dT}{dt} - \frac{dA}{dt} \tag{6}$$

- $L(t) > 0$: the gap is widening—prematurity is deepening.
- $L(t) < 0$: the gap is shrinking—the system is moving toward viability.
- $L(t) = 0$: the system is stuck in persistent misalignment.

In plain terms: $L(t)$ tells you which direction the sector is heading. If technology keeps advancing while market and institutional architectures stay flat, $L(t)$ is positive and the architecture gap widens. More R&D funding, more demonstration programs, more technology subsidies—all increase dT/dt without affecting dA/dt, and therefore widen the gap. Counterintuitively, but very important, that in this scenario the sector is not getting closer to viability, it's actually getting further away.

Definition 8 (Continuous Viability Function)**.** *The viability function provides a continuous* *measure of market formation potential:*

$$V(t) = 1 - \frac{\max\{0,\ T(t) - A(t) - \theta\}}{1 - \theta} \tag{7}$$

This yields $V(t) = 1$ for a fully viable sector, $V(t) = 0$ for a fully premature sector, and smooth decline as the architecture gap widens past the threshold.

3.3 Architecture Dynamics

The three architectures evolve at different rates, driven by distinct processes.

Proposition 2 (Architecture Velocity Ordering)**.** *In emerging sectors, the three architectures follow a strict velocity ordering:*

$$\frac{dT}{dt} > \frac{dM}{dt} > \frac{dI}{dt}$$

Technical architecture advances fastest, driven heavily by R&D competition and engineering optimism. Market architecture advances more slowly, requiring cost stability, modularity, and the emergence of repeatable demand. And institutional architecture advances the slowest of the three, as it is constrained by government regulatory processes, standards formation, public-infrastructure sequencing, and most of all—political coordination.

In plain terms: technology moves fast because engineers can iterate in labs and on launch pads rapidly. Markets move slower because they require customers, supply chains, and stable pricing—things that only emerge through repeated real-world transactions. Unsurprisingly, institutions move slowest of them all because they require consensus, legal processes, and political will; a structural problem with technology commercialization, as standards bodies do not sprint and public infrastructure takes decades. The problem arises when policy assumes all three can be accelerated by pushing the fastest one—typically technology—harder.

Argument. Each architecture contributes a necessary but not sufficient condition for market formation: technical architecture enables capability; market architecture enables repeatability and scale; and institutional architecture enables risk-sharing, and long-horizon investment. Misalignment across these trajectories—and not insufficient performance in any single layer—is the dominant failure mode in frontier technologies.

Corollary 1 (Gap Widening Under Technology-First Policy)**.** *Under the velocity ordering in Proposition 2, any policy intervention that accelerates dT/dt without commensurately accelerating dM/dt and dI/dt will widen the architecture gap $\Delta(t)$ and deepen prematurity.*

In plain terms: funding more launch vehicles and more technology demonstrations—the dominant mode of current space investment—makes the gap worse*, not better. this is because it pushes $T(t)$ further ahead while $M(t)$ and $I(t)$ remain constant or flat. The sector may look more technologically impressive, but it is in fact becoming less economically viable. This is not a failure of individual programs,. it is a structural consequence of investing in the wrong layer.*

4. Architecture Readiness in the Space Economy

To apply the model empirically, we score each architecture layer using observable indicators on a normalized 0–100 scale. These are the author's assessment based on publicly available data.

Table 1: Architecture Readiness Index (ARI) for the space economy. Author's analysis based on publicly observable indicators.

Architecture	**ARI**	**Status**	**Key Indicators**
Technical	87/100	High readiness	Reusable launch demonstrated at scale; autonomous docking operational; on-orbit robotics demonstrated; surface mobility prototypes tested.
Market	28/100	Shallow	Non-elastic demand; no pricing for downstream services; no repeat customers outside Starlink; no standard product categories; no cost curves.
Institutional	12/100	Largely absent	No resource rights framework; no on-orbit liability regime; no interface standards; no logistics infrastructure; no jurisdiction framework for orbital data or services.

Table 2: Space economy indicators, 2025. Sources: SpaceX manifest data, Payload Space analysis, Jonathan McDowell launch database.

Indicator (2025)	**Value**
SpaceX orbital launches	165
Functioning lunar economy segments	≈ 0
In-orbit servicing markets with repeat customers	≈ 0
Priced resource-extraction transactions	\$0
Starlink missions / total SpaceX launches (2024)	89 / 139
Estimated SpaceX internal cost per Falcon 9 launch	∼\$36M
Published price for external customers	\$67M

Computing the architecture gap with equal weights ($w_M = w_I = 0.5$):

$$A(t) = 0.5 \times 0.28 + 0.5 \times 0.12 = 0.20 \tag{8}$$

$$\Delta(t) = 0.87 - 0.20 = 0.67 \tag{9}$$

For any reasonable viability threshold—which in sectors that have successfully industrialized (semiconductors, onshore wind, commercial aviation) tends to fall between 0.15 and 0.30—the space economy is deeply premature. The architecture gap of 0.67 exceeds even the most generous threshold by a factor of two to four.

> *The gap is widening for commercial space industries, as technical architecture continues to advance rapidly while market and institutional architectures are barely moving. This is the defining signature of a premature sector.*

5. Cost Curves Require More Than Just Technology

A persistent confusion in space economics is the assumption that learning curves will emerge naturally from repeated launches, however the empirical record in other sectors shows why this assumption is structurally incorrect.

5.1 Wright's Law and Its Architectural Prerequisites

Definition 9 (Learning Curve). *Wright's Law holds that unit costs decline by a constant fraction each time cumulative production doubles. If $C(n)$ is the cost of the nth unit and*

$C(1)$ *is the cost of the first:*

$$C(n) = C(1) \cdot n^{-b} \tag{10}$$

where $b > 0$ *is the learning parameter. The learning rate* $r = 1 - 2^{-b}$ *measures the percentage cost reduction per doubling of cumulative output.*

In plain terms: every time you double the number of units produced, the cost should fall by a fixed percentage. Semiconductors show a learning rate of roughly 35–45 percent, while onshore wind turbines show roughly 15–23 percent. These are among the most powerful cost-reduction dynamics in economic history, and they all share a common structural prerequisite in that the deployments from which learning accumulates must be comparable.

Proposition 3 (Architectural Prerequisite for Learning Curves)**.** *Wright's Law operates only when deployments are standardized enough to produce comparable data. Formally, the learning parameter* b *in equation* (10) *is endogenous to market architecture. Let* $\mu(t) \in [0, 1]$ *denote the degree of deployment standardization at time* t*. Then:*

$$b(\mu) = b_{\max} \cdot \mu(t) \tag{11}$$

where $b_{\max}$ *is the sector's maximum learning parameter under full standardization. As* $\mu(t) \to 0$ *(fully bespoke deployments),* $b(\mu) \to 0$ *and no cost curve forms regardless of cumulative output.*

In plain terms: clearly, you cannot learn from an experience if every experience is different. Consider a semiconductor fab that is running the same process on the same wafers with the same specifications learns from every run. Now think about an in-orbit servicing mission that is custom-designed from scratch, with bespoke interfaces, unique power configurations, and one-off integration—what is it learning from? Nothing that applies to the next mission, because the next mission is equally bespoke. The exponent in the cost curve is not a property of the technology, as with other industries, but a property of the system around the technology.

Argument. Semiconductors developed their learning curve because fabrication was standardized: the same wafer sizes,me photolithography processes, and interface specifications across an entire industry. Thus, each doubling of cumulative production produced comparable data from which process improvements could be extracted. Similarly, onshore wind developed its learning curve because turbine designs converged on standardized platforms with modular components and interchangeable subsystems. In both cases, standards, supply chains, and repeat customers *co-evolved* with the technology. Learning could accumulate because deployments were comparable.

Space downstream services show no learning curve yet, which is not because the technology is immature but because there are no comparable, repeated deployments from which learning can accumulate. This is largely because every mission is bespoke, and without standardization, $\mu(t) \approx 0$, the exponent b approaches zero; costs do not decline with volume.

Without standardization, there is no cost curve, and hence there is no industry.

5.2 Cross-Sector Evidence

The contrast between the space industry and other economically viable technologies is stark. In semiconductors and wind, the three architectures co-evolved in that technical capabilities advanced alongside market standardization and institutional frameworks (industry standards bodies, certification regimes, grid-connection protocols, renewable energy mandates). The cost curve was not a consequence of technology alone but of a phenomena called *architectural synchrony.* This is the economic equivalent of "it takes a village".

However, in space downstream services, the architecture gap is so wide that the structural prerequisites for learning—comparable deployments, standardized interfaces, modular product architectures, repeat customers—do not exist. Therefore, tech investment continues to increase $T(t)$, but the cost curve cannot form because $M(t)$ and $I(t)$ remain too low to produce the standardization without which learning-by-doing cannot operate.

6. Case Study: Space-Based Data Centres

Space-based data centers illustrate the architecture gap, as the technical case is compelling while the binding constraints are entirely architectural.

6.1 The Pitch

The case for moving compute to orbit rests on four advantages: (1) continuous solar power with no day/night cycle and no grid dependency; (2) no cooling costs, since waste heat radiates directly to space, eliminating roughly 40 percent of terrestrial data center operating expenditure; (3) sovereign compute physically outside any national jurisdiction; and (4) falling launch costs that make the economics tractable for the first time. The EU's ASCEND study priced this opportunity at approximately 10 billion.

6.2 Architecture Readiness Assessment

Table 3: Architecture Readiness Index for space-based data centres. Sources: EU ASCEND study (2021), Lonestar Data Holdings, ISS National Lab.

Architecture	**ARI**	**Assessment**
Technical	48/100	Compute hardware works in orbit—demonstrated on ISS and commercial satellites. But no integration environment exists: no standard rack form factor, validated thermal rejection supply chain at scale, power distribution standards. The system does not exist.
Market	8/100	No pricing mechanism exists for a rack-hour in orbit. No enterprise customer has contracted for orbital compute. Ground station data offload capacity is severely limited—you cannot run a data centre if you cannot retrieve the data reliably. Latency is worse than terrestrial edge for most enterprise use cases.
Institutional	4/100	Whose law governs data hosted in LEO? Which jurisdiction handles a data breach? What constitutes interference with a competitor's orbital data centre? The Outer Space Treaty provides no answers. There is no insurance market, no liability framework, or data residency standard that exist for orbital compute.

6.3 Computing the Gap

Applying the framework:

$$T = 0.48 \tag{12}$$

$$A = 0.5(0.08) + 0.5(0.04) = 0.06 \tag{13}$$

$$\Delta = 0.48 - 0.06 = 0.42 \tag{14}$$

The sector is technically plausible and architecturally premature, and no amount of additional hardware development will create this market. The missing components—pricing mechanisms, jurisdiction frameworks, data offload infrastructure, interface standards—are market and institutional architecture problems.

Proposition 4 (Prematurity of Space-Based Data Centres)**.** *Space-based data centers satisfy the prematurity condition* $\Delta(t) > \theta$ *for any* $\theta \leq 0.30$*. Market formation is not precluded by technological immaturity but by the absence of market architecture (no pricing, no demand, no data offload) and institutional architecture (no jurisdiction, no liability, no standards).*

In plain terms: this is not a sector waiting for better hardware, although investors and governments alike are pushing for it. The hardware is close to being technologically viale, but what does not exist is any mechanism for pricing orbital compute, legal frameworks for governing data in orbit, or standards that would allow a supply chain to form. Until those are built, the sector cannot move from "technically plausible" to "economically viable," regardless of how much additional R&D investment flows into the technical layer.

7. Where the Space Economy Actually Needs to Be Built

If we assume that the architecture gap model is correct, then the investment priorities for the space economy are actually *inverted* from current best practice that we've seen in other industries, in that governments and investors keep funding more launch vehicles and upstream technology, which is at the wrong layer.

The policy objective should be to increase dA/dt—the rate at which system readiness grows—*rather than* dT/dt, which is already high. Closing the architecture gap therefore requires:

$$L(t) = \frac{dT}{dt} - \frac{dA}{dt} < 0 \tag{15}$$

And this condition is only met when market and institutional architectures both progress faster than technology. Current policy does in the space sector in fact does the opposite.

7.1 Four Architecture-Building Interventions

7.1.1 Standards and Interfaces

On-orbit docking, data, and power interfaces need to be standardized before a supply chain can form; yet no private actor has the incentive to do this unilaterally because the returns to standardization are diffuse and non-excludable. This is, therefore, a *public good* problem. The analogy is to the internet's TCP/IP protocols, in that the US government funded the creation of open interface standards that enabled an ecosystem. The space economy needs an equivalent standards mechanism for on-orbit power delivery, data interconnection, and physical docking.

7.1.2 Resource Rights Architecture

You cannot build a lunar economy without legal clarity on who owns the extracted resources, who is liable when things inevitably go wrong, and what constitutes trespassing or interfering during such activities. The Outer Space Treaty (1967) declares that no nation may claim sovereignty over celestial bodies, but says next to nothing about commercial resource extraction, the property rights for extracted materials, or any liability for interference between operators in space. The Artemis Accords represent a partial step towards this need, but they are bilateral agreements between willing parties, and not a binding multilateral framework. To date, no extractive industry in history has scaled without a resource rights regime.

7.1.3 Anchor Demand, Not Technology

The analogy to semiconductor strategy should be carefully considered, as the United States did not build a semiconductor industry by funding chip R&D and hoping demand would follow. Instead, it built a nascent market by becoming a large and predictable buyer through defense procurement and the space program. This institutional demand generated the cadence from which learning curves, supply chains, and eventually commercial markets emerged.

Activities including in-orbit servicing and lunar surface infrastructure need institutional demand to generate cadence, as without it, there is no learning curve; without this, there is no cost curve and therefore no industry.

7.1.4 Price Surface Infrastructure

Power, communications, and mobility at destination are *public infrastructure problems*. No company will build a lunar power grid without a utility-style model that provides predictable returns on infrastructure investment; someone has to price the infrastructure.

The analogy to use here is of terrestrial infrastructure, including roads, power grids, telecommunications networks, and water systems that are priced as public or regulated utilities because no private actor can capture sufficient returns to justify building them unilaterally, yet their absence eliminates all downstream economic activity.

7.2 The Vertical Integration Workaround

SpaceX's response to the architecture gap is extraordinary but non-replicable by competitors that are needed to accelerate the market and institutional aspects of the industry. By building Starlink, SpaceX manufactured its own demand while generating its own cadence, and therefore created an internal architecture that substitutes for the missing system architecture. This is a private solution to a public architecture problem.

Proposition 5 (Non-Generalisability of Vertical Integration)**.** *SpaceX's vertical integration reduces the effective architecture gap* for SpaceX internally *but does not reduce* $\Delta(t)$ *for the broader sector. If a supply chain forms around a transfer price, no learning curve can emerge from its captive constellation's internal cost structure. Thus, a lack of cost curve propagates to the wider market from a vertically integrated monopoly.*

In plain terms: SpaceX solved its own problem very creatively, but it did not solve the sector's problem, as its launch cadence is not available to other firms, nor its cost structure. Its pricing is not a market signal as popularly discussed by industry and investment analysts. In fact, the space economy cannot be built by replicating SpaceX, because SpaceX's advantage is structural and monopolistic, not technological and replicable.

8. Testable Predictions

The architecture gap framework generates a set of empirical predictions that distinguish it from the standard "early-stage" interpretation of the space economy's development trajectory.

Proposition 6 (Persistence of Prematurity Under Technology-First Investment)**.** *Continued investment in technical architecture (*$T(t)$*) without commensurate investment in market and institutional architectures (*$M(t)$*,* $I(t)$*) will widen the architecture gap* $\Delta(t)$ *and deepen the sector's premature status. Thus, no downstream market segment will achieve elastic demand, stable pricing, or a functioning cost curve during this period, regardless of how cheap launch becomes.*

In plain terms: if you keep funding rockets and R&D without building standards, demand architectures, and institutional frameworks, the space economy in 2035 will look exactly like the space economy in 2025; it will be one with extraordinary technology provided by monopolistic firms, but with zero functioning downstream markets—therefore exhibiting a larger architecture gap with a harder path to closing such a gap.

Proposition 7 (Learning Curve Emergence as Architecture Signal)**.** *The first evidence that the space economy is transitioning from a premature to a viable market will not be a technological breakthrough, as those already exist, but the emergence of a learning curve in a downstream service category, acting as a signal that deployments have become standardized enough to produce comparable cost data. The model predicts this will occur first in the service category with the highest market architecture score, not the highest technical architecture score.*

In plain terms: you will know the space economy is real when someone can show you a cost curve for a downstream service—not a cost curve for launch, but for in-orbit servicing, satellite manufacturing, or orbital logistics. That curve will appear first wherever market architecture is strongest, and where there are repeat customers and standardized interfaces. In other words, one needs to watch the architecture scores, not the technology scores, to see how close the market is to forming.

Proposition 8 (Vertical Integration as Gap Diagnostic)**.** *The prevalence of vertical integration in a frontier sector is positively correlated with the magnitude of the architecture gap. This means that firms will vertically integrate when the market architecture required for specialized exchanges do not exist; and the extent of vertical integration therefore serves as a proxy measure for* $\Delta(t)$*.*

In plain terms: the fact that SpaceX had to build its own constellation to generate its own demand is not evidence of visionary strategy, but of a missing market. If market architecture already existed, then SpaceX would not need a captive constellation of its own. Vertical integration is not a business model choice in this context, it is a symptom of architectural failure.

9. Conclusion

SpaceX did something extraordinary by reducing the cost of reaching orbit by a factor of twenty, conducting 165 orbital launches in a single year, and vertically integrating the past the architecture gap by building its own downstream market.

However, this is not a model for a space economy, it is simply a workaround to having to create any space economy at all.

Why? Because it does not reduce $\Delta(t)$ for the broader sector. SpaceX's vertical integration creates an internal architecture that substitutes for the missing system architecture, but no other firm can replicate it. This further means that no supply chain can form around it and, importantly, no cost curve can emerge from it. Thus the internal transfer price is little more than a structural artefact of a monopoly architecture.

The architecture gap model explains why 165 launches and a 20× cost reduction produced one incredible company, but with zero functioning downstream markets: because the binding constraint was never the cost of launch but the absence of the market and institutional architectures—without which economic activity cannot form. The gap between technical architecture ($T = 0.87$) and system readiness ($A = 0.20$) is 0.67, whcih is well beyond any plausible viability threshold.

Closing this gap requires investment in the layers that are boring and that nobody with an interest in high-tech rockets actually wants to build: standards and interface specifications, resource rights frameworks, anchor demand programmes, and priced surface infrastructure. These are unglamorous parts of the industry and they produce neither photogenic launches nor billion-dollar valuations. However, they are the structural prerequisites without which the space economy will remain what it is today: a technologically extraordinary, commercially premature sector in which one company has found a private solution to a public architecture problem.

> *In short, we have a very impressive launch service; alas we do not have a space economy.*

References

Adner, R. (2006) 'Match your innovation strategy to your innovation ecosystem.' *Harvard Business Review*, 84(4), pp. 98–107.

Baldwin, C.Y. and Clark, K.B. (2000) *Design Rules: The Power of Modularity.* Cambridge, MA: MIT Press.

EU ASCEND Study (2021) *Advanced Space Cloud for European Net zero emission and Data sovereignty.* European Space Agency.

Hughes, T.P. (1983) *Networks of Power: Electrification in Western Society, 1880–1930.* Baltimore: Johns Hopkins University Press.

Jacobides, M.G., Cennamo, C. and Gawer, A. (2018) 'Towards a theory of ecosystems.' *Strategic Management Journal*, 39(8), pp. 2255–2276.

Langlois, R.N. (2003) 'The vanishing hand: the changing dynamics of industrial capitalism.' *Industrial and Corporate Change*, 12(2), pp. 351–385.

Mazzucato, M. (2013) *The Entrepreneurial State: Debunking Public vs. Private Sector Myths.* London: Anthem Press.

North, D.C. (2005) *Understanding the Process of Economic Change.* Princeton: Princeton University Press.

O'Sullivan, S. (2025a) 'Market formation as a systems engineering problem.' CEI Working Paper No. 02.

O'Sullivan, S. (2025b) 'Institutions as coordination architectures.' CEI Working Paper No. 01.

O'Sullivan, S. (2025c) 'The subsidy trap.' CEI Working Paper No. 04.

Simon, H.A. (1962) 'The architecture of complexity.' *Proceedings of the American Philosophical Society*, 106(6), pp. 467–482.

Teece, D.J. (1986) 'Profiting from technological innovation.' *Research Policy*, 15(6), pp. 285–305.

Wright, T.P. (1936) 'Factors affecting the cost of airplanes.' *Journal of the Aeronautical Sciences*, 3(4), pp. 122–128.