

SpaceX's Launch Price Isn't a Market Price

Why No One Can Compete With SpaceX, and Why That's Everyone's Problem

Sinéad O'Sullivan

IDEA IN BRIEF

THE ILLUSION

SpaceX has made launch dramatically cheaper, and the industry assumes this is a cost curve that everyone will eventually ride down. Similar beliefs include that reusability works, market competition will follow, and access to space is being democratised.

THE REALITY

SpaceX's price is not a *market price*, it is a *transfer price* set by a vertically integrated company whose largest customer is itself. Starlink generates more internal launch demand than the rest of the commercial market combined, amortizing fixed costs across a cadence no competitor can replicate. Hence, other providers are failing because the benchmark is set by a company playing a structurally different game.

THE CONSEQUENCE

Governments awarding contracts on price are rationally optimizing today and destroying their strategic options for tomorrow. A sustainable launch market requires deliberate intervention, or there will be one provider left, and it won't stay cheap.

The launch industry has a story it likes to tell, repeated often and breathlessly:

SpaceX made launch cheap; the cost curve is coming down; competition will follow; reusability works; access to space is being democratised.

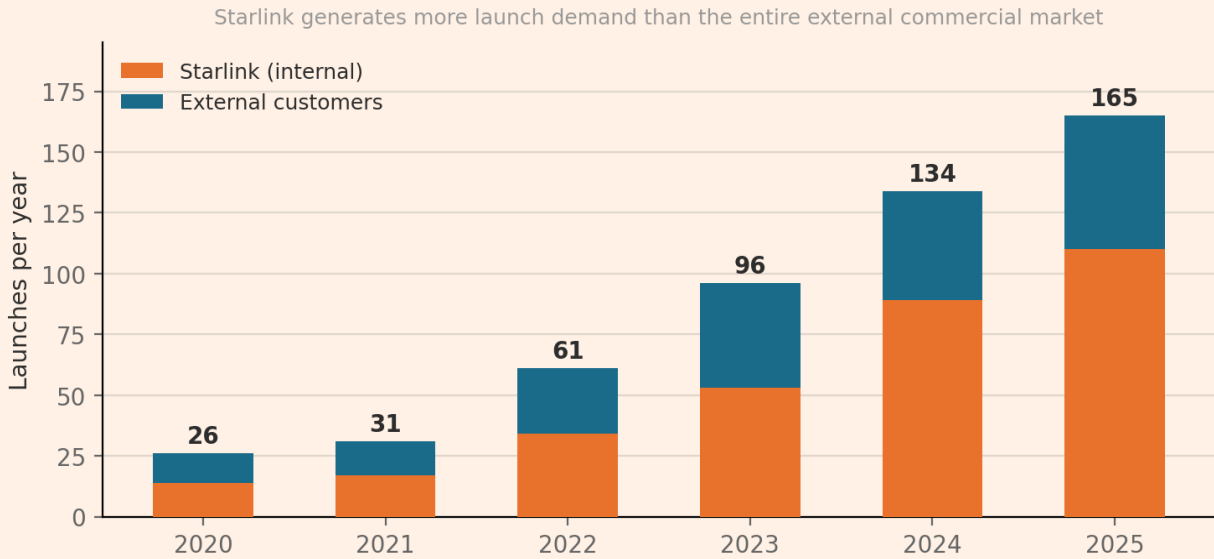
The first part is certainly true, that SpaceX has dramatically reduced the cost of reaching orbit. Falcon 9 launches routinely at a fraction of what legacy providers charge, and in 2025, SpaceX conducted one hundred and sixty-five orbital launches—more than twice the total of China's entire space programme and roughly half of all

launches worldwide. So yes, the technology is real and the achievement is extraordinary.

But the rest of the story? It's just wrong.

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. Starlink—SpaceX's satellite internet constellation—now accounts for roughly two-thirds of the company's launches, and in 2024, eighty-nine of SpaceX's one hundred and thirty-four Falcon missions carried Starlink payloads. Only eighteen were standard commercial launches, meaning that no external customer base in history has generated this kind of cadence demand as Starlink has. Thus, SpaceX did not find demand at this scale, it manufactured it.

SpaceX Launch Cadence: Internal vs External Demand



Source: SpaceX manifest data, Payload Space analysis. Starlink missions are zero-revenue launches that generate the cadence underpinning SpaceX's cost structure.


That changes what the price actually represents, because when SpaceX offers a commercial or government customer a Falcon 9 launch at sixty-seven million dollars, the figure reflects economics no other provider can replicate. SpaceX's fixed costs—factory, workforce, launch infrastructure, mission control—are amortized across a launch cadence three to five times higher than any competitor's. Consider that its reusable boosters have flown twenty or more times each, spreading hardware costs across missions in a way expendable or early-reuse vehicles cannot match. And its vertical integration—in-house engines, avionics, structures, and software—eliminates the margins traditional contractors pay to their supply chains.

Recent estimates suggest SpaceX's internal cost per Falcon 9 launch may be as low as thirty-six million dollars, meaning that the published price of sixty-seven million already undercuts every major competitor. Therefore the actual cost advantage is nearly twice as large as the price gap suggests.

However, if Starlink were removed from the equation, SpaceX's external launch cadence would resemble that of a conventional provider—perhaps forty-five to fifty-five launches per year. Substantial, certainly, but the fixed cost base does not shrink with the disappearance of Starlink missions. Under those conditions, the pricing structure would look very different.

The economics reflect a particular organizational structure, not a universal cost curve.

Published Launch Prices vs SpaceX Internal Cost



Published launch prices and estimated SpaceX internal cost. Competitors benchmark against the list price (\$67M). The real structural gap is against the internal cost (~\$36M). Sources: SpaceX, Ariespace, JAXA, ULA public pricing; Motley Fool estimates for internal cost.

WHY COMPETITORS AREN'T FAILING ON TECHNOLOGY

The standard explanation for why European, Japanese, and other launch providers are struggling is that they are behind on technology, that they need reusability, to iterate faster, and to behave more like SpaceX.


Some of this is true, but it misses the structural point that even if Ariane, H3, or any emerging provider achieves full reusability tomorrow, they will still face a fundamental asymmetry: they do not have a captive constellation generating sixty to a hundred launches per year.

Importantly, launch economics are driven by cadence. This means that a reusable rocket that flies three times a year cannot achieve the same unit economics as a reusable rocket that flies a hundred

and sixty-five times. The marginal cost of an additional launch drops dramatically with volume—maintenance crews stay employed, launch pads stay utilised, manufacturing lines stay warm, and institutional knowledge compounds. Thus, a provider launching five to ten times per year is operating a fundamentally different business than one launching over a hundred, regardless of whether both vehicles are technically reusable.

Ariane 6—Europe's new heavy-lift vehicle—managed seven launches in 2025, its first full year of operation. Even assuming it reaches its design cadence of ten to twelve per year, its fixed-cost base will be spread across an order of magnitude fewer missions than SpaceX's. The resulting per-launch cost simply cannot converge with Falcon 9's pricing. This is not an engineering problem. The limiting factor is very much scale economics rather than engineering capability.

The Cadence Arithmetic



Illustrative model of effective cost per launch as a function of annual cadence. Fixed costs (factory, workforce, launch infrastructure) are spread across each year's missions. At 5-10 launches per year, fixed costs dominate regardless of vehicle type. At 165 launches, they nearly disappear.

The same logic applies to Japan's H3, to ULA's Vulcan, and to virtually every new entrant. They are competing against a price that is structurally

unreplicable because the demand architecture behind SpaceX cannot be matched.

*They are not competing against a price.
They are competing against an architecture.*

WHAT THE LAUNCH INDUSTRY ACTUALLY IS

If a market is defined by multiple suppliers, price discovery, customer switching, and the possibility of competitive entry, then the launch industry today is not a market. Indeed, it is little more than a *procurement relationship with SpaceX that the world has agreed to call a market.*

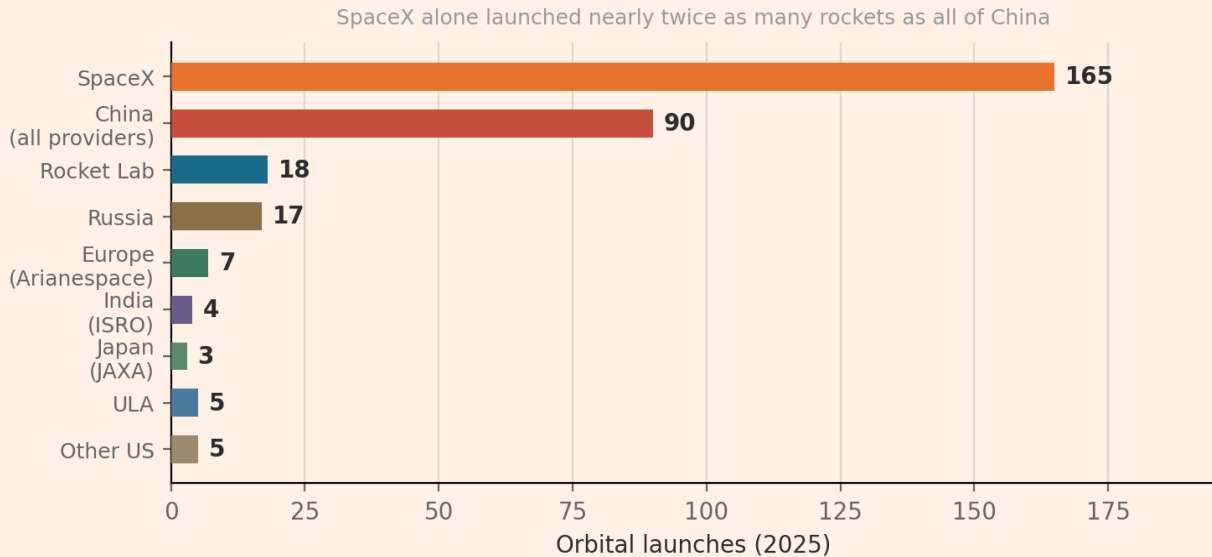
SpaceX currently launches the overwhelming majority of global commercial payloads. As such, the company sets the price and customers have limited alternatives—and the alternatives that exist are significantly more expensive, less proven, or both. Therefore there is no meaningful price discovery

because there is no meaningful competition: there is SpaceX's price, and there is the price everyone else wishes they could match.

This is, of course, not inherently a problem if you believe SpaceX will remain reliable, available, and reasonably priced indefinitely. But that belief requires trusting that a single private company, with a single point of leadership, will continue to serve all customers—including those whose interests may conflict with SpaceX's own commercial ambitions—without leveraging its monopoly position.

The history of monopoly pricing is not encouraging on this point.

Global Orbital Launches by Provider, 2025



Global orbital launches by provider, 2025. SpaceX alone conducted more launches than all non-US, non-Chinese providers combined. Europe's entire launch output (7) was fewer than a single month of SpaceX Starlink missions. Source: Jonathan McDowell, Payload Space.

THE STRATEGIC AUTONOMY PROBLEM

Every government that relies on SpaceX for national security launches has a dependency it wishes did not exist.

The United States has partially mitigated this through the National Security Space Launch programme, which maintains multiple certified providers, but the gap between SpaceX and its competitors is widening, with the economic logic of awarding contracts to the cheapest provider creating persistent pressure to consolidate toward a single source.

For Europe, the situation is more acute. Ariane 6's delays and cost overruns have left European institutional customers with limited sovereign access to space. ESA member states therefore face a stark choice of continuing to fund a domestic launch capability that cannot compete on price with SpaceX, or accepting dependence on an American provider for sovereign missions. Neither option is desirable; the first acts as an expensive subsidy for an uncompetitive industry, while the second means a NATO ally's access to space runs through a single company's manifest. Meanwhile, ESA has already

launched several of its own science missions on Falcon 9—a development that would have been unthinkable a decade ago.

Japan faces a similar calculus with H3, as does India, though its lower cost structure and growing domestic demand provide more insulation from this undesirable trade-off. For smaller spacefaring nations, the question is not even open; they have already defaulted to SpaceX because no alternative exists at a comparable price.

THE PROCUREMENT TRAP

Despite procurement being one of the least interesting aspects of the exciting launch industry, it is the mechanism that deserves the most scrutiny. Government procurement—in the US, in Europe, and increasingly elsewhere—is designed to select for lowest cost or best value on a per-mission basis. This is standard practice and, in most industries, an entirely sensible approach.

In the launch industry, however, it is self-defeating.

When the lowest-cost provider's pricing reflects a structural position that no competitor can replicate, awarding on price does not create competitive

pressure. In fact, it does the very opposite of creating competitive extinction, as every contract awarded to SpaceX on cost grounds reduces the launch cadence available to alternative providers, further degrading their unit economics, and making them less competitive on the next contract. The cycle is self-reinforcing.

This behavior is not a market that is disciplining inefficient providers. In contrast, it is a *monopoly architecture* using captive demand to set prices that extinguish competition—not necessarily through predatory intent, but through structural advantage, even if the effect is the same.

Every contract awarded on price alone reduces the cadence available to alternative providers, degrades their economics, and makes them less competitive next time. The cycle is self-reinforcing.

WHAT A SUSTAINABLE LAUNCH MARKET REQUIRES

It is important to note that the launch industry will not self-correct into a competitive market, as the structural advantages of vertical integration plus captive demand are too large and compounding over time.

Thus, a sustainable launch market requires deliberate architectural interventions:

First, anchor demand for non-SpaceX providers. Governments must commit to multi-year launch contracts that give alternative providers the cadence they need to bring unit costs down. Without guaranteed volume, no business case closes, no manufacturing line is efficient, and no learning curve forms. Cadence is not a reward for competitiveness; it is a prerequisite for it.

Second, procurement frameworks that price strategic value. A launch from a sovereign or allied provider is worth more than a launch from a monopolist, because it preserves future optionality, and procurement should reflect this reality. As such, the cheapest launch is not the best launch if it eliminates the possibility of alternatives.

Third, honest accounting of what SpaceX's price actually represents. It is not the cost of launch; it is the cost of launch to a company that has built a demand architecture no one else possesses. Benchmarking competitors against this price is benchmarking against a structural anomaly, not a market signal.

Fourth, investment in demand architectures—not just launch vehicles. Europe does not need a better rocket, it needs a reason to launch frequently. Constellation programmes, in-orbit servicing, and Earth observation networks—these generate the cadence that makes launch economics work. Without demand, reusability is just an engineering achievement. With demand, it can become a business, and eventually an industry.

THE DEPENDENCY

The launch industry in 2026 has all the appearances of a functioning market: customers, providers, contracts, prices. But what it lacks is the structural foundation that makes a market *actually* sustainable: multiple viable suppliers, price discovery, and the possibility of competitive entry.

Until that foundation is built, what we have is a dependency. And dependencies, eventually, get priced accordingly.

The frameworks in this article—architecture lag, demand architectures, and the structural conditions for market formation—are developed formally in two companion papers by Sinéad O'Sullivan. Copies available on request: s@sinead.co