The $1.8 Trillion Question

Inside the Economics of SpaceX

Dr. Arjuna Sathiaseelan
Space Economy Initiative, Cambridge Judge Business School
a.sathiaseelan@jbs.cam.ac.uk

June 2026

Abstract

In June 2026, SpaceX completed the largest initial public offering in history, entering the public market at a valuation close to $1.8 trillion against annual revenue of roughly $18.7 billion. This paper asks what set of assumptions is required to justify that figure. It traces the economics of the launch industry from the high fixed costs of the expendable era, through the procurement reforms at NASA that enabled fixed-price commercial spaceflight, to the reusable, high-cadence model that has lowered the advertised cost of reaching orbit by roughly a factor of twenty. It then argues that the valuation is not a claim about the launch business, which is too small to support such a number, but a bet on its second-order consequences: a dominant satellite-broadband utility in Starlink, the promise of fully reusable heavy lift in Starship, and a nascent thesis around orbital computing following the acquisition of xAI. Separating the demonstrated business from the speculative options, the paper concludes that the valuation is conditional, defensible only if a demanding sequence of technical and commercial milestones is met, and notes that independent discounted cash-flow analyses place fair value materially below the offering price.

Keywords: SpaceX; space economy; launch costs; reusability; Starlink; valuation.

1 Introduction

On 11 June 2026, Space Exploration Technologies Corp. priced what financial press described as the largest initial public offering in history, selling around 555.6 million shares at $135 each, raising about $75 billion, and stepping onto the public market at a valuation just under $1.8 trillion [1, 2]. That single figure exceeded the proceeds of any previous IPO, and reports noted that the deal immediately placed SpaceX among the largest listed companies in the United States, ahead of Tesla and in the same tier as Microsoft and Alphabet by market capitalisation [1, 2].

For a business founded to build rockets, this is an extraordinary place to arrive. A valuation near $1.8 trillion against 2025 revenue of about $18.7 billion implies a price-to-sales multiple in the mid–90s [3]. By comparison, Nvidia, the defining hardware story of the AI era, has in recent years traded around 30 times revenue, and high-growth cloud infrastructure companies typically sit between roughly 15 and 20 times.[1] SpaceX did not simply price at a premium. It priced at a multiple that has almost no precedent at this scale.

This paper is written to answer one question in plain terms: what has to be true for $1.8 trillion to make sense?

To get there, we have to understand the company as three businesses braided together. There is the launch business, which sells rides to orbit and has already rewritten the cost structure of an entire industry [4]. There is Starlink, the satellite broadband network that now generates the

[1] Representative multiples based on public market data for major US technology stocks as of mid-2026.

majority of revenue and essentially all of the operating profit [3]. And, since the acquisition of xAI in early 2026, there is an artificial intelligence layer that the company increasingly uses to frame its long-term story around orbital computing and a total addressable market it describes in the tens of trillions of dollars [3].

The tension running through the whole document is this. The hard, demonstrated achievement is a launch company that, by standard NASA cost comparisons, has lowered the advertised cost of reaching orbit by roughly a factor of twenty relative to the Space Shuttle [4]. The valuation, however, is not built on launch alone. It is built on what cheap, reliable, high-cadence launch makes possible: a broadband utility, and beyond it, the promise of computing infrastructure in space. The market is paying for the second-order consequences of an engineering breakthrough, not for the breakthrough itself.

We will work through that logic step by step. We begin with why space was so expensive for so long, because the size of SpaceX's achievement only makes sense against the cost of the old way of doing things.

2 The Old Price of Orbit

For most of the space age, getting to orbit worked like a toll road with a single operator and no competition. The price was whatever the operator needed it to be, and there was nowhere else to drive.

The physics sets a hard floor. Reaching low Earth orbit means accelerating to roughly 28,000 kilometres per hour while climbing out of most of the atmosphere.[2] That demands enormous quantities of propellant and a vehicle engineered to the very edge of what materials can survive. None of that is negotiable. But the physics was never the expensive part. The expensive part was the economic model layered on top of it.

Consider the benchmark that dominated the industry for three decades. NASA analysis indicates that the Space Shuttle cost on the order of $1.5 billion to place about 27,500 kilograms into low Earth orbit, which works out to roughly $54,500 per kilogram [4]. At that price, a one-litre bottle of water delivered to orbit would cost more than a luxury car. Every kilogram of a satellite, every bracket and bolt, carried a launch cost that dwarfed the value of the hardware itself.

Three features of the old model kept prices high.

The first was expendability. Rockets were built, flown once, and destroyed. Imagine an airline that scrapped a 747 after a single flight from London to New York. No ticket price could ever be reasonable, because the entire cost of manufacturing the aircraft had to be recovered from one trip. Launch worked exactly this way. The most sophisticated machines humans had ever built were treated as packaging.

The second was low flight rate. Because each vehicle was thrown away, and because each was hand-built, launches were rare. A handful of flights a year meant the fixed costs of factories, launch pads, and standing armies of engineers were spread across almost nothing. Low volume kept unit costs high, and high unit costs suppressed demand, which kept volume low. It was a trap that reinforced itself.

The third was the procurement structure, which we examine in detail in Section 4. Governments bought rockets on cost-plus contracts, paying the contractor its documented costs plus a guaranteed margin. Under that arrangement, spending more money was rewarded, not penalised. There was no commercial pressure to make anything cheaper, because cheaper meant lower revenue for the supplier.

[2] Standard orbital mechanics

The result was a market that stayed small for decades. Satellites were engineered to last fifteen or twenty years not only because replacement was hard, but because launching a replacement was ruinously expensive. Ambitious ideas like large satellite constellations, orbital manufacturing, or routine human spaceflight were dismissed as science fiction. They were not technically impossible. They were economically impossible at $50,000 per kilogram [4].

The toll road had one more feature worth naming. The customers were almost entirely governments and a small number of large telecommunications and defence operators. There was no broad commercial market, because the price excluded everyone else. Space was a place you visited if you were a superpower or a satellite-television company, and almost no one else.

To break this open, someone had to attack all three features at once: stop throwing rockets away, fly far more often, and replace the procurement model that rewarded waste. That is the story of the sections that follow.

3 A Company Built Differently

SpaceX was founded in 2002 with a goal that sounded absurd in a financial pitch meeting: to reduce the cost of access to space enough to make humanity multi-planetary. The commercial logic came second. The mission came first, and the business was built as the engine that would fund it [5].

The early years were defined by the small Falcon 1 rocket and a sequence of failures that nearly ended the company. The first three orbital launch attempts failed. By 2008, the company was close to running out of money, and the fourth Falcon 1 flight was effectively the last chance; it reached orbit, and shortly afterwards NASA awarded a Commercial Resupply Services contract that gave SpaceX the cash and the credibility to survive [5]. The company has been candid since that it was weeks from collapse.

Several decisions made in those early years explain almost everything that came later.

The first was vertical integration. Where the traditional aerospace model relied on sprawling networks of subcontractors, each adding cost and margin, SpaceX chose to build as much as possible in-house: engines, airframes, avionics, and eventually the satellites themselves [5, 4]. This is unusual and capital-intensive, but it gave the company control over cost, quality, and iteration speed. When you make your own engines, you can redesign them on your own schedule rather than negotiating with a supplier whose incentives run the other way.

The second was a culture of rapid, iterative engineering borrowed from software rather than from traditional aerospace. Build a version, fly it, learn from how it fails, and build the next version quickly. Established aerospace prized getting everything right on paper before any hardware flew, which is safer and far slower. SpaceX accepted visible failures as the price of fast learning. The Starship programme, with its public sequence of explosions followed by steady progress, is the clearest expression of this approach [5].

The third was the insistence on reusability as a first-class design goal rather than an afterthought. The company spent years and many crashes learning to land an orbital-class booster, achieving the first successful landing in December 2015 [5, 4]. Most of the industry regarded this as a stunt. It was not. It was the single change that would eventually break the toll-road economics described in the previous section.

The fourth, and least appreciated at the time, was the decision to become its own largest customer. Rather than waiting for a commercial market to materialise at lower launch prices, SpaceX created the demand itself by building Starlink. This solved the volume problem from the inside. If the world would not yet buy enough launches to justify a high flight rate, SpaceX would manufacture the need by launching its own satellites by the thousand [3, 6].

Each of these choices looked questionable in isolation and risky in combination. Together they produced a company structurally unlike its competitors, one that controlled its own supply chain, learned faster, reused its hardware, and generated its own demand [4, 3]. The valuation the market assigned in 2026 is, in large part, a judgement about how durable that combination of advantages will prove.

4 The Procurement Revolution

It is tempting to tell the SpaceX story as a pure tale of private enterprise outrunning a sleepy government. That version is wrong, and missing it leads to a misreading of the whole industry. The most important enabling decision was not made at SpaceX. It was made at NASA, when the agency changed how it bought launch services [4].

The traditional model was the cost-plus contract. The government specified what it wanted in great detail, the contractor built it, and the government paid the contractor's documented costs plus an agreed profit margin. This model has a logic for genuinely novel, one-of-a-kind systems where no one can estimate costs in advance. But it contains a corrosive incentive. The contractor's profit grows with its costs, so there is no reward for efficiency and a quiet penalty for delivering anything cheaply. Schedules slip and budgets swell, and the structure absorbs it all [4].

Starting in 2006 with the Commercial Orbital Transportation Services programme, NASA tried something different. Rather than a cost-plus contract, it used a Space Act Agreement structured around fixed milestone payments. SpaceX received an initial award of about $278 million, later rising to roughly $396 million, in exchange for developing and demonstrating a cargo transport capability across a set of defined milestones [4]. The agency paid for results, not for effort. If development cost more than expected, that was the company's problem, not the taxpayer's.

The contrast in numbers is stark. NASA ran the Falcon 9 development through its traditional cost model, the kind used for cost-plus procurement, and estimated that developing the vehicle that way would have cost around $3.6 billion [4]. The actual development cost, as reconstructed by NASA analysts, was closer to $300–400 million. The fixed-price commercial structure did not simply save a margin. It changed the order of magnitude.

This new approach carried through into crew transport. The Commercial Crew Programme, which restored the ability to fly astronauts from American soil after the Shuttle's retirement, was also built on firm fixed-price contracts. Because most of the cost risk sat with the commercial partner, delays cost the companies rather than ballooning the government's bill [4]. Estimates put the effective cost to NASA of a seat on Crew Dragon at on the order of tens of millions of dollars, markedly below the roughly $90 million per seat the United States had been paying for access to the International Space Station on Russian Soyuz vehicles, and far below the true per-seat cost of the Shuttle [4].

There are several reasons this procurement shift mattered so much.

It transferred risk to the party best able to manage it. A company that controls its own engineering can manage development risk far better than a government office can manage a contractor at arm's length. Fixed prices made that company bear the consequences of its own choices, which sharpened every decision.

It allowed private capital to enter. Because SpaceX kept the intellectual property and could sell the same rockets to other customers, the NASA contracts acted as anchor revenue rather than as the whole business. NASA's own analyses found that its commercial cargo programme leveraged additional private dollars for every public dollar spent, a multiplier that simply does not exist under traditional procurement [4].

It amortised cost across many customers. The same Falcon 9 that flew NASA cargo could fly commercial satellites, national security payloads, and Starlink. Spreading fixed costs across a high flight rate is the mechanism that drives unit costs down, and only a commercial model that allowed non-government sales made it possible [4].

The lesson is not that government got out of the way. It is that government changed what it was buying, from cost-reimbursed effort to fixed-price outcomes, and in doing so it created the conditions for a commercial launch industry to exist at all. The old cost-plus world has not vanished. Large legacy programmes still run that way, with per-launch costs in the billions. But the centre of gravity moved, and SpaceX moved it.

5 How Launch Got Cheap

Lowering the cost of reaching orbit is often described as a single breakthrough. In reality it is the product of several reinforcing changes, each of which matters, and none of which works fully without the others. This section walks through the cost stack from the bottom up.

Reuse the most expensive part. In an expendable rocket, the first stage, with its engines, structure, and avionics, represents the largest share of the manufacturing cost and is destroyed on every flight. Recovering and re-flying that stage is the headline change. NASA's cost analysis, using SpaceX list prices, indicates that a Falcon 9 can deliver payload to low Earth orbit at around $2,720 per kilogram, versus roughly $54,500 per kilogram for the Shuttle, largely because the booster, the most expensive part, is reused rather than discarded [4]. Spreading one manufacturing cost across many flights is the same trick that makes air travel affordable, applied to rockets for the first time at scale.

Fly often enough to matter. Reuse only pays off if you fly. A reusable booster that sits in a hangar saves nothing. In recent years SpaceX has flown at a cadence no other operator approaches, with dozens of Falcon 9 and Falcon Heavy launches per year supporting commercial customers, government payloads and Starlink [4, 7]. High flight rate spreads the fixed costs of factories, pads, and personnel across many missions, and it shortens the turnaround time on each booster, which in turn raises the value of reusing it. Cadence and reuse are not two separate advantages. They are one advantage feeding itself.

Build it yourself. Vertical integration removes layers of subcontractor margin and, more importantly, gives the company direct control over design and iteration. Making its own Merlin and Raptor engines, its own avionics, and its own satellites lets SpaceX redesign for manufacturability and drive costs down continuously rather than negotiating each improvement with a supplier whose profit depends on the status quo [4, 5].

Design for production, not just for performance. Traditional rockets were engineered as exquisite, low-volume objects. SpaceX increasingly engineers for the factory, choosing designs that are cheaper to build in quantity even when they are not the absolute optimum on paper. This is the logic that turned satellite manufacturing into something closer to a production line, with the company building thousands of Starlink satellites a year [3, 6].

Accept iteration and failure as cost-saving. Spending years perfecting a design on paper is expensive in both time and money. Building, flying, and learning quickly, even when that means public failures, compresses development timelines and surfaces problems that no amount of analysis would have found. The savings show up as development programmes that cost a fraction of their traditional equivalents, as the Falcon 9 cost comparison in the previous section showed [4].

Put these together and the numbers move dramatically. Falcon 9 delivers payload to low Earth orbit at roughly $2,720 per kilogram, against the Shuttle's roughly $54,500 per kilogram, a reduction of about twenty times [4]. Falcon Heavy pushes the theoretical figure lower still if it flies close to its advertised maximum payload, because the same core engineering is spread across a larger mass to orbit [4]. A reusable Falcon 9 mission, at published list prices, is widely understood in industry analysis to leave substantial margin that helps fund everything else [4].

It is worth being precise about what has actually been demonstrated and what is still aspirational. The Falcon 9 figures are real and proven over hundreds of flights. The Starship figures are not. Company targets for Starship envisage costs well below $100 per kilogram and per-flight costs in the low millions of dollars, which would be another order-of-magnitude reduction [4]. As of mid-2026, Starship has flown a series of test flights, has reached space, and has successfully recovered its booster on several attempts, but it has not yet flown a real payload to orbit or demonstrated the full, rapid reusability that those cost targets assume.[3] The distinction between the proven Falcon 9 economics and the promised Starship economics is central to the valuation debate in Section 8.

6 The Real Value of Reusability

Reusability is easy to admire as engineering. Its real significance is economic, and it runs deeper than the obvious saving on hardware. This section unpacks why reuse changes not just the price of a launch but the structure of the entire market.

Start with the direct saving, which is the part everyone sees. If the first stage is the most expensive component and you fly it ten or twenty times instead of once, you spread that cost across many missions. But if direct hardware saving were the whole story, reuse would lower prices by some percentage and the industry would look much as it did before, only cheaper. The more important effects are indirect [4].

Reuse unlocks cadence, and cadence unlocks everything. When you no longer have to build a new rocket for every flight, the constraint on how often you can launch shifts from your factory's output to your ability to refurbish and refly. A reusable fleet can sustain a flight rate that an expendable model could never reach without an impossibly large factory [4, 7]. High cadence is what allows a company to deploy a constellation of thousands of satellites, to offer customers reliable and frequent launch slots, and to spread fixed costs thin. The constellation business in the next sections simply cannot exist without it.

Reuse changes what satellites can be. When launch is scarce and expensive, every satellite must be gold-plated to last fifteen or twenty years, because replacing it is a major event. When launch is cheap and frequent, satellites can be smaller, simpler, cheaper, and shorter-lived, because replacing them is routine. This inverts decades of design philosophy. Starlink satellites are built to operate for a limited number of years and then deorbit and burn up, to be replaced by newer, better units [6]. That is only rational when getting the replacement to orbit is inexpensive. Cheap launch turns satellites from monuments into consumables.

Reuse converts a capital event into an operating rhythm. Under the old model, deploying a network meant a series of enormous, infrequent, high-risk launches, each a bet-the-mission event. Under reuse, deployment becomes a steady cadence of routine flights, more like running a delivery service than staging an expedition. This smooths cash flow, reduces the consequence of any single failure, and makes the whole enterprise financeable in a way it never was before [4].

Reuse compounds with vertical integration. Because SpaceX builds its own rockets and its own satellites and flies them itself, the savings from reuse flow straight into the cost of deploying its

[3] Status as of June 2026 based on SpaceX webcasts and major media coverage.

own network rather than being captured as margin by a launch supplier. The company is, in effect, getting wholesale launch for its constellation. A competitor that has to buy launches on the open market deploys the same kind of network at a structurally higher cost. This is the quiet advantage that makes Starlink's economics so difficult for rivals to match [3, 6].

There is a second-stage question worth flagging honestly. Falcon 9 reuses its first stage but expends its upper stage. Full reusability, recovering and reflying both stages, is the prize that Starship is designed to claim, and it is genuinely hard. The upper stage has to survive reentry from orbital velocity, which is far more punishing than the suborbital return of a booster [4]. If Starship achieves rapid full reusability, the cost of orbit could fall by another order of magnitude and the economics of everything downstream would shift again. If it does not, the proven Falcon 9 model still represents a roughly twenty-fold improvement over what came before, which is itself transformative [4].

The honest summary is that reusability has already changed the industry decisively at the level Falcon 9 has demonstrated, and stands to change it again, more dramatically, if Starship delivers. The valuation, as we will see, leans heavily on the second of those.

7 The Field of Competitors

By 2026, SpaceX does not so much lead the launch market as define it. Industry estimates indicate that the company handles a majority of global commercial launches measured by mass, and accounts for the large majority of United States orbital launches [4, 7]. Total mass delivered to orbit worldwide more than quintupled between 2020 and the mid-2020s, largely driven by Starlink and other constellation deployments [4, 7]. The company is both the largest launch provider and its own largest customer, a position no competitor occupies.

It is worth looking at the field honestly, because the valuation question partly depends on how contestable that dominance is.

United Launch Alliance, the Boeing and Lockheed Martin joint venture that long held the United States government launch franchise, has moved from its legacy Atlas and Delta vehicles to Vulcan Centaur, which has been certified for national security launches [8]. ULA is reliable and trusted for high-value government payloads, but its model is not built around the high-cadence, low-cost, reusable approach that defines SpaceX. It competes on assurance and sovereignty more than on price.

Blue Origin, founded by Jeff Bezos in 2000, brought its heavy-lift New Glenn rocket to initial flights in the mid-2020s. New Glenn is designed for partial reusability and Blue Origin also supplies the BE-4 engines that power ULA's Vulcan [9]. The company is deliberate, well-funded by Bezos, and building for the long term, but it remains years behind SpaceX on flight heritage and cadence, and it has its own constellation ambitions to feed.

Rocket Lab carved out the small-satellite niche with its Electron rocket, the most frequently flown dedicated small launcher in the Western world, and is moving up-market with the medium-lift, partially reusable Neutron, targeting commercial service around the mid-2020s [10]. Rocket Lab is also a significant supplier of spacecraft components, which gives it a second revenue stream.

Arianespace and ArianeGroup, Europe's launch champions, brought Ariane 6 into service in 2024 to restore the continent's independent access to space [11]. Their value proposition is sovereignty for European institutional payloads more than head-to-head price competition.

China's state and commercial launchers, including the operators behind the planned Guowang and Qianfan broadband constellations, represent some of the most strategically significant long-term competition. China is deploying its own mega-constellations and developing reusable vehicles, and it operates in a largely separate market governed by export controls and national policy rather than open commercial competition [12]. For investors, China is less a competitor for SpaceX's existing customers than a parallel system racing toward the same capabilities.

The pattern across all of these is consistent. Each competitor is strong in a specific niche - government assurance, heavy-lift, small-lift, regional sovereignty, or is backed by enough capital and patience to matter eventually. None currently matches SpaceX on the combination that drives its economics: reusability, cadence, vertical integration, and an in-house constellation that guarantees demand [4, 3]. The medium-lift workhorse market, the heart of the commercial business, has been close to a single-vehicle monopoly.

For the valuation, two things follow. First, SpaceX's launch dominance is real and durable in the near term, because catching up requires not one breakthrough but a whole reinforcing system, and that takes years. Second, dominance in launch is not what justifies $1.8 trillion. The launch market, even fully captured, is too small. The valuation rests on what SpaceX builds on top of its launch advantage, which brings us to the central question.

8 What Has to Be True

This is the section the whole paper has been building toward. A valuation near $1.8 trillion against about $18.7 billion of 2025 revenue implies a price-to-sales multiple in the mid–90s [3]. The company also reported a consolidated net loss of roughly $4.9 billion in 2025, compared with a profit the year before, as it absorbed heavy investment in launch, Starlink and AI infrastructure [3]. No conventional aerospace or telecom valuation framework gets close to justifying that number on current financials. To make sense of it, you have to identify the specific assumptions the market is pricing in, and then ask how likely each is to hold.

Here are the load-bearing assumptions, roughly in order of importance.

Assumption one: Starlink becomes a globally dominant broadband utility. This is the foundation, and it is the part with the most evidence behind it. SpaceX's IPO filing groups Starlink within a "Connectivity" segment that generated about $11 billion of revenue in 2025, growing roughly 50% year on year, and accounting for around 60% of total company revenue [3]. The same filing indicates that SpaceX reported about 10.3 million paid Starlink subscriptions in the first quarter of 2026, more than double the level a year earlier [3]. Contemporary commentary and company materials suggest Starlink serves customers in over 150 countries and markets [3, 6].

Average monthly revenue per user in connectivity fell from roughly $86 to about $66 over the twelve months to early 2026, as the company expanded into lower-income markets and introduced cheaper plans [3]. Dominance is therefore being bought partly with lower prices per user, and the model only works if total volume and high-margin enterprise, maritime, aviation and government revenue more than compensate. The bull case is that as the constellation densifies, the marginal cost of serving each new subscriber falls while higher-value customer segments expand, producing a structurally dominant broadband business with strong margins [3].

For this to support the valuation, Starlink has to keep compounding subscribers and extend deeply into the broader global telecommunications market, which is measured in the trillions of dollars in annual revenue.[4] That is a large, but not inherently impossible, leap.

Assumption two: Starship works, and works cheaply. The valuation implicitly assumes that Starship achieves rapid, full reusability and drives launch costs toward the low figures the company has outlined in public remarks and regulatory filings. This matters in two ways. It would sharply reduce the cost of deploying and replenishing Starlink, which is the largest single

[4]Telecom market size based on standard industry estimates from global operator and analyst reports.

capital commitment on the SpaceX balance sheet, and it would open new markets that are uneconomic at current Falcon 9 and Falcon Heavy price points [4].

As of mid-2026, Starship has conducted multiple integrated flight tests, reached space, and successfully recovered its Super Heavy booster on several occasions, but it has not yet flown a paying payload to orbit, demonstrated orbital propellant transfer, or proven the kind of rapid, aircraft-like turnaround that would underpin the most aggressive cost-per-kilogram targets.[5] Orbital refuelling, which NASA's lunar architecture depends on in the form of many tanker flights in quick succession, remains an unproven and technically demanding milestone. This is the single largest gap between what has been demonstrated and what the valuation assumes.

Assumption three: the AI and orbital computing story is real. Following the acquisition of xAI in early 2026, SpaceX has increasingly framed its long-term narrative around artificial intelligence and orbital data centres, describing a total addressable market in the tens of trillions of dollars for space-based computing and connectivity [3]. The argument is that cheap launch plus a vast satellite network plus AI compute could create a new layer of computing infrastructure in orbit, with Starlink as the distribution network. This is the most speculative assumption and also, arguably, the one doing the most work in pushing the valuation from the high hundreds of billions toward the low trillions. It rests on technologies and markets that are only nascent today.

Assumption four: the competitive moat holds for a decade or more. The valuation assumes that no rival meaningfully erodes SpaceX's launch dominance or matches Starlink's cost structure over the long period during which all this value is supposed to materialise. As the previous section noted, United Launch Alliance, Blue Origin, Rocket Lab, Arianespace and Chinese state and commercial operators all have credible programmes, but none currently matches the combination of high-cadence reusability, vertical integration, and an in-house constellation that drives SpaceX's economics [4, 3]. That is a reasonable assumption in the near term; it is much less certain over ten years, as competitors' reusable vehicles and constellations mature [8, 9, 10, 11, 12].

Set against these assumptions is the bear case, which deserves a fair hearing. Independent equity research providers that have reviewed SpaceX's S-1 have published discounted cash-flow analyses implying fair values substantially below the IPO pricing, even while acknowledging that the stock could trade above those levels for some time because of investor enthusiasm for growth and AI exposure [13]. The multiple is extreme even on forward estimates: assuming revenue in the mid-$20 billion range in 2026, the forward price-to-sales ratio would still sit at many tens of times sales, far above almost all other listed aerospace, telecoms, or infrastructure businesses [3]. The company is loss-making at the consolidated level. And several of the load-bearing assumptions - full Starship reusability and the orbital-AI thesis in particular, remain unproven.

The way to read the $1.8 trillion figure, then, is not as a valuation of what SpaceX is, but as a bet on what it might become. Stripped down, the market is paying roughly a high-hundreds-of-billions valuation for the demonstrated business - launch dominance plus a rapidly growing, profitable connectivity segment, and adding a very large premium on top for the options: Starship economics, the global telecom opportunity, and the orbital-AI narrative [3, 4]. Whether that is visionary or excessive depends entirely on how many of those options pay off, and over what timeframe.

[5] Flight-test status as of June 2026 based on public SpaceX webcasts and launch coverage in major news outlets.

9 Conclusion

The achievement at the centre of this story is concrete and proven. SpaceX took an industry that, in NASA's own comparisons, delivered payload to orbit at roughly $54,500 per kilogram on the Space Shuttle and brought the advertised figure down to around $2,720 per kilogram on Falcon 9, a roughly twenty-fold reduction by attacking the three pillars of the old toll-road model at once: it stopped throwing rockets away, it flew often enough for reuse to pay, and it benefited from a procurement shift at NASA that rewarded outcomes instead of effort. That is a genuine, durable, world-changing accomplishment, and it is not in dispute.

The valuation is a different and more contestable thing. At roughly $1.8 trillion, the market is not pricing the launch company. The launch market, even completely dominated, is too small to support that number. It is pricing the cascade of consequences that cheap, reliable, high-cadence launch sets in motion. First among these is Starlink, which has already grown into the company's profit engine and a genuine global broadband utility, and which is the most defensible part of the valuation [3]. Beyond Starlink lie the bets: that Starship will achieve the dramatic further cost reductions it promises, that the global telecommunications market is there to be taken, and that an orbital computing and AI layer will become a real and enormous business. The first of these has strong evidence behind it. The others range from plausible to genuinely speculative.

So the honest answer to the $1.8 trillion question is conditional. The valuation can be justified if a demanding sequence of things goes right: if Starlink continues to compound and expand its margins, if Starship crosses the gap from promising test vehicle to cheap, fully reusable workhorse, if the orbital-AI thesis turns into revenue, and if no competitor closes the gap for a long time. Remove any one of those and the number looks stretched. Remove two and it looks indefensible. Independent valuation work putting fair value materially below the IPO level is a reminder that serious analysts, looking at the same disclosures, see only a portion of the headline figure as the defensible core, with the rest as a bet on the future.

There is a deeper point worth ending on. Every previous step in this story followed the same pattern. Reusability looked like a stunt until it reset the cost of orbit. A government-backed constellation looked like manufacturing demand out of thin air until it became the company's profit centre. SpaceX has a track record of making the implausible routine. That history is precisely why the market is willing to pay for options that have not yet paid off, and precisely why a rational, evidence-driven analyst should remain cautious about pricing them as if they already have.

The company changed the economics of getting to space. Whether it has correctly priced the economics of being in space is the question its new public shareholders have just agreed to answer with their money. The engineering case is closed. The financial case is still in flight.

References

[1] Bloomberg News. "SpaceX IPO Raises $75 Billion in Biggest Debut of All Time". In: (2026). URL: https://www.bloomberg.com/news/articles/2026-06-11/elon-musk-s-spacex-set-to-make-history-with-record-breaking-ipo.

[2] Yahoo Finance. "Musk Becomes World's First Trillionaire as SpaceX Shares Soar". In: (2026). URL: https://uk.finance.yahoo.com/news/spacex-ipo-set-liftoff-record-024830673.html.

[3] PCMag. "SpaceX IPO Filing Offers First Glimpse at Starlink Subscriber Numbers and Financials". In: (2026). URL: https://www.pcmag.com/news/spacex-ipo-filing-offers-first-glimpse-at-subscriber-numbers-financials.

[4] Dana G. Andrews. *The Recent Large Reduction in Space Launch Cost*. Tech. rep. NASA Ames Research Center, 2018. URL: https://ntrs.nasa.gov/citations/20200001093.

[5] SpaceX. *SpaceX Company History and Mission*. 2024. URL: https://www.spacex.com/about.

[6] SpaceX. *Starlink Progress Report*. 2025. URL: https://www.starlink.com/progress.

[7] Space Stats. *Orbital launches by year*. 2026. URL: https://spacestatsonline.com/launches/.

[8] United Launch Alliance. *Vulcan Centaur Overview*. 2025. URL: https://www.ulalaunch.com/rockets/vulcan-centaur.

[9] Blue Origin. *New Glenn Heavy-Lift Launch Vehicle*. 2025. URL: https://www.blueorigin.com/new-glenn/.

[10] Rocket Lab. *Neutron Rocket*. 2025. URL: https://www.rocketlabusa.com/rockets/neutron/.

[11] ArianeGroup. *Ariane 6: Europe's New Heavy-Lift Launcher*. 2024. URL: https://www.ariane.group/en/launcher/ariane-6/.

[12] *Chinese Constellation Operators Submit Plans for Almost 200,000 New Satellites*. 2026. URL: https://www.china-in-space.com/p/chinese-constellation-operators-submit.

[13] Andreas Haaker. *Is SpaceX worth $1.75 trillion?* 2026. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6640738.