The Window Wedge: SpaceX, Mega-IPOs, and the Difference Between a Hot Listing and a Market Top

FCRI Research

June 2026

Abstract

On 11 June 2026, public investors were offered SpaceX at a reported $1.77 trillion valuation through a reported $75.0 billion IPO. That is not a small financing round with a rocket logo attached. It is a capital-allocation event large enough to ask a harder question than whether the company matters.

The question is price.

This paper studies that question with a return-observable prototype panel of 39 large IPO and direct-listing events from 2006 through 2026. The central object is the *Window Wedge*: the extra price a seller can obtain when an unusually receptive issuance window meets an asset investors badly want to own. The wedge itself is not structurally identified; the measured proxy is realized benchmark-relative aftermarket performance. The result is narrow, but it is not soft. IPOs raising at least $3 billion lagged their local benchmarks by 14.1 percentage points over the next twelve months, while those same benchmarks rose by 8.1 percent on average. SpaceX fits the timing side of the windows-of-opportunity view: it arrived near Nasdaq highs, after an AI-led rally, and at a scale 2.55 times the prior largest deal in the panel. Mega-IPOs are not reliable market-top bells. They are more often warnings about the price of the thing being sold.

1 The SpaceX Trap

On 11 June 2026, the market was not being asked to fund a sketch on a napkin. It was being asked to fund SpaceX.

That distinction matters. SpaceX is a launch company, a satellite company, a communications company, a space-infrastructure company, and a defense-adjacent platform. It is also scarce. For years, public investors who wanted that bundle could mostly stand outside the fence and watch the launches on video. Then the fence opened, at a reported $1.77 trillion valuation.

Here is the trap. The company can be extraordinary and the entry price still poor.

A reported $1.77 trillion valuation asks public buyers to capitalize years of future execution before much of the long-run cash-flow map is visible. It asks them to value launch cadence, satellite economics, communications scale, strategic optionality, regulation, geopolitical importance, and technical dominance at once. It also arrived after a sharp technology rally, with Nasdaq close to recent highs. Large issuers do not usually choose cold rooms for their offerings.

The SpaceX listing therefore forces two questions that investors often blur into one. Is the business real? And did the moment help set the price?

The old proverb says that big IPOs ring the bell. It is a useful proverb because it contains a truth. Sellers like strong markets. Bankers like full order books. Buyers become less demanding when the thing for sale is scarce, glamorous, and hard to replicate. But the proverb packs three claims into one sentence, and only some of them survive contact with the evidence.

- Large IPOs should arrive when equity markets are strong.
- Buyers of newly listed large IPOs should often lag comparable benchmarks after the deal.
- Large IPO dates should be followed by weak forward returns for the whole market.

The first two claims can hold even when the third fails. A new issue can be expensive relative to the market while the market itself keeps rising. The public listing can absorb the enthusiasm without exhausting the bull market.

That is the spine of the paper. The historical IPOs raising at least $3 billion underperformed their local benchmarks by 14.1 percentage points over the next twelve months, while those same benchmarks rose by 8.1 percent on average. The evidence points less to a market-top bell than to a new-issue warning.

The bell that rings may be local.

2 The Window Wedge

Think about a house. The roof may be sound, the street may be quiet, the school district may be excellent, and the seller may still get a higher price when ten buyers arrive with preapproved mortgages than when two buyers arrive in the rain. The house is not imaginary. The bidding conditions are not imaginary either.

A mega-IPO has the same structure. The seller sees a rare chance to sell. The banker sees demand. The buyer sees access to an asset that was previously locked away. When those three incentives meet in a strong market, the price contains more than the asset alone. It contains the window.

The listing price can be written as:

$$P_{i,t} = F_{i,t} + W_t S_i.$$

Here $P_{i,t}$ is the listing price or early trading price of issuer i at time t. $F_{i,t}$ is fundamentals: revenue, cash flow, assets, technology, contracts, regulation, reserves, or market position. W_t is the issuance window: market-wide willingness to fund risky or long-duration equity. Long-duration equity is equity whose expected cash flows sit far in the future. S_i is sensitivity to that window.

One force does the work: time-varying willingness to fund distant growth. The causal chain is:

$$\begin{aligned} \text{strong market} &\rightarrow \text{greater willingness to fund distant growth} \\ &\rightarrow \text{higher listing price for window-sensitive issuers} \\ &\rightarrow \text{weaker relative aftermarket return if the window fades.} \end{aligned}$$

No villain is required. Issuers and selling shareholders prefer high prices. Underwriters prefer deals that can be placed. Public buyers want access to scarce growth assets. When risk appetite rises, those incentives line up cleanly; the issuer supplies equity, the underwriter sells the transaction, and buyers accept a price that embeds more future success.

The mechanism gives four implications.

1. Large offerings should cluster in strong markets because W_t is high.
2. Newly listed stocks should often underperform their benchmarks after listing if the window component fades.

3. Story-heavy issuers should be more fragile than mature issuers because more of their value depends on outcomes that cannot be verified quickly.

4. Broad-market returns need not be weak after the IPO date because the wedge can sit mainly in the new issue.

The fourth implication is the quiet one, and the important one. A mega-IPO can be a warning about the IPO price without being a reliable alarm for the whole market.

The wedge itself is not observed. The measured proxy is benchmark-relative post-close performance:

$$x_{i,h} = r_{i,h}^{stock} - r_{i,h}^{benchmark},$$

where $x_{i,h}$ is the stock's excess return from the first close to horizon h. A negative value means the IPO lagged the benchmark. The realized relative shortfall is:

$$RS_{i,h} = -x_{i,h}.$$

For example, if an IPO falls 5 percent from the first close over twelve months while its local benchmark rises 15 percent, the excess return is:

$$x_{i,12} = -5\% - 15\% = -20\%.$$

The realized relative shortfall is therefore:

$$RS_{i,12} = 20\%.$$

That number does not prove that 20 percent of the offer price was sentiment. It says that the public buyer lost 20 percentage points relative to the available benchmark after the listing.

This is the right frame for SpaceX. Strategic scarcity, launch economics, satellite communications, defense relevance, and long-duration technological optionality all make the issuer important. Those same features also make the valuation sensitive to public-market willingness to capitalize future dominance. The Window Wedge is not a fraud story. It is a price-sensitivity story.

3 The Evidence We Can Actually See

The dated event is a large IPO or direct listing. The return object is public-market performance after trading begins. That sounds simple; the hard part is knowing what is missing.

The investigation built a return-observable prototype panel of 39 large IPO and direct-listing events from 2006 through 2026. The panel combines daily adjusted closing prices for issuers and local benchmarks with manually curated proceeds, offer prices or reference prices, listing dates, and issuer classifications. The main historical aftermarket sample is restricted to pre-2025 events with proceeds of at least $3 billion and observable return horizons.

The word prototype earns its keep. This is not a completed global IPO census. The candidate-universe audit requested by the strongest referee critique has not yet been built, so the number of eligible global deals missing from the observable panel is not known. That prevents population-level language such as "all global mega-IPOs." The claim here is narrower and cleaner: it is about the deals we can observe in this panel.

The main threshold is nominal proceeds of at least $3 billion. It was chosen because it isolates very large public offerings while preserving enough observations to study post-listing returns. A $10

Table 1: Empirical scope and uncompleted referee tests

Question	Measured in this paper	Not completed in this paper
Candidate universe	Return-observable prototype panel of 39 large listings	Exhaustive global candidate table with missing-deal reasons
Issuance wave	SpaceX scale and current market state	Global IPO denominator and rolling issuance intensity
Issue aftermarket returns	Stock buy-and-hold returns versus local benchmarks	Multi-benchmark, factor, sector, and currency adjustments
Broad-market timing	Forward benchmark returns after IPO dates	Matched hot-date non-IPO placebo test
Issuer type mechanism	Story-versus-mature split	Row-level accounting audit and classification sensitivity table
Transaction type	IPOs and direct listings identified	Full event-type split for IPOs, privatizations, ADRs, direct listings, and secondary-heavy sales

billion threshold is also reported where the data support it. The threshold is not inflation-adjusted in the current results.

Local benchmarks are used because a Saudi listing, a Chinese platform ADR, a European privatization, and a U.S. technology listing do not share one natural index. A local benchmark is a same-market counterfactual. It does not remove sector, style, factor, currency, or size exposures. It is a useful yardstick, not a perfect twin.

The first-day return is measured from the offer or reference price to the first traded close:

$$r_i^{FD} = \frac{P_{i,0}^{close}}{P_i^{offer}} - 1.$$

For direct listings, this is a reference-price-to-close return, not IPO underpricing. A direct-listing reference price is not the same object as an underwritten offer price. Direct listings are therefore included in the broader public-listing panel but should not be pooled mechanically with conventional IPOs when studying offer-price underpricing.

The SpaceX arithmetic illustrates the calculation. If the offer or reference price is $135 and the first available close is 19.2 percent higher, the first available close is:

$$135.00 \times (1 + 0.192) = 160.92.$$

That example is only arithmetic. It is not evidence on SpaceX's one-month, six-month, or twelve-month aftermarket performance.

Table 1 states the boundary of the evidence.

Table 1 is the guardrail. The evidence can show that SpaceX fits a hot-window context and that historical large IPOs often underperformed. It cannot support a complete claim about all global mega-IPOs. It also cannot decompose SpaceX's reported valuation into a measured fundamental component and a measured sentiment component.

The regulatory boundary matters because the move from private to public markets changes disclosure, trading, and investor-protection obligations. Greene et al. emphasize the need for a comprehensive approach to capital markets regulation, which is relevant when very large private firms enter public markets and become tradable by a broader investor base [1].

4 SpaceX Enters a Hot Window

The first empirical question is whether SpaceX arrived in a hot tape.

Table 2: Mid-2026 market state

Market proxy	QTD return (%)	YTD return (%)	One-year return (%)	Interpretation
Nasdaq	19.9	11.4	31.7	Near recent highs
QQQ	25.0	17.6	35.8	Strong large-cap technology
QQEW	21.1	8.2	15.2	Equal-weight technology positive, less extreme
RSP	10.3	11.0	19.9	Broad market positive

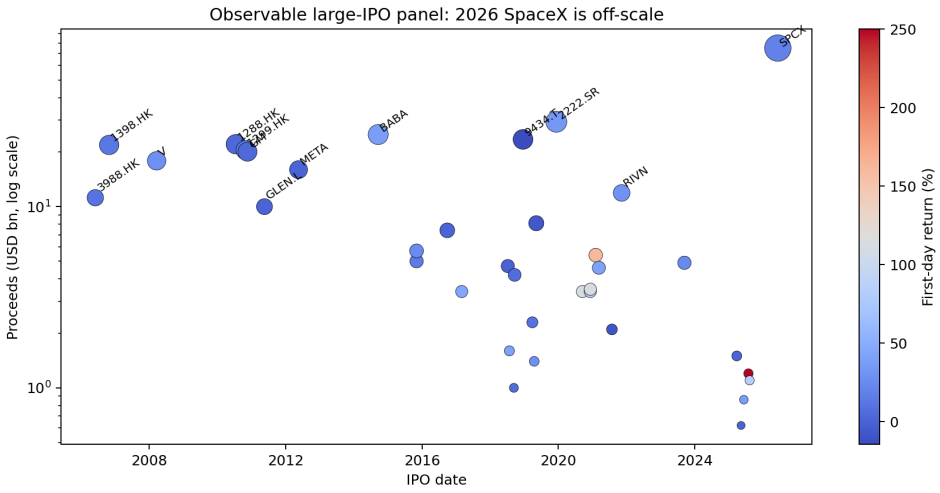

Figure 1: Large IPO and direct-listing timeline. The horizontal axis is listing date and the vertical axis is proceeds in billions of U.S. dollars.

It did.

As of 12 June 2026, Nasdaq was only 4.45 percent below its 252-trading-day high and in the 92.5th percentile of its trailing 252-day range. Against its trailing three-year range, Nasdaq was also 4.45 percent below its high and in the 97.5th percentile. Nasdaq returns were 19.9 percent quarter-to-date, 11.4 percent year-to-date, and 31.7 percent over one year.

Those are not abstract market conditions. They are the kind of conditions in which a long-duration growth asset can be sold with a straight face at a very large number.

Table 2 compares capitalization-weighted and equal-weight proxies. Equal-weight proxies are included because they show whether the rally was broad or concentrated in the largest names.

The market was strong, but not uniformly euphoric. QQQ rose 35.8 percent over one year, while the equal-weight Nasdaq proxy rose 15.2 percent. The rally was real; it was also concentrated in the large technology and growth complex most likely to reward distant optionality. That is exactly where a high-S_i issuer would want the window open.

For SpaceX, this matters because the company is not being valued like a mature utility with all its cash flows visible. Its most attractive possibilities are large, strategic, and far away. Warm risk appetite is worth more to that kind of seller.

Figure 1 places SpaceX against the other large offerings in the observable panel. The vertical axis is proceeds in billions of U.S. dollars, so the main visual fact is scale.

SpaceX is off-scale in the observable panel. Its reported $75.0 billion raise ranks first and is 2.55 times Saudi Aramco's $29.4 billion offering, the largest prior panel deal. At that size, even a modest window premium would be measured in tens of billions of dollars. Scale turns a valuation

Table 3: SpaceX scale in the observable panel

Measure	Value	Interpretation
Reported SpaceX proceeds ($bn)	75.0	Largest deal in the panel
Largest prior panel deal, Saudi Aramco ($bn)	29.4	Prior scale benchmark
SpaceX proceeds divided by prior largest deal	2.55	SpaceX is off-scale historically
Reported SpaceX valuation ($tn)	1.77	Contextual input, not computed from returns

Table 4: Large-IPO cohorts in the observable panel

Cohort	Deals (N)	Proceeds ($bn)	Observable aftermarket pattern
2006 to 2008 pre-GFC financials	3	51.0	Mean 12m excess return of +13.0%
2010 to 2012 privatizations and platforms	5	88.7	Mean 12m excess return of -17.0%
2014 to 2016 China and platform listings	4	43.1	Mean 12m excess return of -21.3%; all underperformed
2017 to 2019 apps, China, and Aramco	11	79.6	Mean 12m excess return of -28.1%; 81.8% underperformed
2020 to 2021 pandemic growth listings	9	34.3	Mean 12m excess return of -20.6%; 88.9% underperformed
2022 to 2024 reopen context	1	–	Arm only; mean 12m excess return of +91.3%
2025 to 2026 SpaceX-led context	6	80.28	Full 12m returns mostly not yet observable

question into a capital-allocation question.

Table 3 states the scale directly.

The episode should therefore be described with care. The panel captures a concentrated SpaceX-led reopening dominated by one issuer, not a completed global issuance census. SpaceX is not merely large inside the sample. It bends the sample around itself.

The historical analogues are shown in Table 4. These cohorts matter because they separate different environments: pre-crisis financial listings, post-crisis privatizations and platforms, later China and platform clusters, app and strategic listings, pandemic-growth issues, and the present SpaceX-led reopening.

The table gives both confirming examples and counterexamples. The 2006 to 2008 pre-GFC financial cohort had a positive mean twelve-month excess return of 13.0 percent. Arm also delivered a large positive twelve-month excess return in the sparse 2022 to 2024 context. These cases are useful precisely because they are inconvenient. They stop the lazy rule from becoming a superstition.

The weaker cohorts had a different flavor: privatization and platform deals after the crisis, China and platform listings, app and strategic listings, and pandemic-growth issues. These were not necessarily bad businesses. Many were plausible, even compelling. The problem was that plausibility had already been capitalized.

Figure 2 shows the same pattern by cohort. The vertical axis reports mean twelve-month excess return, so negative bars mean the cohort lagged its local benchmarks.

The cohort pattern does not say that strong markets are bad markets. It says something more specific: in several platform, app, China, and pandemic-growth cohorts, large new listings failed to keep up with the benchmarks available at the same time.

SpaceX fits the timing side of the windows-of-opportunity view. It arrived near market highs, in a rally led by large-cap technology and growth, and at a scale far beyond prior observed mega-IPOs. The window was open.

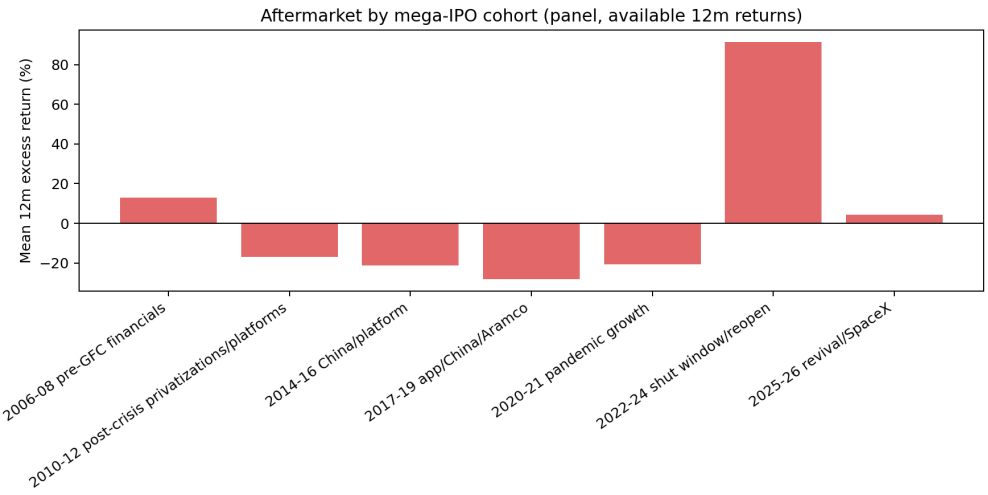

Figure 2: Cohort twelve-month excess returns. The horizontal axis groups IPOs by historical cohort and the vertical axis reports mean twelve-month stock return relative to the local benchmark.

5 The Listing Is the Weak Link

The direct investor question is not grand. After the first close, did the IPO beat the market?

Table 5 reports the answer for historical large IPOs in the observable panel.

The first public trading day looked generous. Historical IPOs raising at least $3 billion had a mean first-day return of 31.1 percent and a median first-day return of 17.1 percent. Then the bill arrived. The mean six-month excess return was -15.1 percentage points, and the median was -14.8 points. The mean twelve-month excess return was -14.1 points, and the median was -20.5 points. In counts, 18 of 25 observable historical IPOs in this size group underperformed their local benchmark over twelve months, or 72.0 percent.

Put less statistically: most buyers would have been better off owning the local benchmark for the next year.

At thirty-six months, the mean excess return was -5.4 percentage points and the median was -19.2 points. The typical large listing therefore struggled to keep up with the benchmark after the initial public trading day. The first-day pop was not a free lunch. It was, often enough, dessert served before the meal.

The uncertainty is large because the sample is small. A deal-level bootstrap repeatedly resamples the observed deals and recomputes the mean. It is used here because the sample is small and the return distribution is uneven. Table 6 reports the resampling evidence.

The bootstrap keeps the inference in proportion. The six-month relative weakness is statistically cleaner. The twelve-month relative weakness is the main economic magnitude. Both point in the same direction, but neither turns this small event panel into a law of markets.

The second row of Table 5 matters for SpaceX. Among historical IPOs raising at least $10 billion, the mean twelve-month excess return was -9.6 percentage points, the median was -11.5 points, and 58.3 percent underperformed. The very largest offerings were weak on average, but less uniformly bad than the broader $3 billion group. Size alone is not the full mechanism.

Size plus window sensitivity is the fragile mix. SpaceX is large enough to clear the size test by a distance; the live question is how much of its reported $1.77 trillion valuation depends on outcomes public investors cannot yet verify.

Table 5: Historical large-IPO aftermarket performance

Observable historical sample	Mean first-day return (%)	Mean 6m excess (pp)	Mean 12m excess (pp)	12m underperformance
Proceeds at least $3bn	31.1	-15.1	-14.1	18 of 25
Proceeds at least $10bn	–	–	-9.6	58.3%

Table 6: Deal-level bootstrap inference for historical IPOs of at least $3bn

Outcome	Mean excess return (pp)	95% interval (pp)	Interpretation
Six-month excess return	-15.1	[-26.2, -2.6]	Interval wholly below zero
Twelve-month excess return	-14.1	[-27.1, +0.1]	Interval barely reaches zero

6 Story Prices and Mature Prices

Meta is a useful humiliation for anyone who likes easy IPO morals. The listing became a cautionary tale quickly, then the business later became much more than the early trade suggested. Arm gives the opposite warning: a large offering can beat the pattern. These examples matter because the question is not whether a famous IPO must disappoint. It is which kind of price has the least room for ordinary disappointment.

The Window Wedge gives one answer. Public buyers value a new issue using two kinds of information. One kind can be checked now: current revenue, cash flow, assets, regulation, reserves, contracts, or proven market share. The other kind must be imagined before it can be verified: future market size, future margins, future network effects, future technical dominance, or future strategic indispensability. In hot windows, the second kind becomes easier to capitalize at a high price.

That creates an asymmetric problem. Optimistic narratives can be priced immediately. Their verification takes years. The first traded price can therefore move faster than the fundamentals that must eventually justify it.

This is not an argument that narrative firms are bad firms. Some narrative firms become exceptional businesses. The issue is timing. When the public-market window is unusually receptive, a firm can be both fundamentally important and priced with little margin for near-term disappointment.

The classification used in the panel is best described as an auditable judgment rule, not a fully automated accounting classifier. A listing is coded as more story-sensitive when listing-date information points to venture-style growth, negative or thin current earnings, platform or network economics, a long-duration technology option, or a valuation that depends heavily on future market creation. A listing is coded as more mature when listing-date information points to a long revenue history, observable profitability or cash-flow base, state ownership or privatization, regulated infrastructure, utility-like assets, or established financial-market infrastructure.

Table 7 splits historical IPOs of at least $3 billion into story and mature groups.

This is the sharpest cross-sectional pattern in the panel. Story IPOs had a mean first-day return of 50.6 percent, compared with 13.2 percent for mature IPOs. Their mean twelve-month excess return was -23.9 percentage points, compared with -5.1 points for mature IPOs. They also underperformed more often over the next twelve months: 83.3 percent of story IPOs lagged their benchmark, compared with 61.5 percent of mature IPOs.

The median split is even starker. Historical story IPOs had a median twelve-month excess return of -39.4 percentage points and a median thirty-six-month excess return of -60.9 points. Historical mature IPOs had a median twelve-month excess return of -12.2 points and a median thirty-six-month excess return of +0.9 points. The largest losses appear where the valuation depends most

Table 7: Story versus mature large IPOs

Group	Deals (N)	Mean first-day return (%)	Mean 12m excess (pp)	12m underperformance (%)
Story IPOs	12	50.6	-23.9	83.3
Mature IPOs	13	13.2	-5.1	61.5

Table 8: Realized relative shortfall by issuer type

Group	Deals (N)	Mean 12m excess (pp)	Realized relative shortfall (pp)	12m underperformance (%)
All historical IPOs at least $3bn	25	-14.1	14.1	72.0
Story IPOs at least $3bn	12	-23.9	23.9	83.3
Mature IPOs at least $3bn	13	-5.1	5.1	61.5
IPOs at least $10bn	–	-9.6	9.6	58.3

heavily on future realization. That is what the Window Wedge predicts.

A Welch test for equality of story and mature twelve-month excess returns has a p-value of 0.204. The split is therefore economically large but not statistically decisive in this small sample. The classification should be read as a mechanism check, not as a settled taxonomy.

For SpaceX, this is the relevant aisle of the evidence. The company has real assets and strategic importance, but much of the reported valuation necessarily leans on future scale, future market structure, and future dominance. Those are story-sensitive inputs even when the story is plausible.

Table 8 translates the excess-return evidence into realized relative shortfall. This is the closest quantified answer in the current design to the sentiment-versus-fundamentals question.

A 23.9 percentage-point realized relative shortfall for story IPOs does not mean that exactly 23.9 percent of their offer prices was sentiment. It means that after public trading began, the story group gave up 23.9 percentage points relative to local benchmarks over twelve months. That is consistent with hot-window pricing risk being larger for story-heavy issuers.

First-day return is a tempting sentiment proxy. It is also noisy. For conventional IPOs, it can reflect investor enthusiasm, offer-price conservatism, underwriter strategy, bookbuilding incentives, litigation risk, and allocation rents. For direct listings, it is not underpricing at all. It is the move from a reference price to the first traded close.

Figure 3 plots first-day return against twelve-month excess return. The horizontal axis is the first-day return from offer or reference price to first close. The vertical axis is twelve-month stock return relative to the local benchmark.

Figure 3 shows why first-day return cannot carry the mechanism by itself. Across historical IPOs and direct listings with observable twelve-month returns and proceeds of at least $1 billion, the Spearman rank correlation between first-day return and twelve-month excess return was 0.045, with a p-value of 0.814. In a small regression of twelve-month excess return on first-day return, log proceeds, a story dummy, and benchmark percentile, the first-day coefficient was positive but insignificant, with a p-value of 0.433. The story dummy was negative but also insignificant, with a p-value of 0.317.

Quartiles tell the same story. The top first-day-pop quartile had a mean first-day return of 79.3 percent and a mean twelve-month excess return of -34.2 percentage points. But the bottom quartile was also weak, at -27.7 points. First-day excitement can be part of the Window Wedge, but it is not a sufficient statistic for it.

The cross-section gives the useful nuance for SpaceX. A mega-IPO is not all sentiment. Visa, Arm, and eventually Meta after its initial drawdown show that fundamentals can overwhelm window-timing concerns. The issue is conditional. The pricing risk appears largest when the

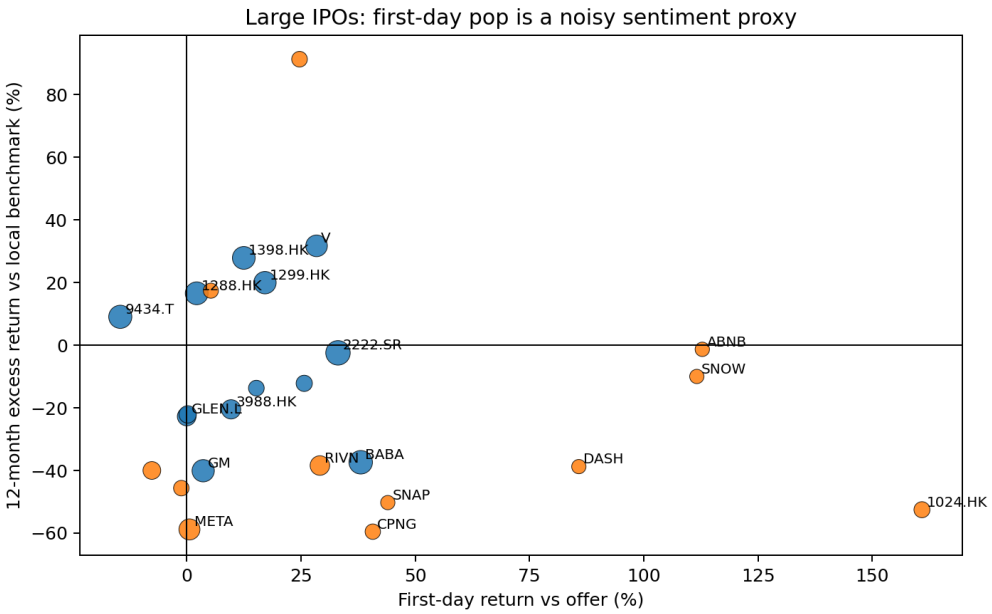

Figure 3: First-day return and twelve-month excess return. The horizontal axis reports first-day return from offer or reference price to first close. The vertical axis reports twelve-month stock return relative to the local benchmark.

deal is both large and narrative-heavy.

7 The Market-Top Temptation

The broad-market version of the story is more dramatic, which is why people like it. It says that mega-IPOs arrive near market highs and are followed by weak market returns. The first half is visible in the panel. The second half is not.

For historical IPOs raising at least $3 billion, the mean local benchmark drawdown from its 252-day high at the IPO date was -6.5 percent. In counts, 56.0 percent of deals occurred within 5 percent of a 252-day high. Issuance windows do tend to open when markets are strong. Sellers are not known for choosing bad weather voluntarily.

The forward-return evidence is weaker. Table 9 reports benchmark returns after the large-IPO dates.

This is the cleanest rejection of the simple market-top version in this sample. Historical large IPOs occurred in strong markets, but those markets did not reliably fall afterward. The mean forward twelve-month benchmark return was +8.1 percent. The median was +6.5 percent. The mean percentile of each benchmark's forward return in its own history was 46.1 percent, which is roughly ordinary rather than crisis-like.

For SpaceX, that distinction is not academic. The listing can be expensive without making Nasdaq a short. The buyer's problem may be the price of SpaceX, not the calendar date of the next bear market.

The hard placebo test has not been completed. The right next test would match each mega-IPO date to non-IPO dates in the same benchmark history with similar trailing twelve-month return, drawdown from high, volatility, and valuation or rate regime where available. If forward returns after mega-IPO dates are materially worse than after matched hot-market non-IPO dates, the market-top claim becomes stronger. If they are similar, the Window Wedge remains issue-level.

The current evidence points to the second interpretation. Mega-IPOs are stronger evidence of a hot issuance window and issue-level valuation risk than of an imminent broad-market top.

Table 9: Forward benchmark returns after historical large IPOs

Historical sample	Mean forward 12m benchmark return (%)	Median forward 12m benchmark return (%)	Interpretation
IPOs with proceeds at least $3bn	8.1	6.5	Positive, not crash-like
Story IPOs with proceeds at least $3bn	10.6	–	Also positive on average
Forward-return percentile in benchmark history	46.1	–	Roughly ordinary

8 SpaceX Through the Window Wedge

SpaceX sits at the extreme corner of the framework.

On the window side, the timing facts are strong. The listing came when Nasdaq was only 4.45 percent below its 252-day high, in the 92.5th percentile of its trailing one-year range, and after a 31.7 percent one-year Nasdaq return. QQQ was even stronger, with a 35.8 percent one-year return. The deal also arrived in a market where large technology and growth assets had led the rally.

On the scale side, SpaceX is not merely the largest deal in the panel. It is off-scale. The reported $75.0 billion raise is 2.55 times the prior largest panel deal. The reported $1.77 trillion valuation is a contextual news input, not computed from the price panel, but it explains why the transaction is economically large enough to test the windows story.

On the fundamentals side, the simple bubble label is too weak. SpaceX has strategic scarcity, infrastructure relevance, national-security optionality, satellite communications exposure, and launch economics. These features can justify a large fundamental component.

On the sentiment side, those same features create long-duration valuation sensitivity. Public buyers are not only buying current earnings. They are buying future dominance in markets whose eventual size, margins, regulation, and geopolitical role are uncertain. That makes SpaceX a high-S_i issuer in the Window Wedge framework.

The actual SpaceX aftermarket test is not yet available. Only one post-listing close was available in the run. The first available close was recorded as +19.2 percent versus the $135 offer or reference price. That is context only. It is not evidence on one-month, six-month, twelve-month, or thirty-six-month performance.

The historical prior is therefore conditional. If SpaceX behaves like the story-heavy large IPOs in the panel, the issue-level risk is meaningful. If it behaves like the strongest mature or strategic exceptions, fundamentals may overwhelm the window. The Window Wedge says where to look. It does not decide the case before returns and fundamentals arrive.

9 What Would Make the Result Stronger

The strongest criticism of this paper is also the most useful one: the distinction is interesting, but the evidence is still closer to an investigation than a settled discovery. I think that is the right standard. Four additional empirical blocks would decide how far the Window Wedge can travel.

First, the candidate universe must be audited. A global table should identify all large listings from 2006 through 2026 above nominal $3 billion, nominal $5 billion, nominal $10 billion, and CPI-adjusted $3-billion-in-2026 thresholds. It should include proceeds, primary and secondary components where available, valuation, exchange, security type, ticker mapping, inclusion flag, and exclusion reason. Without that table, the present sample cannot claim population completeness.

Second, the issuance-wave denominator must be built. A true wave requires annual or quarterly global IPO counts and proceeds. Mega-IPO proceeds should be scaled by total IPO proceeds and by global equity-market capitalization. Without that denominator, SpaceX can be called off-scale, but the whole 2025–2026 period cannot be called a formally measured global mega-IPO wave.

Third, the return tests should be rebuilt by event type and benchmark. Conventional IPOs, state privatizations, ADR or foreign listings, direct listings, and secondary-heavy monetizations are economically different. Direct listings should be excluded from underpricing statistics. Benchmarks should include local broad indexes, country ETFs, global equity proxies, sector or style proxies, and USD-investor returns where currency conversion matters.

Fourth, the broad-market claim needs the matched hot-date placebo. This is the strongest missing identification test. Mega-IPO dates should be compared with non-IPO dates that look similarly hot before the event. That design would distinguish a true market-top signal from the ordinary fact that IPOs happen in strong markets.

These tests are not cosmetic robustness checks. They decide whether the Window Wedge is a broad regularity or a useful description of this observable panel. If the aftermarket underperformance survives universe completion, event-type separation, inflation-adjusted thresholds, multi-benchmark returns, and matched-date placebo tests, the paper becomes a stronger claim about mega-IPO pricing. If it does not, the current result should remain a narrower descriptive fact.

10 Limitations

The main limitation is simple: the panel is observable, not exhaustive. Public price observability is required. Delisted firms, ticker changes, ADR terminations, local-market failures, restructurings, mergers, and hard-to-map securities are likely underrepresented. This selection likely biases long-run performance upward if failed listings are harder to observe, but the size and direction of the bias are not measured.

Three other limits matter. The paper does not construct total global IPO counts or proceeds by year, rolling twelve-month window, or rolling twenty-four-month window, so it cannot prove that 2025–2026 is a global issuance wave in the strict denominator-based sense. The main cutoff is nominal $3 billion, not inflation-adjusted and not scaled by total IPO proceeds, issuer valuation, local equity-market capitalization, or global equity-market capitalization. A $3 billion transaction in 2006 is not equivalent to a $3 billion transaction in 2026.

Benchmark and classification choices also carry weight. Local broad indexes are transparent same-market comparators, but they do not adjust for sector, style, currency, quality, profitability, duration, or factor exposure. Conventional IPOs, privatizations, ADRs, direct listings, and secondary-heavy monetizations do not have the same economics. The story-versus-mature split is economically intuitive but judgmental; a stronger version would publish row-level profitability, revenue growth, age, sector, valuation multiple, primary-versus-secondary proceeds, state ownership, and narrative classification sources, then run sensitivity checks on borderline names.

The final limitation is identification. The paper documents associations. It does not causally identify sentiment, fundamentals, or an exogenous issuance-window shock. The matched hot-date placebo test is the key missing design for the broad-market-return question.

The assumption most likely to be wrong is that the observable panel is representative of the relevant mega-IPO universe. Evidence that missing large IPOs performed much better than included IPOs would weaken the aftermarket conclusion. Evidence that missing failed or delisted IPOs performed worse would strengthen it.

That is the bargain. The result is economically large and worth taking seriously; it is not yet a completed map of the global IPO universe.

11 Reproducibility Note

The calculations use daily adjusted closing prices for issuers and benchmarks, manually curated listing dates, proceeds, offer or reference prices, and issuer classifications. Returns are buy-and-hold returns from the first close to each horizon. Excess returns subtract the matched benchmark return over the same dates. First-day returns use the offer or reference price and the first traded close. Summary statistics, cohort means, bootstrap intervals, rank correlations, and the small regression are deterministic given the input table. No random seed affects the point estimates; resampling intervals are reproducible from the saved event-level data and resampling procedure in the accompanying replication package.

12 Conclusion

The 11 June 2026 SpaceX listing has the classic look of a hot-window event. It arrived near Nasdaq highs, after a sharp technology rally, and at a size that dwarfs prior large IPOs in the observable panel. If one wanted a specimen for the windows-of-opportunity theory, SpaceX would do. It is almost too well cast.

But the historical lesson is not the melodramatic one. Mega-IPOs in this panel do not reliably ring the bell at the top of the market. The stronger pattern is narrower and more useful: large new listings often struggle after public trading begins. Historical IPOs raising at least $3 billion lagged their local benchmarks by 14.1 percentage points over the next twelve months, even though those same benchmarks rose by 8.1 percent on average.

The cross-section explains where the risk is largest. Story-heavy issuers had a mean twelve-month excess return of -23.9 percentage points, compared with -5.1 points for mature issuers. That does not prove that a fixed share of any mega-IPO price is sentiment. It does show that the public-market window matters most where the valuation depends most on future realization.

For SpaceX, the right posture is neither contempt nor prophecy. The company can be extraordinary. The window can still matter. The bell that rings at a mega-IPO is not necessarily the bell at the top; more often, it is the sound of a window closing behind the buyer.

References

[1] Greene, Edward F., Douglas W. Arner, and Iris H.-Y. Chiu. 2022. "The Need for a Comprehensive Approach to Capital Markets Regulation." *Columbia Business Law Review*. doi:10.52214/cblr.v2021i2.8637.