INVESTMENT FRAMEWORK — AN EXPERIMENT IN APPLIED AI – NOT INVESTMENT ADVISE – WINK WINK!

From the Industrial Revolution to the Space and AI Revolution

A four-stage conviction framework for investing in Space and AI via the upcoming IPOs

---

Building on the instinct that Space and AI will have a similar impact as the Railways did during the Industrial Revolution, I will be investing in these IPOs and in a certain way. If you are thinking of investing this thought process and execution may help. Inspired from the *"Investing in 100K in SpaceX GanjaGang Chat"*

Reader's Map — The Four Stages – linked so you can skip.

THE UNIVERSAL INVESTMENT CONVICTION

1. The Thesis: A Once-in-a-Generation Inflection

Every two or three generations, humanity goes through a foundational shift in how value is created. The steam engine and the factory system reorganized the 19th century. Electrification and the assembly line reorganized the 20th. The internet and mobile reorganized the early 21st. We are now at the start of the next shift — and it is not one shift, but two converging at the same time.

Space becomes infrastructure

For sixty years, space was a sovereign affair — flags, prestige, and a small number of government missions. That is no longer true. Reusable rockets have collapsed the cost of getting to orbit by more than 90% in fifteen years, and the next generation of vehicles is targeting another order of magnitude. Once launch becomes a commodity, everything that depends on it — global broadband, earth observation, defense, eventually orbital compute and manufacturing — moves from speculative to commercial. That is the same pattern as railroads in the 1850s and electrification in the 1920s: an enabling layer becomes cheap enough that entire industries are built on top of it.

AI becomes general-purpose technology

In parallel, large language models have crossed a threshold where they can do meaningful cognitive work — drafting, coding, reasoning, customer service, research synthesis — at a price that makes them economically deployable across every knowledge industry. Enterprise adoption is happening at a pace no software category has matched. Annualized revenue for the leading AI labs has gone from millions to tens of billions in roughly thirty-six months. Whether or not the most ambitious capability claims pan out, the productivity layer is being rebuilt.

Why the convergence matters — and the honest caveat

The thesis is that space and AI are becoming intersecting investment narratives — with orbital data centers, AI-controlled satellite constellations, and global compute that does not depend on terrestrial power grids representing the long-horizon convergence point. SpaceX and Anthropic are positioned at this intersection. Importantly: orbital compute commercialization remains technologically unproven and should be treated as forward optionality, not established trajectory. The convergence framing strengthens the case for taking exposure to both companies, but the specific orbital compute thesis is one that may or may not materialize over the 10-year hold horizon.

2. The Case for SpaceX

SpaceX is the dominant launch provider on Earth, with more orbital launches in a single recent year than every other provider combined. Through Starlink, it is also a top-tier global telecommunications operator with an exceptionally rare combination — vertical integration, recurring subscription revenue, and a cost structure competitors cannot replicate without rebuilding the entire stack.

The numbers, fact-checked

- **Starlink subscribers:** More than 10 million globally as of early 2026, growing at roughly 20,000 net additions per day during peak periods.
- **Starlink revenue:** Approximately $11.4 billion in 2025, with consensus 2026 forecasts in the $18–20 billion range. Segment EBITDA margins are approximately 50–55% (note: the 60%+ figure that circulates in press coverage typically refers to contribution margin before allocated overhead, not true EBITDA).
- **Cost per kilogram to orbit:** Reduced from roughly $15,600 in 2008 to under $1,000 today; Starship targets sub-$100 in the next decade.
- **Rumored IPO target:** Approximately $1.75 trillion valuation, June 2026 listing window, with 21 banks engaged. Confidential S-1 filed April 1, 2026.

Why this is investable — the moat analysis

An economic moat is what keeps competitors from taking a profitable business's profits. SpaceX has what investors call a compound moat — multiple moat types reinforcing each other. This is rare and is the single biggest reason a 109x revenue multiple is even arguable.

- **Scale economies (the deepest moat):** SpaceX launches rockets at a per-kilogram cost no competitor can match. The Falcon 9 reusability program took 15 years and billions in R&D. The cost gap is currently roughly 5–10x versus traditional providers, and the advantage compounds — cheaper launches mean SpaceX can deploy Starlink satellites at a fraction of any competitor's cost.
- **Regulatory and orbital scarcity:** Low Earth Orbit is finite real estate. Spectrum allocations are finite licenses. SpaceX moved first and now occupies orbital slots and spectrum bands that competitors literally cannot acquire because they're already taken.
- **Network effects (Starlink):** More satellites mean better coverage and lower latency, which attracts more subscribers, which funds more launches. This is a genuine, accelerating network effect.
- **Switching costs (defense and enterprise):** Pentagon contracts built around Starshield, maritime fleets equipped with Starlink terminals, airlines integrated with in-flight Wi-Fi — switching to a competitor is operationally painful.

Valuation scenarios — contextualizing the $1.75T target

The IPO valuation reflects one specific path through several plausible 10-year outcomes. Three scenarios bracket the range:

Scenario	Implied Valuation	What has to be true	Probability assessment
Bear	~$250B	Starship stalls, Starlink TAM caps at developed markets, regulatory erosion compounds, xAI burn forces dilutive raise	Real but lower probability — requires multiple negative events to converge
Base (IPO target)	~$1.75T	Starlink continues subscriber growth, Starship reaches commercial scale by late decade, Starshield contracts grow, no major regulatory shock	Most plausible single outcome but still requires significant execution
Bull	~$4T	Orbital compute commercializes, Starship enables space-based manufacturing, Starlink saturates emerging markets, xAI integration produces meaningful revenue	Possible but assumes near-perfect execution across multiple unproven trajectories

> **THE VALUATION GRAVITY PROBLEM — THE SINGLE LARGEST RISK**
>
> At a $1.75 trillion target on roughly $15–16 billion of 2025 revenue, the multiple is approximately 109x. This is higher than profitable SaaS comparables — and SpaceX is a hardware and logistics company with massive ongoing capital requirements for Starship development, satellite replenishment, and ground infrastructure. The market is not paying for 2026 revenue. It is paying for a 2030–2035 thesis that includes Starship operating at sub-$100 per kilogram, orbital data centers generating compute revenue, and Starlink reaching saturation in dozens of new markets. A 12-month delay on any of those — a regulatory pause, a catastrophic test failure, a slower-than-expected Starship ramp — does not change the long-term thesis but can compress the multiple by 30–50% with no offsetting fundamental change. This is the single largest risk to a 10-year hold, and it is the risk that most retail investors most reliably underestimate.

Why a thoughtful investor still pauses — the structural risk inventory

- **Valuation gravity (single largest risk):** The 109x multiple is pricing perfection over a decade. See callout above.
- **Key-person fragility:** Elon Musk's 78% voting control means public shareholders have zero governance leverage. Investors are betting on the long-term stability of one person's reputation, attention, and political and legal trajectory. Institutional controversy clauses can trigger forced

selling regardless of fundamentals — a sufficiently large reputational event could cause a 30–50% drawdown that has nothing to do with Starlink growth or Starship cadence.

- **Regulatory erosion (not regulatory binary):** The dangerous outcome is not a single shutdown ruling — it is gradual TAM compression as India, the EU, China, and the FCC tighten access. Cumulatively, these can take 20–30% off the long-term TAM that supports the current multiple.
- **Convergence and capital intensity:** Consolidated EBITDA closer to 50% than press coverage suggests, combined with continued Starship and xAI burn, points to a likely dilutive secondary offering within 24–36 months post-IPO. Additionally, a meaningful portion of the long-term valuation depends on orbital compute commercialization that remains technologically unproven — if that optionality fails to materialize by 2029–2030, a key pillar of the bull thesis dissolves. Competition (Project Kuiper, AST SpaceMobile, state-backed programs) is real but not the proximate threat.

3. The Case for Anthropic

Anthropic is one of two leading frontier AI labs, founded in 2021 by former OpenAI researchers around an explicit safety-and-alignment thesis. Its Claude family of models is the only frontier model available across all three major cloud platforms (AWS, Azure, Google Cloud), and its enterprise traction is the strongest evidence that the AI productivity thesis is becoming real economic activity rather than speculation.

The numbers, fact-checked

- **Annualized revenue:** Approximately $30 billion as of early April 2026, up from roughly $1 billion at the end of 2024 and $9 billion at the end of 2025.
- **Enterprise composition:** Approximately 80% of revenue from business customers; over 300,000 business customers; eight of the Fortune 10 are Claude users; over 500 customers spending more than $1M annually.
- **Most recent valuation:** $380 billion (Series G, February 2026); investor offers reportedly received at $800 billion in April 2026.
- **Rumored IPO target:** October 2026 window with Goldman Sachs and JPMorgan; raise potentially exceeding $60 billion.

Why this is investable — the moat analysis

Anthropic has moats. But they are categorically different from SpaceX's — narrower, more recent, and more contestable on shorter time horizons.

- **Talent concentration:** Roughly 200 researchers including a meaningful percentage of the people in the world capable of training frontier models. Real but fragile — researchers move, and talent moats erode faster than physical ones.
- **Enterprise switching costs:** Once a Fortune 500 company integrates Claude into core workflows, switching is operationally painful. Real switching cost, but measured in months not decades.
- **Distribution position:** Claude is the only frontier model available across all three major hyperscalers. Real distribution advantage, but structurally fragile because the same hyperscalers are also competitors.
- **Safety brand:** Real differentiator in finance, healthcare, government, and legal verticals where compliance officers can defend the choice on safety grounds.

> **THE AGI SAFETY PARADOX — A GENUINE TRADEOFF, NOT AN UNAMBIGUOUS MOAT**
>
> Anthropic's safety-first positioning is real and creates real enterprise trust. Regulated industries explicitly value the alignment focus. But safety in an AGI race is not unambiguously a moat. It can also be a capability liability. If a less constrained competitor (OpenAI, Google DeepMind, or a

> leading Chinese lab) achieves a meaningful capability step-change first, the enterprise customers who chose Claude for safety may find their procurement criteria shifting toward whoever has the demonstrably more capable model. Safety as enterprise trust moat versus safety as capability lag — both can be true, and which one dominates determines the long-term valuation outcome.

Why a thoughtful investor still pauses — the structural risk inventory

- **Profitability is genuinely far away and the moat is contestable:** Anthropic plans to spend approximately $10–12 billion on training and inference in 2026 — close to its full revenue. Gross margins have fallen to roughly 40% as inference costs ran 23% over plan. This is an arms race with negative operating margins, and the moat is structurally fragile: Anthropic depends on hyperscalers (AWS, Google Cloud, Azure) for compute, and those same hyperscalers are also competitors building their own foundation models.
- **Open-source convergence risk:** Meta's Llama and Mistral are closing the gap on reasoning benchmarks at every release cycle. If a smaller, cheaper, self-hostable model can do 80–90% of what Claude does at 10–20% of the cost, enterprise procurement will move — that is how Linux displaced proprietary Unix and Postgres displaced Oracle. The lag may be three years or seven; the directional risk is real.
- **Valuation discipline and the AGI safety paradox:** Doubling from $380B to $800B in two months — with no change in revenue trajectory in those two months — is the kind of price action that historically precedes resets. Compounding the risk: see callout above on safety as enterprise trust moat versus safety as capability lag.

4. The Universal Conclusions: 80/20 Split and Hold Horizons

The two preceding sections establish a side-by-side moat comparison. SpaceX has a compound moat across four reinforcing categories — scale economies, regulatory and orbital scarcity, network effects, and switching costs. Three of those four are physical assets. Physical moats erode in decades. The Suez Canal still has a moat after 150 years. Standard Oil's pipeline moat survived multiple antitrust eras.

Anthropic's moats — talent concentration, enterprise switching costs, distribution position, safety brand — are real but categorically different. They are intangible, recent, and constantly under attack from open-source convergence and hyperscaler competition. Intangible moats erode in years. Yahoo had one of the strongest intangible moats in 2000 — brand, distribution, talent. By 2010 it was gone.

The 80/20 split is a direct conclusion from this asymmetry

Given the structural asymmetry between a compound physical moat and a narrower intangible moat, the position should be weighted toward the more durable asset. Weighting 80% toward SpaceX and 20% toward Anthropic reflects this asymmetry honestly — meaningful exposure to the faster-growing but more contestable position, with the larger commitment to the more durable one. This is not a personal preference. It is the natural conclusion of comparing a moat measured in decades against a moat measured in years.

> **WHY 80/20 AND NOT ANOTHER SPLIT — READ CAREFULLY**
>
> Reasonable investors might land at 90/10 or 75/25 depending on their existing portfolio composition. But the structural argument for SpaceX-heavy weighting is robust: an investor who weights more than 30% toward Anthropic is effectively saying that intangible AI moats are as durable as compound physical moats, which is a claim the historical evidence does not support. An investor who weights less than 10% toward Anthropic may be reasonable if they already hold significant AI exposure through NVIDIA, Microsoft, or Google. The 80/20 default is the structurally honest expression of the moat comparison; deviations should be conscious responses to specific portfolio circumstances, not arbitrary preferences.

Hold horizons follow the same logic

Physical moats compound over decades. The timeframe required to capture that compounding is roughly a decade — long enough for orbital infrastructure investments to mature into recurring revenue at scale, for Starship economics to play out, and for the network effects of a fully-deployed satellite constellation to fully manifest. A 10-year hold for SpaceX is the natural match between the asset's compounding period and the investor's time horizon.

Intangible moats in fast-moving competitive landscapes require periodic re-underwriting. AI competitive cycles run on 12–24 month rhythms — a major model release, a regulatory shift, a benchmark inflection.

Five years allows the investor to observe two to three full competitive cycles before deciding whether the original thesis still applies. A 5-year mandatory review for Anthropic — explicitly a review, not a forced exit — is the natural match between the asset's competitive cycle and the discipline required to hold it.

> **THE UNIVERSAL CONCLUSION OF INVESTMENT CONVICTION**
>
> The argument up to this point is that any investor convinced by the thesis should take exposure weighted approximately 80% to SpaceX and 20% to Anthropic, holding SpaceX for a decade and reviewing Anthropic at year five. These conclusions are universal — they flow directly from the moat analysis and do not depend on the individual investor's circumstances. What does depend on individual circumstances is how much capital to commit. That is the subject of Our Peer Group Sizing and 80/20 Split, where we narrow to our peer group and land the sizing range.

The Contrarian Stress Test — Reading the Bear Case Charitably

Before moving to commitment, the framework demands an honest engagement with the strongest version of the bear case across both companies. Confirmation bias is the most dangerous failure mode for high-conviction investors, and it is best neutralized by deliberately steel-manning the contrary view rather than dismissing it. The bear case in its sharpest form: SpaceX's Starship program faces multi-year delays that compress the multiple by 40%+ before commercial reusability is proven; Elon Musk's attention divides further across xAI, Tesla, X, and political activity, and a major reputational event triggers institutional ESG divestment; orbital compute commercialization fails to materialize by 2030, eliminating a key pillar of the bull thesis; the post-IPO capital raise within 36 months dilutes early shareholders by 5–10%. Simultaneously: Anthropic faces capability leapfrogging by a less constrained competitor, open-source models close the 80–90% performance gap at 10–20% the cost, and enterprise procurement shifts away from Claude as the safety brand becomes secondary to capability. Both scenarios are coherent. Both are plausible. The framework's response is not to dismiss them but to size the position so that being wrong is recoverable — which is precisely what the 1.5–3% sizing range in Our Peer Group Sizing and 80/20 Split enforces.

OUR PEER GROUP, THE SIZING RANGE, AND THE 80/20 AT $100K

5. The Commitment — Take the Position, 80/20, Sized to Your Portfolio

We are all somewhere in the $1M–$10M+ net worth range. Investment Conviction Thesis established the conviction: 80% SpaceX, 20% Anthropic, 10-year hold on SpaceX, 5-year review on Anthropic. This section asks the next thing: take the position now, in that structure, sized to your portfolio. The 80/20 split stays fixed at any sizing level. Only the absolute dollars change with your net worth.

The 80/20 commitment — what we are actually doing

At a $100K commitment (2% of $5M, the median answer for our peer group), the position is $80K SpaceX and $20K Anthropic. That is the conviction-driven weighting derived from the moat asymmetry — a compound physical moat measured in decades against a contestable intangible moat measured in years. Within that $100K, every subsequent rule in this document carries a concrete dollar value. Day 1 deployment of 40% means $32K SpaceX and $8K Anthropic. The -10% drawdown trigger commits another $24K SpaceX and $12K Anthropic. The -20% trigger commits the final $24K SpaceX. The 5% concentration trim trigger fires when SpaceX appreciates from $80K to $250K — roughly 3x. These numbers are memorable, executable, and don't require recalculation in real time.

Position	$ at $100K total	Why this weighting
SpaceX (80%)	$80,000	Compound physical moat — scale economies, orbital scarcity, network effects, switching costs. Erodes in decades.
Anthropic (20%)	$20,000	Narrower intangible moat — talent, switching costs, distribution, safety brand. Contestable in years.
Total	**$100,000**	2% of $5M — within the 1.5–3% defensible range for our peer group

The sizing math — 1.5% to 3%, with 2% as the median

For our peer group convinced by the thesis, the right allocation falls between 1.5% and 3% of net worth. Below 1.5% is symbolic — too small to matter if the thesis plays out, too small to teach you anything if it doesn't. Above 3% crosses from investment into speculation — a complete loss of more than 3% starts

to materially affect financial trajectory rather than being absorbed by ordinary returns. Where you land within that range depends on five factors: liquidity composition (net worth often includes home equity and retirement accounts that aren't relevant for sizing); existing macro exposure (significant NVIDIA/Microsoft/Google holdings amplify rather than diversify); emotional tolerance (the dollar amount you can watch go to zero without damaging your judgment); competing opportunities (this should be sized as a fraction of your discretionary risk budget if other high-conviction bets are competing); and life constraints (a 10-year lock-up has different implications based on dependents, education timing, and near-term liquidity needs).

Most of us will land at 2%. Some will land at 1.5% with significant existing tech exposure or near-term liquidity needs. Some will land at 3% with no existing AI exposure and strong conviction. At $1.5M net worth, that range is $22.5K to $45K. At $5M, $75K to $150K with $100K as the median. At $10M+, $150K to $300K, with the upper end becoming more defensible as net worth grows because the absolute-dollar ceiling for ‘speculation' rises with available capital.

THE CORRELATION ILLUSION — A SIZING CONSIDERATION

The 80/20 split between SpaceX and Anthropic looks like diversification across two industries. It is not. Both positions are proxies for the same macro factor: high-multiple, high-growth, tech-heavy, AI-tailwind exposure. In a 2022-style rate-driven tech selloff, both drop together. In a regulatory crackdown on Big Tech, both face headwinds. This allocation is a concentrated single-theme bet, not a diversified position. If your broader portfolio is already 40%+ in growth tech, push your sizing toward the 1.5% floor of the range rather than the 3% ceiling.

THE LIQUIDITY STRESS TEST — RUN THIS BEFORE COMMITTING

Beyond the foundational precondition below, run this specific simultaneity test: imagine a scenario in which both SpaceX and Anthropic experience a 60% drawdown at the same time AND you face an unexpected personal liquidity need (medical event, job loss, education funding pulled forward, real estate opportunity). Can you absorb both simultaneously without selling at the bottom of the drawdown? If yes, your sizing is appropriate. If no — if a 60%/60% combined drawdown plus a liquidity event would force you to sell at the worst possible time — size down to 1.5% rather than 2% or 3%. The correlation illusion above means you cannot count on one position offsetting the other in a stress scenario.

FOUNDATIONAL PRECONDITION BEFORE ANY SIZING DECISION

Before committing capital: confirm you have 6 months of emergency savings, maxed retirement contributions for the year, no high-interest debt outstanding, and no other immediate large capital needs (down payment, education funding, medical reserves). This framework assumes those

> foundations are in place. If they are not, this allocation should wait until they are.

MY EXECUTION APPROACH

6. Phase One — The Portfolio Rebalancing Opportunity

WHAT THIS PHASE IS FOR

A high-conviction IPO does not cause the need for portfolio cleanup — it creates the forcing function for one. Every long-only portfolio drifts. Positions appreciate at different rates. Conviction levels shift. Tax-loss harvesting opportunities accumulate. The IPO is the catalyst for rebalancing the disciplined investor should be doing periodically anyway, with the added benefit of a clear destination for the freed capital. This section describes how I am personally approaching that rebalancing. The framework applies broadly; the specific positions and tax mechanics differ for every reader.

The thinking — rebalancing as opportunity, not burden

Most investors hold portfolios that are long overdue for honest review. Stocks held out of inertia rather than conviction. Once-core positions that are no longer core. Dividend payers in declining industries kept for the dividend even though the underlying business has weakened. Beat-up names retained on hope of recovery rather than analysis. These holdings are not generating conviction-weighted returns; they are generating market beta with idiosyncratic downside. They are exactly the right source of capital for a high-conviction position — and they should probably be sold even if you weren't planning to fund this IPO.

THE HONEST DISCIPLINE CHECK

The rebalancing framing is more powerful than 'liquidation' precisely because it refuses to be used as cover for undisciplined selling. The test for every candidate position: would I be selling this if I weren't planning to invest in SpaceX and Anthropic? If yes, it is real rebalancing. If no, it is liquidation with a flattering name. The framing only works if the discipline is honest. Apply this check to every position before selling.

THE STATE TAX REALITY — CRITICAL FOR HIGH-TAX-STATE RESIDENTS

The federal long-term capital gains rate is roughly 23.8% (20% federal + 3.8% Net Investment Income Tax). That is not the full picture. California residents face an additional 13.3% state rate on capital gains. New York City residents face approximately 14.776% combined state and city. New Jersey, Oregon, Minnesota, and several other states stack 8–11% on top of federal. For a high-net-worth investor in California liquidating long-term winners, the total tax drag on each dollar realized can be 35–40%. That is not a tweak to the math — it is a different math entirely. Readers in high-

tax states should explicitly consider whether single-event rebalancing is even viable. For many, the only sensible funding mechanism is monthly cash flow over 12–24 months, building the position gradually rather than through a one-time portfolio restructuring. For investors with significant charitable intent at the upper end of the peer group, donating appreciated SpaceX or Anthropic shares directly to a Donor-Advised Fund post-lockup can capture fair-market-value deductions while avoiding the LTCG hit entirely — discuss with your CPA before assuming this works for your situation.

THE AMT STRESS TEST — RUN BEFORE LIQUIDATING WINNERS

Before executing any rebalancing of long-held appreciated positions, run an Alternative Minimum Tax simulation with your CPA. Large concentrated capital gains realized in a single tax year can trigger AMT in ways most retail investors do not anticipate, particularly for those exercising stock options, holding ISOs, or having complex deduction profiles. If AMT is triggered, the right response is almost always to switch to the 12-month cash-flow funding path rather than concentrate the realization event into a single tax year. This check is non-optional and goes before the rebalancing execution, not after.

My generalized approach — the steps anyone can adapt

- **Step A — Inventory:** Pull a full holdings statement and identify positions where you no longer have an active investment thesis. "Held since 2018" is not a thesis. Be honest about which holdings are conviction-driven versus inertia-driven.
- **Step B — Cost-basis and full tax audit:** Run each candidate position through a CPA review with full federal AND state tax modeling. Separate into three buckets: clear losses (harvest immediately), modest gains (sell if conviction shift justifies after-tax cost), and large long-term gains (evaluate carefully).
- **Step C — Apply the discipline check:** For every candidate, ask: would I sell this even if I weren't funding the IPO? If yes, it stays a candidate. If no, remove it. The IPO does not justify selling positions you would otherwise keep.
- **Step D — Funding decision:** Based on the audit, decide whether single-event rebalancing works or whether cash flow funding over 12–24 months is the better mechanism. For high-tax-state residents with mostly long-term winners, the cash flow path is often the only viable one.
- **Step E — Document the logic:** Write down one sentence per position explaining why you are selling. Forces honest engagement and creates an audit trail for tax season.
- **Step F — Execute:** Sell the candidate positions. The proceeds become your IPO allocation cash, with reserves for Phase 2 mechanical drawdowns and Phase 3 monitoring.

My specific rebalancing candidates are positions in my portfolio I no longer hold conviction in. Yours will be different. The framework — inventory, audit, discipline check, decide, document, execute — applies regardless of which specific positions you are evaluating.

7. Phase Two — Pure Asymmetric Deployment

The deployment framework that follows is meant to be followed as written. Pure Asymmetric Deployment is a discipline mechanism designed to protect against universal cognitive biases — particularly the impulse to chase rallies and to deploy more capital at peak enthusiasm. The four rules are structurally interdependent. Bending one bends them all.

The Four Rules

Rule 1 — Day 1 Entry (40% of position)

On the IPO day, deploy 40% of the planned allocation. At the $100K working example: $32K into SpaceX and $8K into Anthropic. This establishes skin in the game and ensures meaningful participation if the stock runs from open without retracing — while preserving 60% of capital for the more common scenario of post-IPO retracement within 90 days. Use a market order or post-auction limit order; limit orders set against the IPO opening auction often do not fill correctly. Hot offering price contingency: if the IPO prices more than 15% above current secondary market levels, reduce the Day 1 tranche by half.

Rule 2 — Mechanical Drawdowns (-10% / -20% from offering price)

Deploy 30% of capital at -10% from the IPO offering price; another 30% at -20%. At the $100K working example: $30K total at -10% ($24K SpaceX + $6K Anthropic), then $30K total at -20% (SpaceX-only — $24K, since Anthropic only has two tranches). Both triggers reference the offering price. Use Good-Till-Canceled limit orders so execution doesn't require real-time monitoring. Three execution overrides apply:

- **The wait rule:** When a trigger fires, wait 3 trading days for company-specific drops, 5 days for ambiguous events. Filters out flash crashes.
- **Systemic event override:** If the trigger is firing because the broader market is falling more than 15% (S&P), deploy without delay.
- **Thesis-break override (do NOT deploy):** Do NOT deploy if the trigger fires alongside confirmed bad news that materially damages the thesis. The price triggers exist to deploy on volatility WITHOUT a thesis change.

Rule 3 — Time Decay with a Hard Ceiling (18 months)

If neither drawdown trigger fires within 18 months post-IPO, the remaining undeployed capital deploys in equal monthly tranches over the following 12 months — approximately 5% of total allocation per month, or roughly $5K per month at the $100K working example. The hard ceiling: time decay deployment pauses automatically if the stock is trading more than +40% above the IPO offering price. You never deploy cash during bubble euphoria. If the stock retraces below +40%, deployment resumes

from where it paused; if it never retraces below the ceiling, capital remains undeployed permanently — by design.

Rule 4 — No Rally Triggers

The framework includes no triggers that deploy capital based on positive price momentum. If the stock runs to +50% above offering price, you hold your initial 40% and sit on your hands. The opportunity cost of missing the absolute top is mathematically outweighed by the risk of deploying dry powder at peak valuations.

> **WHY NO RALLY TRIGGERS — READ THIS CAREFULLY**
>
> Most investors instinctively want a rule that lets them deploy more when the stock is running. That instinct is the cognitive bias the framework is specifically designed to defeat. Capital deployed at peak enthusiasm has worse forward returns than capital deployed at peak fear. A framework that fires hardest at peak excitement systematically worsens average entry prices. The honest tradeoff: in the scenario where the stock runs to +60% and stays there for years, this framework will permanently underdeploy. That is not a flaw — it is the framework working as intended.

Contingency Protocols — Edge Cases the Standard Rules Don't Cover

Two scenarios warrant explicit deviation from the standard four-rule framework. Both are bounded modifications, not new rules — the underlying discipline (asymmetric deployment, no rally triggers, hard ceiling on euphoria) remains intact.

- **Missed Day 1 Protocol:** If you fail to execute the 40% Day 1 entry due to technical issues, brokerage delays, or hesitation, do not collapse the entire framework into accelerated deployment. Instead: deploy 30% at the -10% trigger (catching up to where standard Tranche 2 would have placed you), 30% at the -20% trigger (standard Tranche 3), and let the remaining 40% follow the standard 18-month time decay schedule from the original IPO date. This preserves the discipline without compressing the timeline; missing Day 1 means you have more capital to deploy thoughtfully, not less.
- **Overpriced IPO Protocol:** If the IPO prices more than 30% above current secondary market levels, defer Day 1 entry entirely. Convert the 40% Day 1 tranche into a third drawdown tranche at -25% from offering price. Standard -10% and -20% triggers remain unchanged. If the IPO prices more than 50% above secondary, additionally reduce total allocation to the 1.5% floor of the sizing range — the premium is large enough that forward expected return no longer justifies the upper end of conviction sizing.
- **Day 181 Lock-Up Watch:** The standard insider lock-up expires 180 days post-IPO. Predictable selling pressure typically materializes in the days surrounding expiration. If a price trigger fires within 5 trading days before or after Day 181, pause and observe rather than deploy — selling

pressure from lock-up expiration is mechanical, not thesis-driven, and waiting for it to clear typically produces a better entry price.

Anthropic-specific notes

Anthropic deployment is contingent on an actual public IPO announcement and pricing. As of this writing, no firm date is confirmed. If Anthropic does not IPO by December 2027, the allocation rolls into SpaceX as additional time-decay deployment or back into the broader portfolio. The same four-rule structure applies, with one adjustment: the Day 1 tranche is 40%, but the second drawdown tranche (at -10%) is sized at 60% of remaining (rather than the SpaceX 30/30 split). The larger second tranche reflects the heightened competitive risk in AI.

8. Phase Three — Monitoring, Governance, and the Measured Exit

Most investment plans fail not at the entry but in the holding period — the position drifts, the world changes, and the investor stays in by inertia. Phase 3 is the governance layer that makes a long hold an active decision rather than a default.

Quarterly review cadence

Once per quarter, spend thirty minutes reviewing both positions. The review is not about the price — it is about whether the original thesis still holds. Five questions to write down each quarter:

- Is the core revenue engine still growing? (Starlink subscribers and ARPU for SpaceX; enterprise revenue and customer count for Anthropic.)
- Has anything fundamental changed in the competitive landscape, regulation, or leadership?
- If I were not already in the position, would I buy it today at this price?
- Orbital Compute Viability (SpaceX-specific): Has SpaceX achieved credible milestones on orbital data center commercialization? Track successful orbital test deployments, cost-per-watt benchmarks for space-based compute, and customer pilots with meaningful contract values.
- Capital Raise Watch (SpaceX-specific): Is SpaceX showing signs of needing additional capital? Track quarterly cash burn, Starship development milestones, and any signals of secondary offering preparation.

Concentration discipline — the refined trim rule

If either position grows substantially through appreciation, partial trimming preserves discipline without forcing complete exit from a parabolic phase. The rule: when a position reaches 5% of net worth through appreciation, trim 25% of the position. When it reaches 7.5% of net worth, trim another 25%. When it reaches 10%, trim back to 5% — this is the hard cap. At the $5M / $100K working example: SpaceX trim trigger fires when the position reaches $250K (roughly 3x appreciation from the original $80K), then again at $375K, with a hard cap at $500K.

Cross-position harvesting — tax optimization between SpaceX and Anthropic

The two positions will likely diverge in performance. When divergence occurs and one position is at a meaningful unrealized loss while the other is at a meaningful gain, harvest the loss to offset the gain without changing the strategic 80/20 exposure. The mechanic: sell the losing position, immediately purchase a similar exposure (a sector ETF or related public company) for 31 days to satisfy wash-sale rules, then re-establish the original position. This preserves the strategic 80/20 weighting while capturing tax efficiency on the divergence. Coordinate with your CPA on timing — this is most valuable in years where you have other large capital gains to offset.

Binary thesis-break criteria — exit immediately if any trigger

SpaceX

- **Regulatory:** A major ruling that prevents Starlink or Starshield from operating in core markets, OR cumulative regulatory erosion that takes more than 25% off the long-term TAM assumption.
- **Operational:** Two consecutive years of revenue decline post-IPO.
- **Leadership:** Elon Musk steps away from or is removed from SpaceX.
- **Technological:** A competitor achieves successful orbital reusability at less than 50% of SpaceX's cost-per-kilogram for two consecutive years.

Anthropic

- **Competitive:** A competitor's model demonstrably outperforms Claude on enterprise benchmarks for two consecutive quarters AND Anthropic's enterprise revenue growth slows below 30% year-over-year.
- **Open-source convergence:** An open-source model achieves performance within 10% of frontier closed-source models across enterprise benchmarks AND total cost-of-ownership less than 25% of Claude's enterprise pricing.
- **Margins:** Gross margins fall below 30% for two consecutive quarters with no credible path to recovery.
- **Operational:** Loss of access to a major hyperscaler distribution channel.
- **Five-year mandatory review (review, not exit):** At year 5, formally re-underwrite the position from scratch. AI moves fast enough that a position bought in 2026 is operating in a different industry by 2031. The review is a discipline mechanism. The default outcome of a positive review is to continue holding.

The measured exit

When a thesis-break trigger fires or the natural hold horizon arrives, exit in three tranches over 60 days rather than a single sale. The exception is a hard regulatory or leadership event, in which case exit is immediate.

9. Closing the Loop — The Multi-Decade Thematic Discipline

This section reframes everything that came before it. Exiting SpaceX after the 10-year hold or Anthropic after the 5-year review is not the end of the thesis. It is the end of the first expression of it. The Space and AI revolution is a multi-decade structural shift on the scale of the Industrial Revolution itself — and what we are doing with these two IPOs is taking a position in the first generation of investable companies that emerge from that shift. The next generation is already forming.

The next layer of the Space/AI stack

As SpaceX and Anthropic mature, the investable surface area of the underlying thesis expands. The 2030s will produce the next layer of public companies in this theme — companies that today are early-stage, private, or not yet founded. Three categories worth tracking specifically:

- **Orbital manufacturers:** Companies producing semiconductors, pharmaceuticals, fiber optics, and specialized materials in microgravity environments. Today these are early experiments. By the early-to-mid 2030s, with Starship-class launch costs, several will reach commercial scale. The first IPO in this category may well happen during your SpaceX hold.
- **Robotics integrators:** The companies that turn frontier AI models into embodied capability — autonomous warehouse systems, surgical robotics, agricultural automation, defense applications. Anthropic and OpenAI build the cognitive layer; the integrators turn it into physical-world value. The category is emerging now and will produce major public companies through the late 2020s.
- **Specialized data centers:** Whether terrestrial (next-generation power, cooling, and AI-optimized architecture) or eventual orbital (the convergence thesis you are partially betting on with SpaceX). The infrastructure layer that supports continued AI scaling will produce significant winners over the 2030s.

How to think about the next position

When you exit SpaceX or Anthropic — whether by hold completion, thesis-break trigger, or measured exit — the proceeds are not random capital looking for a home. They are thematic capital looking for the next expression of the same conviction. The discipline you built here applies again with no architectural changes:

- **Same conviction architecture:** Compound moats over intangible ones. Physical scarcity over commodity competition. Verified moat analysis before commitment.
- **Same sizing range:** 1.5–3% of net worth per high-conviction position, with 2% as the median. Multiple thematic positions across the Space/AI stack are reasonable; concentration in any single name is not.
- **Same Pure Asymmetric Deployment:** 40% on Day 1, 30% at -10%, 30% at -20%, time decay with +40% ceiling, no rally triggers. The framework was built to be repeatable.

- **Same governance:** Quarterly review, trim at concentration thresholds, binary thesis-break criteria, measured tranched exit.

> **THE REAL COMMITMENT**
>
> Reading this document and committing to SpaceX and Anthropic is not the commitment that matters most. The commitment that matters is the multi-decade one — recognizing that the Space and AI revolution will produce a series of investable expressions over the next 20 years, and committing to the discipline framework that lets you take thoughtful exposure to each one as it emerges. SpaceX and Anthropic are the first IPOs you will encounter in this theme. They are not the last. The framework you are using here is the one you will use again — for the orbital manufacturer that goes public in 2031, the robotics integrator that goes public in 2033, the specialized data center that goes public in 2035. Build the muscle now.

This is the closure that makes the entire framework worth building. A document that just tells you how to invest in two IPOs is a tactical guide. A document that tells you how to take disciplined exposure to a multi-decade structural shift, starting with two IPOs, is a strategic framework. The work was always the second one.

10. Final Notes

This is not a prediction that SpaceX or Anthropic will outperform. Historical research on IPO performance (Jay Ritter, University of Florida) consistently finds that a substantial majority of high-multiple IPOs trade below their offering price at some point during their first three years, and that high-multiple tech IPOs underperform broader market benchmarks more often than not over their first three years post-listing. The thesis is that a small, disciplined, time-staggered position in two companies sitting at the convergence of two structural shifts is a defensible use of 1.5–3% of net worth — and even if both positions disappoint, the portfolio outcome is unaffected at these sizing levels.

Before executing: complete a CPA review with full federal AND state tax modeling on rebalancing candidates; complete a fee-only fiduciary review of the overall plan; honestly assess whether your existing portfolio's tech exposure pushes your sizing toward 1.5% rather than 3%; commit in writing to the four rules of Pure Asymmetric Deployment, especially Rule 4. The rule that bends in one direction will bend in the other. And acknowledge that the conviction may turn out to be wrong — the position size is what makes that survivable.

The investors who did well during the railroad boom were not the ones who were most certain — they were the ones who were sized appropriately for being wrong. The right way to invest in a once-in-a-generation shift is not to bet the portfolio on it. It is to take a position large enough to matter if you are right, and small enough that being wrong is recoverable. This framework is built around that principle, with the explicit expectation that the discipline applied here will be applied again to the next layer of the Space/AI stack as it emerges.

> **IMPORTANT — READ FIRST & HOW TO READ THIS DOCUMENT**
>
> This document - created by MogamboAI guided and edited by Amit - is a personal investment framework shared between friends. It is not financial advice, not a recommendation, and not a solicitation. Its an Experiment! Every reader should consult a fee-only fiduciary financial advisor and a CPA before acting on anything in this document. The framework was developed in dialogue with Amit, MogamboAi (GPT+Anthripic+Qwen+Mistral+nVidia inspired LLMs). Mogambo is learning hence any advice to improve experimentation is appreciated.