


Macfarlane Crusade Fund L.P

Abundance of Underdeliver & Overpromise

Tesla Inc. (\$TSLA)



Disclosure: The author holds a short position in TSLA at the time of publication. This report is published for informational purposes only and does not constitute investment advice or a solicitation to buy or sell any security. All opinions are those of Macfarlane Research. Past performance is not indicative of future results. Investing involves risk, including the possible loss of principal.

1. Executive Summary

376x earnings. \$1.5 trillion market cap. Priced as though Tesla has already won the race in autonomous driving, humanoid robotics, and energy storage. However, the underlying business tells a very different story and the gap between narrative and reality has never been wider.

The core automotive business is deteriorating. EPS has collapsed 75% since 2023. Revenue fell for the first time since IPO. Deliveries declined for the second consecutive year. Market share has been halved in the U.S. and Europe. China domestic sales dropped 45% year-over-year. BYD now outsells Tesla by over 600,000 vehicles annually. This is the business that funds everything else.

The two product lines justifying roughly \$1.2 trillion of implied value have generated zero disclosed revenue. The Austin robotaxi fleet has 1–2 functionally active unsupervised vehicles on any given day, a crash rate 3.5x worse than the commercial insurance threshold, and no driverless authorization from any state DMV.

Optimus produced a few hundred units against a 10,000-unit 2025 target. Material robotaxi revenue arrives in the 2029–2031 window. The market is pricing 2027. The core of our thesis is a timeline mismatch. Consensus embeds material Cybercab revenue by 2027–2028. We model 2029–2031. The gap is not a matter of opinion, it is a regulatory cascade with specific dates: Texas AV permits do not open until May 28, each new city requires its own approval process, and Tesla has not filed the federal Part 555 exemption required to sell a vehicle without a steering wheel. Musk announced seven expansion cities by June 30, 2026; today Tesla has 1–2 actively operating unsupervised vehicles in one city. Even if the technology works, Tesla needs 16,500 Cybercabs just to generate \$1 billion in annual revenue. We assign zero robotaxi revenue in 2026 and \$25–50 billion in option value in our SOTP. Sell-side targets embed \$300–500 billion. That gap is the trade.

Tesla is entering its most aggressive spending cycle in company history at the exact moment its revenue base is shrinking. Capex exceeds \$20 billion in 2026, more than double 2025, producing negative free cash flow of \$7–8 billion. The CFO is already in discussions with banks about financing. That \$20 billion excludes the TeraFab semiconductor facility which Musk announced March 14, 2026. The rescue thesis requires three things simultaneously: crash rates improving 3.5x to clear the commercial insurance threshold, regulatory approval for driverless operation (not granted anywhere), and capital sufficient to fund it all. None are on track.

The people who would have to deliver the autonomy story are leaving: head of Autopilot, VP of Engineering, CFO, and multiple senior FSD engineers since January 2025. The next six months are a series of hard checkpoints: NHTSA crash data submission (missed twice already), Q1 earnings with the capex ramp live, Texas AV permits on May 28, seven expansion cities by June 30. Each is a test the stock has not yet had to pass. If they fail, this is the third consecutive year of broken autonomy timelines.

Tesla has \$44 billion in cash, a real energy storage business at ~30% margins, and Musk has defied skeptics before. The thesis is simply that at 376x earnings the stock prices near-perfect execution across

multiple unproven business lines. The evidence says execution is falling dramatically short. A compression to 100–150x implies 40–60% downside.

2. Declining Automotive Business & Market Share

Tesla's automotive segment still accounts for roughly 70% of total revenue. It is this business that funds everything else, the robotaxi pilot, the Optimus R&D, and the AI compute buildout. And it is this business that is deteriorating across every measurable dimension.

The Financial Damage

Diluted EPS peaked at \$4.30 in 2023, fell to \$2.04 in 2024, and collapsed to \$1.08 in 2025, a 75% decline in two years and a 70% decline from 2022's \$3.62. Total revenue dropped 3% to \$94.8 billion, the first annual decline since Tesla's IPO. Automotive revenue specifically fell 10%. Deliveries declined for the second consecutive year, falling to 1.64 million units from a peak of 1.81 million in 2023. Operating margin compressed to 4.6%, and net income fell 47% year-over-year to \$3.79 billion. Our 2026 EPS estimate sits at +\$0.87 versus a Street consensus of \$2.25, a 61% discount driven by ASP compression, margin deterioration, and continued delivery stagnation. Even this relatively modest figure depends heavily on \$2.0 billion in below-the-line interest and other income to offset near-breakeven auto operations. The directional message is clear: the trajectory of the core auto business is sharply negative, and the Street is materially too optimistic.

Market Share Erosion

Tesla pioneered the modern EV market and once held near-monopoly positions across the United States, Europe, and China. That dominance has eroded steadily in every region as competition has intensified.

Year	U.S. EV Share	EU BEV Share	China BEV Share
2020	~79%	7–8%	10–12%
2021	~70%	13–14%	7–9%
2022	~75%	14–16%	~7.5%
2023	~55%	~18%	~6%
2024	~49%	~17%	5–6%
2025	~46%	~9%	~5%

In the United States, Tesla's share of EV sales has been cut nearly in half, from approximately 79% in 2020 to 46% in 2025. GM is now the second-largest EV seller at roughly 13% share, and nearly 90 EV models are available in the U.S. market today. Tesla remains the leader, but its dominance is eroding quarter by quarter.

Europe represents the most rapid deterioration. Tesla's BEV market share peaked at approximately 18% in 2023, then collapsed to roughly 9% by early 2025. European sales fell 28% in 2025 while the overall BEV market grew 27%, meaning Tesla didn't just lose share, it shrank in an expanding market. VW Group, BMW, and Chinese entrants like BYD and MG are gaining rapidly.

China is the most alarming recent development. Domestic sales plummeted 45% year-over-year in January 2026 to just 18,485 units. The Model Y fell from the best-selling electric SUV in December to 20th place in January. Xiaomi, a consumer electronics company that launched its first vehicle in early 2024, sold more than twice as many cars as Tesla's Model Y in January alone. Tesla's 2025 annual report showed a staggering structural shift: U.S. profits outweighed foreign pretax profits for the first time, despite the Shanghai factory producing roughly half of global deliveries. BYD, which was dismissed by Musk, has dramatically expanded its product lineup and rapidly expanded internationally, leaving Tesla increasingly on the back foot.

Product Lineup Is Aging and Narrowing

Tesla has the oldest product lineup in the EV industry. The Model 3 and Model Y are the only volume vehicles. The Model S and Model X are to be discontinued in Q2 2026, accounting for less than 3% of deliveries. Musk scrapped the planned \$25,000 "Model 2", the affordable vehicle that would have competed directly in the segment where BYD dominates, in favor of the Cybercab, which requires Level 4 FSD approval to function and has no timeline for that approval. The 2025 Model Y update amounted to a facelift: revised exterior lighting and interior refinements, not a platform redesign. Competitors typically execute full ground-up product cycles on three-to-four year timelines; Tesla's core volume vehicles have not had one.

The Cybertruck is a case study in overpromising. Every major specification missed its 2019 target:

Specification	2019 Promise	Delivered
Starting Price	\$39,900	\$60,990 (+53%)
Top Model Price	\$69,900	\$99,990 (+43%)
Range (Top Model)	500+ miles	320 miles (-36%)
Towing Capacity	14,000 lbs	11,000 lbs (-21%)
Payload	3,500 lbs	2,500 lbs (-29%)
Annual Sales Target	250,000 units	~20,237 (2025)

Sales collapsed 48% in 2025 to just 20,237 units representing the largest single-year sales drop of any electric vehicle in America. The truck has now been recalled more than ten times in its first two years on the market. Tesla holds an estimated \$800 million in unsold Cybertruck inventory, with units stacked across storage lots and empty mall parking lots visible from satellite imagery, and more than 80% of its claimed production capacity sitting idle. In a telling admission of the

program's failure, Musk has resorted to selling approximately 1,000 units to his own company SpaceX to reduce the backlog.

Competitive Pressure & Pricing Power

BYD's Seal sedan offers comparable specifications to the Model 3 at roughly \$8,000 less, with more standard features included. BYD produces its own batteries, giving it a structural cost advantage Tesla cannot match. Tesla is now offering 0% APR financing in the United States, a subsidy that costs Tesla roughly \$2,000–\$3,000 per vehicle in foregone interest income and signals that demand at current prices cannot sustain volume without financial engineering. Tesla has effectively turned South Korea into a dumping ground for unsold inventory. On the last day of 2025, Tesla cut Model Y prices by over \$9,000, pushing Korean prices below what Chinese buyers pay for the same vehicle. These are not the pricing dynamics of a company with robust organic demand.

Compounding Cost Pressures

Tesla is now the most expensive new car to insure in the United States as of 2026. Model Y insurance rates are rising 2.9 times faster than the national average. Rising lithium and battery cell costs are adding thousands of dollars per vehicle to U.S. production costs at precisely the wrong moment. The cancellation of GHG credits will weigh on Tesla's \$2.7 billion regulatory credit revenue stream.

China's phased elimination of VAT export rebates is a compounding headwind. In December 2024, the rebate on vehicle and battery exports was cut from 13% to even 9%, meaningfully increasing the landed cost of every vehicle Tesla ships from Shanghai to Europe and other export markets. In January 2026, Beijing announced further cuts: the battery rebate dropped from 9% to 6% in April 2026 and is to be eliminated entirely by January 2027. This hits Tesla on two fronts, raising costs on Shanghai vehicle exports (roughly half of global deliveries) and increasing the price of LFP cells that Tesla imports from CATL for both vehicles and Megapack energy storage. The trajectory is one-directional: export incentives that once made Tesla's China operations highly profitable are being systematically removed.

3. Robotaxi & Optimus: Hopium vs. Reality

The market prices Tesla as an AI and robotics company, not an automaker. Wolfe Research projects \$30 billion in robotaxi revenue by 2030, rising to \$250 billion by 2035. ARK Invest argues that the robotaxi network could represent 90% of Tesla's enterprise value by 2029. Musk himself has stated that 80% of Tesla's future value will come from Optimus. Combined, these two product lines account for the vast majority of Tesla's \$1.5 trillion market cap. Neither has generated a single dollar of disclosed revenue.

Robotaxi: What Musk Says vs. What Exists

On the Q4 2025 earnings call, Musk claimed "well over 500 taxi vehicles carrying paid customers." As of Early March 2026, Robotaxi Tracker had identified approximately 445 unique vehicles by license plate, 356 in the Bay Area and 89 in Austin. The fleet has grown from roughly 200 in mid-January, which shows some ramp. But the headline number is misleading for two reasons.

First, the Bay Area fleet, which accounts for 80% of the total, is not a robotaxi service in any meaningful sense. Every Bay Area vehicle operates with a human driver at the wheel because Tesla has never applied for driverless testing or deployment permits from the California DMV, permits that Waymo, Cruise, and Zoox have obtained. The service is functionally indistinguishable from any typical ride-hailing service. Counting these as "robotaxis" inflates the number by 4x.

Second, even in Austin, the only city with any unsupervised operation, the 89 tracked vehicles do not translate into 89 operating robotaxis. Only 8 of the 89 vehicles were designated "unsupervised," and tracking data from early March shows that only 1 had been spotted operating in the prior week, with the remaining 7 last observed three to four weeks earlier. The functional unsupervised fleet on any given day appears to be one to two vehicles, not eight. This is the program that Musk said would have one million robotaxis on the road by 2020 and cover 50% of the U.S. population by the end of 2025.

The Crash Data

The Crash Record: 14 incidents, Zero Transparency

Tesla has reported 14 crashes involving its Austin robotaxi fleet to the NHTSA's Standing General Order database since the service launched in June 2025. Based on Tesla's Q4 earnings data showing approximately 700,000 cumulative paid miles through November 2025, the fleet likely surpassed 800,000 miles by mid-January 2026, equating to one crash every 57,000 miles.

We present this data with full context, because fairness strengthens the argument.

The 14 Incidents

The crashes break down as follows: five collisions with other vehicles, five with fixed objects, one with a cyclist, one with an animal, and two categorized as "other." Eight of the 14 occurred at speeds below 6 mph, with five of those between 0 and 2 mph, minor scrapes during low-speed maneuvers like parking and reversing. In multiple reports, the robotaxi was stationary when struck. One crash involved a city bus striking a stopped Tesla, which is almost certainly the bus driver's fault, not the autonomous system's. Another involved the Tesla being rear-ended by a truck at 4 mph.

Only three of the 14 involved the robotaxi proceeding forward and colliding with a moving object or fixed obstacle at meaningful speed. The most concerning was a collision with a fixed object at 17 mph while driving straight, that is difficult to explain away regardless of fault attribution. The cyclist collision and the animal strike at 27 mph also raise questions about object detection and avoidance that Tesla's redacted narratives make impossible to assess.

One additional detail emerged quietly: a crash from July 2025, originally filed with NHTSA as "property damage only," was reclassified five months later to "minor with hospitalization." Someone was injured seriously enough to require a hospital visit, and Tesla did not disclose this for five months. The pattern of delayed and upgraded crash reporting is itself under a separate NHTSA investigation.

The Transparency Problem

Here is what makes the crash data unusable for independent analysis: Tesla fully redacts every crash narrative in its NHTSA filings, citing "confidential business information." It is the only autonomous driving company in the entire NHTSA database to do this. Waymo provides detailed descriptions, vehicle speed, direction of travel, what the other party was doing, who initiated contact, and what the system was attempting. Zoox, Aurora, and Nuro do the same. Tesla provides nothing.

This means we cannot determine fault in any of the 14 crashes. We cannot assess whether the safety monitor intervened too late, whether the FSD system failed to detect an obstacle, whether the crashes were caused by other road users, or whether the system made correct decisions that still resulted in contact. Tesla wants the public to evaluate its safety record while simultaneously making that evaluation impossible. When the Austin Fire Department requested crash video footage to understand root causes, Tesla refused, citing "proprietary AI."

Crash Rate in Context

Even without fault attribution, the raw crash rate tells a story. By Tesla's own published benchmark, its Vehicle Safety Report states that the average American driver experiences a minor collision every 229,000 miles, the robotaxi fleet is crashing approximately four times more often. By NHTSA's police-reported crash data of one per 500,000 miles, the rate is roughly nine times worse. Using the most conservative human benchmark, approximately one crash per 200,000 miles, including unreported fender-benders, the fleet is still 3.5 times worse.

Operator / Benchmark	Miles per Crash	Human Supervision	vs. Tesla (raw)
Tesla Robotaxi (Austin, raw)	57,000	Safety monitor present	Baseline
Waymo (fully driverless)	98,000	None	1.7x safer
Human driver (Tesla's own benchmark)	229,000	Human	4x safer
Human driver (incl. unreported, NHTSA est.)	~200,000	Human	3.5x safer
<i>Tesla Robotaxi (fault-adjusted, ~5 of 14 attributable)</i>	<i>~160,000</i>	<i>Monitor present</i>	<i>Still 1.25x worse than human</i>
<i>Commercial insurance threshold (AV fleet)</i>	<i>200,000+</i>	<i>Unsupervised required</i>	<i>Raw: 3.5x gap; Fault-adj: 1.25x gap</i>

Some of these incidents were almost certainly not Tesla's fault. A crash with a bus while the robotaxi was stationary is not a FSD failure. Low-speed reversing incidents in parking lots may reflect sensor limitations rather than full system failures. A reasonable reading of the 14 incidents might attribute five to the autonomous system itself meaning even on a generous fault-adjusted basis, the fleet crashes

approximately once every 160,000 miles, still 1.25 times worse than the most charitable human benchmark and still below the commercial insurance threshold required for unsupervised deployment. But here is the deeper problem: we cannot know. Tesla redacts every crash narrative in the NHTSA database as confidential business information. Waymo, Zoox, Aurora, and Nuro all provide detailed incident descriptions. Tesla provides nothing. Fault attribution, intervention timing, whether the safety monitor had an opportunity to prevent the crash—all of it is hidden. Tesla is asking the market to price its robotaxi program at roughly \$1.2 trillion in implied value while making independent safety evaluation structurally impossible. That opacity is not a footnote. It is itself a material risk.

The Trend is Accelerating

Perhaps the most damning detail is the trend. Five of the 14 crashes, 36% of the total, occurred in a six-week window spanning December 2025 and January 2026, despite representing only a fraction of the total operational period. Early autonomous vehicle deployments typically show improvement over time as edge cases are identified and corrected. Tesla's crash rate is flat to worsening. This is the opposite of the learning curve the technology is supposed to produce.

The Experiment That Proves Supervision Matters

The reported crash rate also carries a hidden asterisk: safety monitors were present for nearly all of those 800,000 miles and are documented to have intervened at least once to prevent a traffic violation, as captured in footage published by independent tracker Joe Tegtmeier. An unknown number of additional interventions almost certainly occurred and were not publicly reported. The implication is that without monitors, the crash rate would be higher than one per 57,000 miles, possibly materially so.

The Austin-versus-Bay Area split provides a natural experiment. All 14 crashes occurred in Austin, where a safety monitor sat in the passenger seat for the first seven months of operation. In January 2026, Tesla removed the passenger seat monitors from 8 of its 42 Austin vehicles, but video evidence revealed the monitors had simply been relocated to trailing chase vehicles following each robotaxi. As of early March 2026, those chase vehicles appear to have been removed from the 8 unsupervised units, meaning the crash data accumulated under supervised conditions. We do not yet have a statistically meaningful sample from genuinely unsupervised operation. Zero crashes have been reported from the Bay Area fleet, where a human driver sits behind the wheel in the traditional driving position. Same FSD software. Same vehicle. Different seating position for the human.

The Waymo Comparison

Waymo now provides 450,000 weekly fully driverless rides across six U.S. cities with more than 2,500 vehicles, no safety monitor, no chase car, no human backup of any kind. The company has logged over 200 million fully driverless miles as of early 2026. A peer-reviewed study found Waymo reduces injury-causing crashes by 80% and serious-injury crashes by 91% compared to human drivers. Waymo averages approximately one crash per 98,000 miles, with zero human intervention capability.

Tesla's rate of one crash per 57,000 miles, with a trained safety monitor who can intervene at any moment, is worse than Waymo's rate without any human backup at all.

Waymo already provides roughly 20% of all Uber rides in Austin, in Tesla's own backyard. When Waymo vehicles were caught passing school buses, the company filed a voluntary recall within weeks and provided full transparency to regulators. Tesla, by contrast, has received two deadline extensions just to hand over its FSD traffic violation data to NHTSA, and redacts every crash narrative it files.

Musk's Timeline vs. Reality

Musk's Prediction	What Actually Happened
500 robotaxis in Austin by EOY 2025	~42 active in Austin by year-end
1,000+ in Bay Area by EOY 2025	~175, all with human drivers
8–10 metro areas by EOY 2025	2 cities (Austin + Bay Area)
"Half the U.S. population" covered	Parts of Austin only
Unsupervised rides by EOY 2025	8 designated unsupervised as of Feb. 2026
Cybercab production April 2026	No regulatory approvals obtained
"Full self-driving next year" (since 2016)	FSD remains Level 2

The Cybercab Economics & Profitability Timeline

The robotaxi bull case rests on a simple and genuinely attractive premise: remove the driver, who represents 60–70% of a ride-hailing operator's cost structure, and the unit economics of on-demand transportation become extraordinary. Musk has cited gross margins of 70%+. At a steady state and meaningful scale. The problem is the distance between today and steady state, and what it costs to get there.

Unit Economics at Scale

A robotaxi generates revenue by the mile. Community-sourced fare data from Robotaxi Tracker across 1,868 tracked trips shows Tesla's Austin service now priced at \$3.00 base plus \$1.40 per mile, equating to \$2.19 per mile effective at the 3.8-mile average trip length, below the \$2.50–\$4.00 range Waymo achieves with its flat per-mile pricing. Prior community-tracked fare data from 1,868 trips showed an average of \$2.07 per mile under the old pricing structure; total tracked revenue across that prior period was \$8,445. We model \$2.19 per mile as our revenue assumption; the cost stack at steady state looks like this:

Cost Item	Per Mile (Steady State)
Vehicle depreciation	\$0.15-0.18
Insurance	\$0.30–0.40
Maintenance & cleaning	\$0.15–0.25

Charging (electricity)	\$0.05–0.10
Compute & software	\$0.05–0.15
Remote operations & monitoring	\$0.10–0.20
Total Costs	\$0.80–1.28
Gross Margin at \$2.19/mile (current rate)	\$0.91-1.39 (42-63%)

Tesla has two genuine structural advantages here. Its Supercharger network is real infrastructure that materially reduces charging costs relative to a competitor building from scratch. And a \$30,000 target vehicle cost, if achieved, produces dramatically lower depreciation per mile than a Waymo Jaguar I-Pace at \$75,000+. The bull case is not fictional. It just requires a sequence of execution steps that Tesla has not demonstrated the ability to complete on any stated timeline.

Why the Steady-State Math Doesn't Apply Today

Insurance is priced off the safety record, and the safety record is bad. Tesla's Austin fleet is crashing once every 57,000 miles with a trained safety monitor present. Commercial insurers price autonomous vehicle coverage against demonstrated safety data. The market standard for competitively-priced AV insurance is approximately one crash per 200,000+ miles without human backup. Tesla is currently running at 3.5x that threshold. Until the gap closes through genuine improvements in FSD verified over millions of unsupervised miles—insurance alone could consume 20–25% of gross revenue, which structurally eliminates the margin advantage.

Manufacturing cost is nowhere near the target. Musk's sub-\$30,000 Cybercab retail price implies a steady-state manufacturing cost of approximately \$25,000. Early units rolling off the Giga Texas line are almost certainly costing \$40,000–\$50,000 each, roughly 1.5–2x the target, as the unboxed production process is debugged and volumes remain low. At \$40,000–\$50,000 per vehicle over a realistic 5–6 year urban operating life (~165,000–200,000 total miles), depreciation runs \$0.24–0.36 per revenue mile — roughly 1.5–2x the \$0.15–0.18 steady-state assumption in our model above. The unit economics don't approach their steady-state profile until Tesla reaches high-volume Cybercab production, which at current ramp pace implies 2027 at the earliest.

Utilization takes years to reach. The steady-state margin math assumes 60–70% revenue-generating utilization — a level consistent with Waymo's current mature fleet, which Waymo achieved after 6 years of development. Tesla's Austin fleet today operates at roughly 11% utilization. Getting from 11% to 66% requires network density, brand trust, and operational reliability that Tesla has not yet demonstrated in any market. In the interim, fixed costs per revenue mile are substantially higher than the steady-state model implies.

Network density is a prerequisite, not a byproduct. A robotaxi network only produces competitive wait times and therefore usable demand once a city has sufficient vehicle density. Waymo requires roughly 300–500 vehicles in a market the size of Austin to achieve sub-10-minute wait times reliably. Tesla has

8 designated unsupervised vehicles in Austin today. Reaching minimum viable density at target vehicle costs requires \$15–\$40 million per city before generating a dollar of recurring profit. Launching the 10+ cities the bull case requires is a \$150–\$400 million pre-revenue capital commitment.

The Sequential Gate Problem

Before Tesla generates a dollar of robotaxi profit, it must clear a series of binary gates in sequence. Each gate is a hard dependency; you cannot reach Gate 3 without having cleared Gate 2.

Gate 1: Regulatory authorization. The Trump administration has signaled intent to streamline AV regulations, including potentially removing the steering wheel requirement from FMVSS. Bulls cite this as resolving Gate 1. It doesn't, for three reasons: rulemaking takes 18–36 months regardless of political intent; Tesla skipped the NHTSA panel discussions where Waymo, Zoox, and Aurora were shaping the new rules; and even a streamlined standard still requires Tesla to certify safety performance — the same FSD data that is currently under active NHTSA investigation. The regulatory bar may move. Tesla still has to clear it.

Gate 2: Insurable safety record. Regulatory authorization is necessary but not sufficient. Commercial insurers set prices independently of regulators. Tesla needs millions of unsupervised miles at a crash rate approaching human benchmarks before insurance is priced at levels that allow attractive unit economics. This gate cannot be rushed—it requires time, miles, and demonstrated improvement that does not yet exist in the data.

Gate 3: Manufacturing cost at target. The unit economics require a vehicle cost near \$30,000. Getting there requires volume. Volume requires regulatory authorization and network buildout. This gate is downstream of Gates 1 and 2, and the cost curve does not bend until Tesla is producing tens of thousands of Cybercabs annually.

Gate 4: Network density. Profitable unit economics in one city require hundreds of vehicles on the road simultaneously. That requires Gate 3 to be well advanced. A single profitable city does not constitute a business, it constitutes a pilot.

The Honest Timeline: A Falsifiable Checkpoint

On January 28, 2026, Musk announced on the Q4 earnings call that Tesla would “expand Robotaxi to seven new cities in the first half of 2026.” The cities named were Dallas, Houston, Phoenix, Miami, Orlando, Tampa, and Las Vegas. The market interpreted this as a near-term revenue catalyst. The statement is deliberately ambiguous about what “expand” means operationally, and the timeline is mathematically impossible under any credible regulatory reading.

Three definitions of “expansion” are worth separating, because they lead to completely different investment conclusions.

Definition A: Regulatory approval for driverless operations. Self-certification is not a universal skeleton key across the seven cities. The state-by-state picture is materially worse than the headline implies.

Phoenix (Arizona): Tesla completed self-certification in September 2025 and received a TNC commercial permit in November—four days after applying, however, this only authorizes the commercial transaction (charging for a ride). It does not bypass the requirement for a safety driver.

Las Vegas (Nevada): Tesla obtained a testing permit in September 2025, but Nevada’s DMV confirmed that testing authorization is separate from commercial operations. Tesla still needs to complete a distinct “self-certification for operations” process and obtain Nevada Transportation Authority approval. Testing permit obtained; commercial launch not cleared.

Dallas and Houston (Texas): The Texas DMV’s AV permitting system does not go live until May 28, 2026—five weeks before H1 ends. Even under the most optimistic reading, Tesla applies May 28 and receives approval in 60–90 days minimum, putting first deployment in August or September at the earliest. Mathematically outside H1.

Miami, Orlando, and Tampa (Florida): Miami, Orlando, and Tampa (Florida): Florida is among the most permissive AV states in the country. A licensed human driver is not required, and fully driverless operation is explicitly legal under Florida Statute 316.85. Waymo is already operating commercially in Miami. Tesla's obstacle is not regulation, it is that the same safety record and Part 555 constraints follow the Cybercab into every market. Florida's permissive regime does not accelerate Tesla's timeline; it simply removes one excuse.

Realistic outcome by June 30, 2026: at most one city (Phoenix) with arguable regulatory clearance, and even their commercial launch requires operational infrastructure that does not yet exist.

Definition B: Limited test fleet with safety drivers or remote monitoring. Tesla could place 5–15 vehicles in a handful of cities under existing ride-hailing or testing licenses with human supervision in-vehicle or remote. This is operationally trivial—minimal revenue, no unsupervised miles logged—but technically satisfies the word “expand.” Realistic outcome by June 30, 2026: one to three cities, ten to twenty total vehicles under supervised operation. This is what the announcement will likely mean in practice.

Definition C: Revenue-generating unsupervised fleet at scale. This is what Musk implied and what the stock appears to be pricing. Multiple cities with active unsupervised fleets generating material revenue. This would require regulatory approval (impossible in six of seven cities by June 30), sufficient unsupervised safety data, operational infrastructure, and local permits in each market. As of March 12, 2026, Austin has one to two actively operating unsupervised vehicles. Realistic outcome by June 30, 2026: virtually impossible.

The market’s job by June 30 will be to decide which definition Musk meant. If the stock accepts Definition B as a win, Tesla avoids a missed deadline. If the market holds Musk to Definition C—which is what the valuation implies—it is the third consecutive year of broken autonomy timelines. Note that Waymo is already operating commercially in Miami, Dallas, Houston, and Orlando, the same cities Tesla is announcing as targets. Tesla is not pioneering new markets. It is announcing entry into cities where its primary competitor is already generating revenue. The definition mismatch and the Waymo overlap together make the H1 expansion promise one of the sharpest near-term catalysts in the thesis.

What Unit Economics Actually Look Like Year by Year

Grant Tesla every optimistic assumption: Texas deployment approval in Q3 2026 and insurance priced reasonably. What does the economics actually look like at the vehicle level?

In Year 1 (H2 '26 - H2 '27), Tesla deploys 100–200 unsupervised vehicles in Austin. At the current \$2.19/mile effective rate (\$3.00 base + \$1.40/mile at a 3.8-mile average trip), and assuming 12 trips per operating day—a conservative but realistic assumption for an early-stage fleet—each vehicle generates roughly \$100 in daily gross revenue. Against a realistic cost stack of approximately \$155 per vehicle per day, Year 1 is deeply loss-making at any deployment scale. Insurance alone runs \$60–100 per day for an unproven AV fleet, roughly 5–7x what a human rideshare driver pays, reflecting the absence of a verified safety record. Remote operations add another \$35 per day at a 1:15 operator-to-vehicle ratio. Depreciation on a \$40,000 vehicle over 60 months contributes \$22 per day, with charging and maintenance adding a further \$23. Total costs of \$155 per day against \$100 in revenue produces a loss of roughly \$55 per vehicle per day at 12 trips. To break even, Tesla needs approximately 19 trips per day per vehicle—a utilization level that Waymo, after 6 years of development, is only now achieving at mature fleet scale.

In Years 2–3 (H2 '27 - H2 '29), assume Tesla scales to 500–1,000 vehicles across three to five cities and manufacturing cost declines from \$40,000 toward the \$30,000 steady-state target. At \$30,000 per unit, depreciation falls to \$500/month—\$17/day—versus \$667 in Year 1. Insurance pricing improves modestly as a verified safety record begins to accumulate, dropping toward \$40–45/day. Remote operations costs decline to approximately \$20/day as tooling matures and operator ratios improve. At 15 trips per day, daily revenue reaches roughly \$125 against costs of approximately \$103, producing operating margins of 15–20%. That is real progress, but it is not the 70%+ gross margin figure that appears in analyst models and Musk presentations. It is a thin-margin transportation business at an early, fragile stage of its safety and regulatory track record—priced by the market as if it were a software platform.

In Years 4–5 (H2 '29 - H2 '31), with 5,000+ vehicles across ten or more cities and manufacturing cost approaching the \$25k target, per-vehicle margins finally reach 40%+. But by this point Tesla has consumed its \$20 billion capex budget, a significant portion of which funds multi-city deployment infrastructure, and \$1–2 billion in operating losses accumulated from the ramp period. The fleet is now profitable at the unit level. It is not yet material to a \$1.5 trillion market cap: 5,000 vehicles at mature utilization (~39 trips/day, consistent with Waymo's current fleet average) produces roughly \$592 million in annual revenue — still less than 0.1% of Tesla's current market cap implied revenue base. Material segment contribution, by any credible analysis, arrives in the 2029–2031 window. The market is not pricing 2029. The fleet-owned model described above assumes Tesla owns, operates, and depreciates every vehicle—carrying the full capital, insurance, and operational risk. The bull case, and the basis for ARK's \$250 billion revenue projections, assumes something Tesla has never confirmed operationally: a software take-rate on customer-owned Cybercabs, where Tesla collects 25–30% of ride revenue while bearing none of the vehicle cost, depreciation, or insurance exposure. That model produces 70%+ gross margins and requires no incremental fleet capex. The fleet-owned model produces 40-60% margins at

scale and requires continuous capital deployment. Tesla has not publicly committed to either structure exclusively. If the market is pricing take-rate economics and Tesla ends up owning and operating the fleet—as its current Austin operations suggest—the margin compression alone represents a multi-hundred-billion dollar valuation correction.

The \$1 Billion Benchmark

The \$592 million figure above is not cherry-picked to make the math look bad, it is derived directly from Waymo's own mature fleet utilization benchmarks applied to a 5,000-vehicle deployment. But \$592 million deserves a frame of reference. At Tesla's current share price, the implied enterprise value of the robotaxi business embedded in the stock is approximately \$300–500 billion, using sell-side SOTP models. To justify \$300 billion at a 30x revenue multiple, Tesla would need approximately \$10 billion in annual robotaxi revenue. At mature utilization, that requires roughly 85,000 Cybercabs operating simultaneously, a fleet 34 times the size of Waymo's entire current global operation. To justify even \$1 billion in annual revenue, one-tenth of what the valuation implies, Tesla needs approximately 8,500 vehicles at mature utilization (39 trips/day), or 16,500 vehicles if utilization remains near the break-even threshold of 20 trips/day.

For context, Waymo launched commercial driverless rides in October 2020 and had approximately 2,500 vehicles by early 2026 — a 5.5-year ramp. Tesla starts this clock with one to two actively operating unsupervised vehicles. Even assuming Tesla scales at three times Waymo's pace, reaching 8,500 vehicles takes until 2029–2030 at the earliest. Reaching 85,000 vehicles — the fleet implied by sell-side robotaxi valuations — does not happen this decade under any credible scenario.

The valuation gap is not a rounding error. The market is pricing a \$300–500 billion robotaxi business. Our year-by-year model, built on Waymo's own demonstrated unit economics, implies \$25–50 billion in probability-weighted option value.

Milestone	Bull Case	Base Case	Bear Case
Texas regulatory approval	Q3 2026	Q1 2027	Q3 2027
First unsupervised paid rides at scale	Q4 2026	Mid-2027	2028+
Unit economics positive (one city)	Q2 2027	2028	2031+
Five-city network operational	2028	2029	Not this decade
Robotaxi segment profitable	2029	2030	2033+
Material contribution to Tesla P&L	Q4 2029	2031	Mid-2030s

Waymo launched paid driverless rides in Phoenix in October 2020. It crossed 10,000 weekly rides in May 2023, 100,000 in August 2024, and 450,000 across six cities by December 2025 — a 5.5-year

ramp from first commercial ride to current scale. Waymo had a decade of regulatory relationships, safety data, and operational infrastructure before that clock started. Tesla's bull case prices profitability well before Tesla reaches Waymo's current position.

If Tesla launches genuine driverless paid rides in Austin in Q3 2026, the optimistic scenario, and follows Waymo's demonstrated ramp trajectory, it reaches Waymo's current scale around 2029. The Street's implicit modeling assumes meaningful robotaxi revenue by early 2027. The gap between that assumption and a grounded timeline is the single largest driver of our \$25–\$50 billion SOTP value for robotaxi versus the \$300–\$500 billion embedded in current sell-side price targets.

The Regulatory Bottleneck

There Is No Federal Level 4 Standard

The bull case assumes Tesla receives "Level 4 certification" and scales from there. The problem: no such certification exists at the federal level. NHTSA has not issued a performance standard defining what a vehicle must demonstrate to qualify as Level 4. What exists instead is a voluntary safety self-assessment framework that Tesla refuses to participate in, and a patchwork of state DMV permitting regimes each with their own requirements. NHTSA's AV STEP program, which was supposed to create a unified framework, remained in comment period as of March 2026. Tesla cannot certify to a standard that has not been written.

At the state level, the picture is no cleaner. Texas opens its DMV permitting system in May 2026 and requires "Level 4 certification" without defining it operationally; Tesla has not applied. California requires a separate driverless deployment permit distinct from the testing permit Tesla currently holds, and a December 2025 ALJ ruling that FSD is "unambiguously false" advertising creates a hostile regulatory backdrop. Every additional jurisdiction New York, Illinois, Massachusetts adds another approval process. There is no single switch Tesla can flip to go national.

The Legal Wall Tesla Can't Drive Through

The legal exposure facing Tesla's autonomous driving program is substantial, growing, and largely unpriced by the market. What follows is not speculative, it is drawn from court filings, regulatory rulings, and NHTSA investigation documents that are already public.

The California DMV Ruling: A Legal Landmark

In December 2025, a California administrative law judge ruled against Tesla in the CA DMV's three-year enforcement action over Autopilot and Full Self-Driving marketing. The judge's finding was categorical: the name "Full Self-Driving" is "actually, unambiguously false and counterfactual." That is not an analyst's opinion. It is language from a judicial ruling based on evidence submitted under oath.

The proposed remedy included a 30-day halt to Tesla's California sales and production operations, an outcome the DMV director ultimately stayed for 60 days to allow Tesla time to comply with corrective marketing requirements. But in doing so, he made the underlying finding explicit: the vehicles "could not at the time of those advertisements, and cannot now, operate as autonomous vehicles."

More consequential than the penalty is what Tesla's own counsel conceded during proceedings. Tesla's legal team stated on the record that "the currently enabled Autopilot and Full Self-Driving features require active driver supervision and do not make the vehicle autonomous." Statements made by a party's attorneys during litigation constitute judicial admissions, admissible as evidence in any subsequent lawsuit, in any jurisdiction. Tesla's lawyers have now said in open proceedings what the company spent years obscuring in its marketing. That concession travels.

The precedent implications are significant. California, the state where Tesla is headquartered, manufactures vehicles, and holds its AV testing permit, has now formally ruled that FSD does not meet the definition of autonomous. Any state DMV evaluating a Tesla driverless permit application can point directly to that finding. The major ride-hailing markets outside California, New York, Texas, Illinois, and Massachusetts among them, have their own consumer protection frameworks and their own regulators who can reach the same conclusion independently.

The NHTSA Investigation: March 9 Deadline

On October 7, 2025, NHTSA opened Preliminary Evaluation into FSD-related traffic safety violations, including vehicles running red lights, crossing into oncoming lanes, and behaving erratically near other road users. The investigation launched with 58 documented incidents and had grown to more than 80 by December 2025, drawing on driver complaints, Tesla's own filed reports, and media accounts. The probe covers approximately 2.88 million Tesla vehicles equipped with FSD.

NHTSA's data demand was extensive: raw event recorder files, CAN bus logs, onboard camera footage, software version histories, and records of system-generated driver warnings. Tesla missed the submission deadline twice, receiving extensions from January 19 to February 23, and then again to March 9. In requesting the final extension, Tesla described the task as "unduly burdensome," citing the need to manually review more than 8,000 individual records before it could comply.

The same FSD software stack under investigation is the software powering the Austin robotaxi service. This is not a separate system but the same code. If the submitted data reveals systematic failure patterns, the possible outcomes include a formal recall of FSD across 2.88 million vehicles, mandatory over-the-air software changes, or deployment restrictions that would directly undermine the robotaxi expansion timeline.

Separately, Tesla faces concurrent NHTSA probes into delayed crash reporting (Tesla reclassified a July 2025 crash to include a hospitalization five months after the incident) and inoperative door handles on certain vehicle models.

The contrast with Waymo is instructive. When Waymo vehicles were found to have passed stopped school buses, the company filed a voluntary recall within weeks and provided regulators with complete incident documentation. Tesla, facing a far more serious investigation, has required months of deadline extensions to produce its data, and continues to redact every crash narrative it submits to NHTSA.

Mounting Legal Liability

Tesla now faces active litigation across multiple jurisdictions, with exposure spanning consumer fraud, securities law, and product liability.

The California FSD fraud class action involves tens of thousands of plaintiffs alleging Tesla marketed FSD as capable of autonomous operation while knowing it required constant human supervision. Consumers paid between \$8,000 and \$15,000 for the feature. With Tesla's own lawyers having now conceded in open proceedings that FSD is not autonomous, the plaintiff bar has a significantly stronger evidentiary foundation than it did at filing. If the class expands to all U.S. FSD purchasers, total exposure could reach into the tens of billions.

A separate investor securities fraud class action (*Morand v. Tesla*) covers the period from April 2023 to June 2025, alleging that Tesla and Musk made materially misleading statements about FSD capabilities and the autonomous driving timeline, inflating the stock price.

A separate Australian class action involving thousands of plaintiffs pursues similar FSD marketing claims under Australian consumer law.

TeslaDeaths.com has documented 31 verified Autopilot and FSD-related fatalities in California since 2018, representing roughly half of the 62 total U.S. fatalities tracked. Each opens potential wrongful death liability. A jury awarded \$243 million against Tesla in August 2025 for a fatal 2019 Autopilot crash — the first federal jury verdict against Tesla on an Autopilot wrongful death case. Tesla rejected a \$60 million pre-trial settlement. In February 2026, a federal judge upheld the verdict in full. Since then, Tesla has settled at least four additional Autopilot crash lawsuits rather than face further verdicts. Tesla is expected to appeal to the Eleventh Circuit.

A significant population of North American Tesla owners who purchased FSD on older hardware generations face real questions about whether they will ever receive the autonomous capabilities they were promised. A mandatory hardware upgrade program across this installed base would represent a liability running into the billions.

The Fundamental Technical Barrier

All of this legal and regulatory pressure converges on a single technical fact: FSD is classified as SAE Level 2. This means it legally requires a fully attentive driver at all times. NHTSA's own investigation documents characterize it as Level 2. Tesla's California AV testing permit, held since 2014, is for Level 2 testing only, it requires a safety driver at all times and does not authorize driverless commercial operations.

Despite holding a California testing permit for over a decade, Tesla has never applied for the driverless testing or deployment authorizations that Waymo, Cruise, and Zoox have pursued and in some cases obtained. Tesla has also declined to submit a Voluntary Safety Self-Assessment to NHTSA, the standard disclosure mechanism that every credible autonomous vehicle program provides to demonstrate safety maturity. It similarly refuses to file disengagement reports with any state DMV, the primary data source regulators rely on to evaluate whether a given system is actually improving.

Without Level 4 certification, which requires demonstrating that a vehicle can operate safely without any human intervention within its defined operational domain, the Cybercab cannot legally operate commercially in any U.S. jurisdiction. The vehicle ships with no steering wheel and no pedals. If the software fails, a human cannot take over because no mechanism exists for them to do so. Every major U.S. regulatory framework requires either Level 4 certification or a human-operable redundancy system. The Cybercab satisfies neither condition, which means Tesla is not simply waiting on regulatory approval. It is operating in a legal framework that its flagship autonomous product was not designed to meet.

Cybercab Stranded Asset Risk

Tesla began pre-production of the Cybercab at Giga Texas in mid-February 2026, with the first unit rolling off the line on February 17. By early March, drone footage showed approximately 25 units on factory grounds, with reports of hundreds in various stages of build. Volume production remains targeted for April. Musk has warned the early ramp will be "agonizingly slow."

The fundamental problem remains unchanged, indeed, it has become concrete. Tesla is actively manufacturing a vehicle with no steering wheel or pedals that it cannot legally deploy without further regulatory authorization. Tesla holds a Transportation Network Company (TNC) ride-hailing license in Texas, valid through August 2026, which permits it to operate a ride-hailing service with or without safety drivers. However, this TNC license does not classify Tesla's vehicles as autonomous, that designation must come separately from the Texas Department of Motor Vehicles, which Tesla has not yet received.

Under Texas SB 2807, effective September 1, 2025, autonomous vehicle operators must demonstrate Level 4 certification, submit a Voluntary Safety Self-Assessment to NHTSA, and prove Minimal Risk Condition capability, none of which Tesla has completed. The DMV's new permitting computer system does not go live until May 28, 2026. In California, Tesla has never applied for the driverless testing or deployment permits that Waymo, Cruise, and Zoox have obtained, despite holding a Level 2 testing permit since 2014. Tesla's own chairwoman signaled the Cybercab may need a steering wheel, and test units have been spotted with steering wheels installed. Musk himself admitted Tesla needs roughly 10 billion miles of FSD data to achieve safe unsupervised driving, a threshold not expected until mid-2026 at the earliest, with training, validation, and debugging requiring additional time beyond that. The production line is being designed to produce hundreds of Cybercabs per week at scale—a rate that would exhaust NHTSA's Part 555 exemption cap of 2,500 non-compliant vehicles per year in a matter of weeks. Tesla has not filed for that exemption. Every Cybercab produced before autonomous vehicle authorization is inventory that cannot generate revenue as intended.

Optimus: The \$1.2 Trillion Phantom

Musk has publicly stated that "80%" of Tesla's future value will come from Optimus and related AI businesses. At Tesla's current market cap of approximately \$1.5 trillion, that implies roughly \$1.2 trillion attributed to the humanoid robot program. This section examines what \$1.2 trillion is actually buying.

The Production Miss Mirrors Robotaxi Exactly

The timeline of broken promises follows the same pattern as FSD and robotaxi: extraordinary claims, followed by dramatic misses, typically by 80–90%.

From 2021 through 2023, Tesla progressed from a dancer in a bodysuit at AI Day to walking prototypes—genuine engineering progress, but demonstration units in controlled conditions, not products.

The credibility break came at the 2024 “We, Robot” event. The Optimus robots appearing to mingle, serve drinks, and hold conversations were being operated by remote human teleoperation offstage—a fact Tesla did not disclose. This is not a minor distinction: it is the difference between a product and a prop.

In January 2025, Musk said Tesla's "normal internal plan calls for roughly 10,000 Optimus robots to be built this year." By March, the target was revised down to 5,000 during an all-hands meeting. In July, The Information reported that actual production was in the hundreds, less than 10% of the already-reduced target. The Optimus program head, Milan Kovac, who had led the project since 2022, departed in June 2025. Leadership of the program was absorbed by Ashok Elluswamy, Tesla's VP of AI, who also runs FSD and Autopilot, adding a third major program to his responsibilities.

In December 2025, at a Miami showcase meant to demonstrate Optimus capabilities, a robot fell backward on stage while attempting to hand out water bottles. The video went viral.

On the Q4 2025 earnings call in January 2026, Musk admitted that zero Optimus robots are currently doing "useful work." He called the program "still very much in the R&D phase." When asked directly how many robots Tesla actually has, he declined to answer. This is the CEO who attributed 80% of the company's future value to this product, refusing to disclose how many units exist.

What "Gen 3" Actually Means

Tesla plans to unveil "Gen 3" Optimus in Q1 2026. The name suggests a full generational upgrade. The reality is narrower: "Gen 3" refers specifically to upgraded hands with 22 degrees of freedom and 50 actuators (25 per forearm/hand), up from 11 degrees of freedom in Gen 2. The robot body remains the Gen 2 design, 173cm, 57kg. This is a meaningful hand upgrade, but calling it "Gen 3" creates the impression of a completely new robot when the body, legs, torso, and core mobility platform are unchanged.

Musk has acknowledged that Optimus relies on "a completely new supply chain" with "really nothing from the existing supply chain" used in Tesla's vehicles. Every component is built from scratch using bespoke suppliers. This is the opposite of the vertical integration advantage Tesla has in its auto business, where decades of supplier relationships and manufacturing scale drive costs down. For Optimus, the cost curve starts from zero.

Demonstrations Are Not Autonomy

The gap between what Tesla shows in videos and what the robot can actually do autonomously is significant.

Multiple demonstrations have been revealed to rely on teleoperation rather than autonomous AI. At the "We, Robot" event in 2024, robots appeared to interact naturally with attendees, but observers confirmed that human operators were controlling the movements remotely. Tesla was not transparent about this. The distinction matters: a teleoperated robot is a sophisticated puppet, an autonomous robot is a product.

What has been demonstrated autonomously operates reliably only in structured, staged settings, controlled lighting, known objects, pre-mapped environments, and limited failure modes. The October 2025 kung fu demonstration was confirmed as AI-driven and genuinely impressive as a balance and motion control exercise. Factory data-collection tasks reportedly run without human intervention. But these are narrow, repetitive tasks in environments specifically designed for the robot.

Unstructured environments, a random home, an unfamiliar warehouse, an outdoor construction site, represent an entirely different challenge. No demonstration has shown Optimus performing useful work autonomously in an unstructured environment. Consumer sales are targeted for the end of 2027, a timeline that has shifted multiple times and should be treated with extreme skepticism given the pattern documented above.

The Cost Gap Is Enormous

Current estimated unit costs range from \$50,000 to \$100,000. The Gen 3 hands alone are estimated at \$30,000–\$80,000 per set (unofficial figures, as Tesla has not disclosed component costs). Initial commercial units are expected to be priced at \$100,000–\$150,000.

Musk's target price is \$20,000–\$30,000, a figure he has referenced repeatedly as the eventual consumer price point. Reaching that target depends on manufacturing scale that does not exist, a supply chain that is being built from scratch, and cost reductions that have no precedent in robotics hardware. For context, Tesla's own automotive business, which has been manufacturing at scale for over a decade with mature supplier relationships, still carries a cost per vehicle above \$34,000. Achieving a lower cost on a far more complex product with a brand-new supply chain defies industrial experience.

China's recent rare earth export restrictions add a direct supply chain threat. Optimus relies on rare earth magnets for its movement systems, the actuators that power every joint. China controls approximately 90% of global rare earth magnet production and has already restricted exports as a trade policy lever. This is a single-source vulnerability for a critical component.

Competition Is Already Ahead on Cost and Deployment

Musk presents Optimus as if Tesla is the only serious player in humanoid robotics. The reality is that more than 300 humanoid robot companies now exist globally, with over 150 in China alone.

Unitree's G1 humanoid robot is available for under \$20,000, already below Musk's aspirational target price for Optimus, and is shipping to customers today. Boston Dynamics' Atlas is in pilot testing at Hyundai manufacturing facilities for 2026–2028 deployment, focused on practical industrial applications. Figure AI, backed by OpenAI, Nvidia, Microsoft, Intel, and Amazon, has raised over \$1.9 billion and is developing humanoid robots for warehouse and logistics applications. Chinese competitors including Unitree, UBTech, and Fourier Intelligence are focused on production reliability and cost optimization, leveraging China's manufacturing base and favorable rare earth access.

Goldman Sachs estimates the total global humanoid robot market at \$38 billion by 2035, with the most optimistic scenario reaching \$154 billion. Even the bullish estimate falls far short of the trillions Musk has implied. If Optimus captured 25% of Goldman's optimistic \$154 billion scenario, which would make it the dominant global player, that would be roughly \$38.4 billion in revenue. At a generous 20% operating margin, that is \$10 billion in operating profit. At 30x earnings, that is a \$300 billion business in 2035, at the most optimistic end of the range. The market is pricing Optimus at \$1.2 trillion today for a product still in the R&D phase.

Rodney Brooks, the cofounder of iRobot and one of the most respected roboticists in the world, has called the vision of humanoid robots as general-purpose assistants "pure fantasy thinking." Experts cited by Deutsche Welle described aspects of the Optimus project as a "complete and utter scam." Multiple robotics researchers have noted that what Optimus demonstrates was accomplished by competitors years earlier.

The Capital Commitment Is Staggering for an R&D Project

Tesla is converting the Fremont factory's Model S and Model X production lines to Optimus manufacturing, targeting eventual capacity of 1 million units per year. A separate 10-million-unit-per-year facility is planned at Giga Texas. The company is ending production of its oldest and most iconic vehicles, the cars that built the Tesla brand, to make room for a product that its own CEO describes as still in R&D.

These are not small bets. Retooling Fremont requires significant capital expenditure. Building a 10-million-unit facility at Giga Texas would be one of the largest manufacturing projects in U.S. history. And the supply chain for every component must be built from scratch, on a timeline that competes with hundreds of other companies for the same rare earth materials, actuator components, and sensor systems.

All of this is being committed before Optimus has a single external customer, a single dollar of revenue, or a single unit doing useful work outside a staged demonstration.

The \$1.2 Trillion Question

The implied valuation math is simple and devastating. If Musk is correct that 80% of Tesla's future value comes from Optimus, the market is currently assigning approximately \$1.2 trillion to a program that has: zero units doing useful work, zero external customers, zero disclosed revenue, production volumes in the hundreds against a target of 10,000, unit costs 3–5x above the target price, a departed program head,

demonstrations that relied on hidden teleoperation, a supply chain built from scratch with single-source rare earth exposure, and a CEO who admits it is "still very much in the R&D phase" and refused to disclose how many units exist when asked directly.

For \$1.2 trillion, you could buy Toyota, Volkswagen, GM, Ford, BMW, Mercedes-Benz, Hyundai, and Honda, and still have \$300 billion left over. Those companies collectively delivered over 45 million vehicles last year. Optimus delivered zero.

4. The Talent Exodus

The departures are concentrated in exactly the programs the market is paying the highest multiple for:

Andrej Karpathy, Director of AI, departed July 2022. Architect of FSD's vision-based neural network approach; the most credentialed AI researcher Tesla employed.

Zachary Kirkhorn, CFO, resigned August 2023. Managed every capital raise through the survival years; replaced by a CFO now disclosing active bank financing discussions.

Pete Bannon, Chip Architect (Dojo/D1), departed late 2023. Led custom AI inference chip design underpinning the autonomy stack.

Drew Baglino, SVP Powertrain & Energy Engineering, departed April 2024. 18-year veteran; the technical foundation of both the vehicle and energy businesses.

Tom Zhu, Head of Global Manufacturing, reassigned early 2024. Credited with executing the Shanghai ramp that rescued Tesla's 2022 production crisis.

David Lau, Head of Software, departed 2025.

Milan Kovac, Optimus Program Head, departed June 2025. Left while production was running below 10% of internal target; program absorbed by Ashok Elluswamy, who also runs FSD and Autopilot.

Victor Nechita, Cybercab Program Manager, departed early 2026, days after the first production unit rolled off the line at Giga Texas.

Thomas Dmytryk, Director of OTA & Robotaxi Infrastructure, departed March 2026 after 11 years. Built the software backbone of the Austin robotaxi service.

Sendil Palani, VP of Finance, departed March 2026 after 17 years. One of the last pre-IPO executives still in a senior role; joined in 2009 when Tesla had weeks of cash runway.

Tesla has also cycled through four global sales leaders in under two years. What remains is a more centralized organization with Musk at the apex simultaneously running Tesla, SpaceX, xAI, X, The Boring Company, Neuralink, and fewer experienced operators beneath him to execute the most complex buildout in the company's history.

5. Fundamental Cash Burn

Tesla is about to embark on its most aggressive spending cycle in company history, at the exact moment its core business is shrinking. The company generated \$6.2 billion in free cash flow in 2025. It is guiding for capex "in excess of \$20 billion" in 2026. The math does not work.

The \$20B+ Capex Bomb

CFO Taneja stated on the Q4 2025 earnings call that 2026 capex would be "in excess of \$20 billion", suggesting a 135% increase from \$8.5 billion in 2025 and the largest annual outlay in Tesla's history. The spending funds six new factories: a lithium refinery, LFP battery plant, Cybercab production line, Semi factory, Megafactory Houston, and Optimus production line, plus AI compute infrastructure, existing factory expansion, and robotaxi fleet growth. The \$20 billion figure does not include a solar cell fab or the proposed semiconductor "TeraFab" — and this is not a footnote. CFO Taneja stated explicitly on the Q4 call: "none of these numbers which I shared of \$20 billion factor in anything to do with the solar fab or the semiconductor chip fab — those would come later on." On March 14, 2026, Musk announced on X that the "Terafab Project launches in 7 days." Semiconductor fabs of this scale typically cost \$20–\$30 billion and take two to three years from groundbreaking to first wafers, with another one to two years to full yield. "Launch in 7 days" almost certainly means a groundbreaking or formal announcement, not production — but the capital commitment is real and imminent. The \$20 billion capex figure the market is already struggling to absorb is a floor, not a ceiling. Taneja acknowledged that longer-tail infrastructure projects will require "more debt or other means" to fund. TeraFab is exactly that project.

R&D on Top of Capex

R&D surged 41% to \$6.4 billion in 2025, an income statement expense not included in the capex figure. Total 2026 investment spend (capex plus R&D) will approach \$27–28 billion, against operating income of just \$4.8 billion. Tesla is spending roughly \$6.40 on investment for every \$1 it earns. This R&D is spread across more than ten simultaneous categories, with a CEO whose attention is divided across six companies.

Cash Engine Explosion

In 2025: \$14.7 billion OCF, \$8.5 billion capex, \$6.2 billion FCF. The CFO guided capex "in excess of \$20 billion" for 2026. At our OCF estimate of roughly \$12–13 billion, that produces negative free cash flow of approximately \$7–8 billion, the worst in Tesla's history by a wide margin, in a year when EPS is down 75% from its 2023 peak. Tesla could go from \$44 billion in cash to below \$20 billion within 18 months without external financing, which explains why the CFO is already in discussions with banks.

Cash Flow Quality Is Already Weak

The FY2025 10-K reveals that accounts payable contributed \$4.4 billion to operating cash flow, 30% of the \$14.7 billion headline. Tesla generated nearly a third of its cash by stretching payments to suppliers. This has accelerated each year: \$2.6 billion in 2023, \$3.6 billion in 2024, \$4.4 billion in 2025. Days payable outstanding swung wildly: 72 in Q1 (when deliveries collapsed to 337,000), 65 in Q2, 52 in Q3, 61 in Q4. Accrued liabilities grew \$2.6 billion (24%) while revenue declined 3%. Stock-based compensation added \$2.8 billion as a non-cash add-back, up 41% from 2024.

Net income was \$3.9 billion. Operating cash flow was \$14.7 billion. The \$10.9 billion gap is almost entirely non-cash add-backs and working capital timing. Strip out the payables growth and SBC, and the underlying cash generation looks far less comfortable against \$20 billion in planned capex.

Accounting Quality Risk: Deferred FSD Revenue as an Earnings Management Tool. Tesla carries \$3.83 billion in deferred FSD revenue on its balance sheet as of Q3 2025 — accumulated from customers who paid for Full Self-Driving capability over years of vehicle purchases. Tesla's own Q3 2025 filing discloses it expects to recognize \$880 million of this balance within the next 12 months. The mechanism has already been deployed: Q3 2024 operating income was boosted by one-time FSD revenue releases tied to Cybertruck and specific feature unlocks, and Q3 2025 results were explicitly described as weaker year-over-year due to the absence of those releases. Investors analyzing Q1 and Q2 2026 results should normalize for this effect. A \$500–880 million deferred revenue release in H1 2026 would mechanically boost reported automotive revenue and gross margin, potentially masking 200–300 basis points of organic margin deterioration in the core vehicle business. The release is real revenue under GAAP — but it is not recurring, it does not reflect improving unit economics, and it will not repeat in 2027. A headline EPS beat driven by deferred recognition is not validation of the bull case. It is a one-time accounting release from a balance sheet accumulated over nearly a decade.

A Capital Raise Is Coming

CFO Taneja disclosed that Tesla is "in discussions with banks" about financing. The "consistent stream of cash flow" from robotaxis he referenced does not exist. If Tesla raises equity at ~376x P/E, dilution would be significant. The last capital raise was in 2020, doing one now shatters the self-funding narrative.

Capital Allocation Discipline Is Broken

Tesla invested \$2 billion in xAI at a \$250 billion valuation after shareholders had already voted down a \$5 billion proposal to do exactly that. At the time of Tesla's seed investment, xAI was valued at \$0.7 billion — Musk had spent the intervening period diverting Tesla's AI engineers and Nvidia GPUs to xAI rather than investing Tesla capital when the price was low. A Delaware class action arising from the transaction is active. The xAI situation rhymes precisely with the SolarCity acquisition: a related-party transaction, a conflicted board that cannot credibly evaluate deals involving the CEO's other companies, financial projections that did not survive contact with reality, and Tesla shareholders left holding the bag. The pattern is not coincidence — it is the predictable output of a governance structure where the CEO's personal interests and Tesla's corporate interests are structurally misaligned. The misalignment has now become explicit. Tesla has received government clearance to convert its \$2 billion xAI investment into a sub-1% stake in SpaceX, following SpaceX's acquisition of xAI in a share swap. The strategic rationale has fundamentally changed: Tesla shareholders were told the xAI investment was essential to Tesla's AI roadmap. That capital is now a passive minority position in a rocket company. FTC filings show Musk selling his own SpaceX holdings to other investors — including Valor Equity Partners and DFJ Growth — while Tesla simultaneously acquires its stake. Tesla shareholder capital is being deployed into a

Musk-controlled private company in a transaction where Musk sits on both sides, ahead of a SpaceX IPO that would primarily benefit Musk personally. Tesla is guiding for negative \$7–8 billion in free cash flow in 2026, has disclosed bank financing discussions, and just redeployed \$2 billion of shareholder cash into a sub-1% stake in a private company its CEO controls.

6. Energy Growth: Real Business, Wrong Rescue

Tesla's energy business generated \$12.7 billion in revenue in 2025 (+27% year-over-year), with gross margins of approximately 30%, record deployments of 46.7 GWh, and \$4.96 billion in deferred revenue for 2026. It generates 23% of Tesla's total gross profit from just 13% of revenue. This is real, growing, and profitable. Any honest short thesis must grapple with it.

However, the CFO himself flagged margin compression ahead from three factors: tariffs on Chinese LFP cells jumping from 7.5% to 25%+ (with 99% of global LFP production in China and Tesla's domestic capacity covering only "a fraction" of needs); direct competition from CATL and BYD (Tesla's own cell supplier); and policy uncertainty.

Quantifying the Tariff Impact

Metric	2025 Actual	2026 Estimated
LFP Cell Tariff Rate	7.5%	25–54%+
LFP Cells as % of Megapack COGS	~40–50%	~40–50%
Implied Cell Cost Increase	,	+7–9 pp of COGS
Current Gross Margin	~30%	
Estimated Post-Tariff Margin	,	~21–23%
Gross Profit Impact (\$B)	\$3.8B	~\$2.7–3.0B

LFP cells represent approximately 40–50% of Megapack cost of goods sold. A tariff increase from 7.5% to 25% implies a 7–9 percentage point increase in cell costs. Against current gross margins of approximately 30%, this could compress margins to 21–23%, reducing annual gross profit from \$3.8 billion to approximately \$2.7–3.0 billion, a loss of \$800 million to \$1.1 billion in profit. If tariffs stack to 54%+ under the current trade framework, the impact would be significantly worse.

The Structural Math Problem

At 13% of revenue, the energy segment cannot offset a 70% automotive segment that is shrinking. If energy grows 30% in 2026, it adds approximately \$3.8 billion in revenue. Auto revenue alone declined \$7.5 billion in 2025. To fully offset one year of auto decline, energy would need to nearly triple. Energy is a real business with genuine growth. It is not a business that justifies a \$1.5 trillion market cap on its own, and it cannot offset the deterioration in auto.

7. Where We Diverge from Consensus

Street consensus on Tesla is not simply more optimistic than our view. It is optimistic in ways that are internally inconsistent, and in some cases contradictory. Below is a point-by-point breakdown of where the Macfarlane base case diverges most materially from the sell-side, and why we believe the Street is wrong on each dimension.

EPS: \$0.87 vs. Street \$2.25

The Street models 2026 EPS of approximately \$2.25, implying roughly a 2x recovery from 2025's \$1.08. Our estimate is \$0.87, a 61% discount to consensus. The gap is not driven by a single heroic assumption but by three compounding differences. First, the Street assumes ASP stabilization; we model continued pressure as Tesla defends volume against BYD in China and offers 0% financing in the U.S. Second, the Street embeds margin recovery; we see no credible path to margin expansion when capex is rising 135% and R&D is up 41% with no new volume product until the Cybercab, which has no regulatory authorization. Third, the Street treats the \$20 billion capex guidance as a ceiling; we treat it as a floor given the six simultaneous factory programs. On our numbers, auto operations are near breakeven and the positive EPS figure depends entirely on \$2 billion of below-the-line interest income.

Robotaxi: We Do Not Count the Bay Area Fleet

Sell-side models that project robotaxi revenue by 2026 or 2027 implicitly assume the Bay Area fleet counts toward the autonomous vehicle opportunity. It does not. Every Bay Area vehicle operates with a human driver because Tesla has no driverless permit in California and has never applied for one. The genuine robotaxi fleet, vehicles operating without a driver in any seat, numbers 8 in a single city. Wolfe Research's \$30 billion robotaxi revenue projection for 2030 requires a trajectory that is currently running approximately 5 years behind the pace needed. We assign zero revenue to robotaxi in our 2026 model and probability-weighted option value of \$25 to \$50 billion in our SOTP, versus the \$300 to \$500 billion embedded in current sell-side price targets.

Free Cash Flow: Negative \$7 to \$8 Billion vs. Street Near-Breakeven

Consensus models 2026 FCF near breakeven or modestly negative, anchored to the assumption that operating cash flow recovers as margins improve. We model OCF of \$12 to \$13 billion against capex in excess of \$20 billion, producing FCF of negative \$7 to \$8 billion. The Street is not wrong to expect OCF improvement; D&A, SBC, and working capital dynamics support a higher OCF than net income suggests. Where consensus goes wrong is on the capex assumption. The CFO said "in excess of \$20 billion." The Street models closer to \$20 billion as if that phrase means exactly \$20 billion. Six simultaneous factory programs with no precedent in Tesla's history do not stay at the guidance floor. We treat \$20 billion as the starting point.

The Multiple: 376x Is the Thesis, Not a Reason to Pass

The most common pushback from sell-side bulls is that valuation multiples are irrelevant because Tesla is a technology company being valued on optionality. We disagree with the framing. A 376x trailing P/E is not a neutral starting point from which optionality is additive. It is itself the thesis. It prices near-certain execution across robotaxi, Optimus, and energy simultaneously. Every leg of that thesis has to work. The crash data, the regulatory position, the talent exits, the production misses, and the FCF trajectory all point in the same direction: execution is not on track. When a stock is priced for perfection, evidence of imperfection is a catalyst, not a footnote.

Energy Margins: We Model Compression, Consensus Does Not

The sell-side broadly models the energy segment at sustained 28 to 30% gross margins through 2026, extrapolating the 2025 run rate. The CFO explicitly flagged margin compression ahead from LFP tariffs, CATL competition, and policy uncertainty on the Q4 call. We model margins compressing to 21 to 23%, costing approximately \$800 million to \$1.1 billion in gross profit. This is not a bear case assumption. It is arithmetic applied to the CFO's own stated headwinds. Consensus appears to be ignoring a margin warning delivered by management on a public earnings call.

8. Derating Catalyst Framework

Being right on fundamentals is not enough, you need a catalyst. The next six months contain an unusually dense cluster of binary events, any one of which could crack the narrative independently. If one or two legs weaken, the multiple could compress from ~376x to 100–150x, implying 40–60% downside.

Tier 1: Hard Catalysts (Date-Certain)

Now through April: Cybercab Stranded Capital. Pre-production started February 17 at Giga Texas, with ~25 units visible by early March and hundreds in various build stages. Volume production targets April. No steering wheel, no pedals, requires autonomous vehicle classification from the Texas DMV that Tesla has not received. However, Tesla could attempt to self-certify the Cybercab under existing FMVSS frameworks, bearing full liability for any safety failures that result. Tesla holds a ride-hailing license but not AV authorization, and the DMV's permitting system does not go live until May 28. Every unit produced is capital that cannot generate revenue. Musk warned the ramp will be "agonizingly slow." His own chairwoman suggested it may need a steering wheel after all.

March 9, 2026: NHTSA FSD Data Submission. Tesla must deliver crash data covering 80+ FSD violations across 2.88 million vehicles. Two prior deadlines missed. If the data reveals systematic failures, a formal recall or deployment restrictions become possible. The same FSD stack powers the Austin robotaxi service.

April 2026: Q1 Earnings. First quarter of the \$20B capex ramp. If FCF turns negative and bank financing is confirmed, this validates the cash burn thesis. On the autonomy side, the unsupervised fleet data is already damning ahead of the print: as of March 12, 2026, robotaxitracker.com shows 8 total

unsupervised vehicles ever deployed in Austin—Tesla’s only permissive regulatory market—of which 4 are formally deprecated (not spotted in 30+ days) and most of the remainder last seen weeks ago. Realistically, one or two vehicles appear to be actively operating unsupervised in the city today. Musk claimed the combined fleet was “well over 500” and doubling monthly on the Q4 call. If Q1 earnings arrive with no credible fleet count and no acceleration in Austin, the gap between the narrative and the reality becomes impossible to paper over.

May 28, 2026: Texas AV Permits Open. First date Tesla can apply for a driverless permit. Requires Level 4 certification, VSSA, and Minimal Risk Condition proof, none of which Tesla has completed. If Tesla is rejected or fails to apply, it is a tacit admission that FSD is not Level 4.

July 2026: Q2 Earnings. Q2 earnings arrive three weeks after the June 30 deadline on the seven-city H1 expansion promise—making this the first hard public accountability checkpoint. The framework is simple: Musk’s announcement can mean three things. Regulatory approval for driverless operations (Definition A) is mathematically impossible in most cities; Texas permits do not open until May 28, and state-level approvals run 60–90 days minimum. A handful of supervised test vehicles in 1–2 cities (Definition B) is achievable but operationally trivial—a rounding error on Waymo’s 450,000 weekly paid rides. Revenue-generating unsupervised fleets at scale (Definition C) is what the stock prices and what Musk implied—and is virtually impossible when Austin, the only permissive market, has 1–2 actively operating unsupervised vehicles today. By Q2 earnings, the market will know which definition it got. The stock is priced for C. The evidence points to B at best. That gap—arriving on a quarterly earnings call, in front of analysts, with Waymo already operating commercially in four of the seven announced cities—is a derating event.

Tier 2: Narrative Catalysts (Trend-Based)

European market share reports (monthly ACEA data, the 18%→9% collapse continuing). BYD quarterly results (outsold Tesla by 600K+ in 2025, gap widening). Waymo expansion (450K weekly rides, ~20% of Austin Uber rides). Energy margin compression (CFO warned; tariff impact quantified above). China demand deterioration (January domestic sales –45% year-over-year). Brand sentiment decline (DOGE backlash, declining consumer preference among core demographic).

Tier 3: Tail Risks (High Impact, Non-Trivial Probability)

Capital raise announcement (CFO disclosed bank discussions). FSD fatality in unsupervised mode (4x crash rate with monitors being removed). Multi-state regulatory crackdown (CA ruling could spread to NY, MA, NJ, and IL). China supply chain disruption (90% of rare earths, all LFP). Musk governance crisis (French court summons April 20, xAI conflicts, Delaware suit). Tesla/SpaceX merger speculation following SolarCity pattern.

The Derating Scenario

We do not need all catalysts to play out. If even one major leg of the bull case weakens, the multiple could compress from ~376x to 100–150x, implying 40–60% downside. The nearest hard catalyst was March 9. The next six months bring multiple shots on goal in a compressed timeframe.

9. Valuation Framework & Trade Structure

9a. Multiple Compression Scenarios

The simplest way to frame the opportunity: what happens if the market assigns a more realistic P/E multiple? The table below shows implied prices under various EPS and multiple assumptions.

EPS Basis	P/E Multiple	Implied Price	Market Cap	Downside
At 2025 EPS (\$1.08)	376x (Current)	\$406	\$1.5T	
	150x	\$162	\$580B	-60%
	100x	\$108	\$387B	-73%
	50x	\$54	\$193B	-87%
At 2026E Consensus (\$2.25)	150x	\$338	\$1.2T	-16%
	100x	\$225	\$808B	-44%
At Macfarlane 2026E (\$0.87)	150x	\$131	\$471B	-68%
	100x	\$87	\$314B	-79%

At 150x 2025 earnings, still a premium technology multiple, not an automaker multiple, the stock is worth approximately \$162, implying 60% downside. Even at 150x the optimistic Street consensus for 2026, the stock is worth \$338, implying 16% downside. At 150x our own Macfarlane 2026E estimate of \$0.87, the stock is worth \$131, implying 68% downside. At 100x our estimate, it is worth \$87, implying 79% downside. The asymmetry is striking: downside scenarios produce 40–80%+ declines, while upside from 376x is limited because expectations are already near-perfect.

9b. Sum-of-the-Parts Analysis

A more granular approach values each segment independently and assigns probability-weighted option value to speculative businesses. Even our "bull" scenario, which generously assigns \$100–200 billion to robotaxi and \$50–100 billion to Optimus despite neither generating revenue, produces a value well below the current market cap.

Segment	Method	Bear	Base	Bull
Automotive (70% rev)	15–20x op. inc.	\$40–55B	\$55–75B	\$75–100B
Energy Storage (13% rev)	25–35x op. inc.	\$25–40B	\$35–55B	\$50–75B
Services & Other (13% rev)	10–15x op. inc.	\$5–10B	\$8–15B	\$10–20B
Robotaxi Option Value	Probability-weighted	\$0	\$25–50B	\$100–200B

Segment	Method	Bear	Base	Bull
Optimus Option Value	Probability-weighted	\$0	\$10–25B	\$50–100B
Total Enterprise Value		\$70–105B	\$133–220B	\$285–495B
Per Share (3.75B shares)		\$19–28	\$35–59	\$76–132
Downside from ~\$406		–93 to –95%	–85 to –91%	–67 to –81%

The sum-of-the-parts reveals the core problem: Tesla's existing businesses (auto, energy, services) are worth \$70–195 billion depending on assumptions. The remaining \$1.3+ trillion of implied market cap is option value on robotaxi and Optimus programs with zero disclosed revenue, approximately 445 active robotaxi vehicles, and zero useful Optimus units. The market is pricing these options at near-certainty. We believe the probability-weighted value is a fraction of what is embedded in the current share price.

10. Why We Could Be Wrong

The strongest short theses are the ones that take the bull case seriously. Below are the risks that could invalidate this thesis, ranked by probability and impact.

Risk 1: Musk Has Defied Skeptics Before

This is the most important risk and the hardest to model. Tesla was months from bankruptcy in 2018. The Model 3 ramp was called impossible. SpaceX was told reusable rockets couldn't work. In each case, Musk delivered, late, over budget, and with more drama than necessary, but he delivered. The bear case on Tesla has been correct on fundamentals and wrong on the stock for most of the last decade. There is a non-trivial probability that FSD makes a genuine breakthrough, that Optimus finds a commercial application, or that Musk's sheer force of will produces an outcome the data does not currently support.

Why we're short anyway: Past execution was on products with clear demand (electric cars, rockets). The current bets, robotaxis in a patchwork regulatory environment, humanoid robots with no established commercial application, and a semiconductor fab, are fundamentally different. And the valuation has never been this disconnected from earnings: in 2020, Tesla traded at ~100x forward earnings during the Model 3 ramp. Today it trades at ~376x trailing earnings during an earnings collapse. The margin for error is categorically different.

Risk 2: Federal AV Preemption

If Congress passes federal autonomous vehicle legislation, or if the Trump administration issues an executive order preempting state-level AV regulation, the city-by-city permitting bottleneck disappears overnight. Given Musk's relationship with the current administration, this is not a negligible probability. The SAFE DRIVE Act is already in legislative discussion.

Why we're short anyway: The SAFE DRIVE Act faces a Senate filibuster. Even if preemption passes, Tesla still needs to demonstrate that FSD is safe enough to operate without a driver, and the crash data currently says it is 4x worse than humans. Federal preemption removes the regulatory barrier but not the safety barrier. And Waymo, which already has permits, would benefit equally or more from any federal framework.

Risk 3: Energy Storage Acceleration

If AI data center demand drives electricity consumption faster than projected, Tesla's Megapack business could grow 50–70% annually rather than 25–30%. At those growth rates, energy could become 20–25% of revenue within two years, partially offsetting the auto decline.

Why we're short anyway: Even at 50% growth, energy adds approximately \$6.4 billion in revenue in 2026, still less than the \$7.5 billion auto revenue decline in 2025. And margins are compressing due to tariffs on the very cells that make Megapack profitable. Energy is a good business. It is not a \$1.5 trillion business.

Risk 4: The Stock Can Stay Irrational

Tesla's valuation has defied gravity for years. The stock rose 63% in 2024 while EPS fell 53%. Narrative-driven momentum, retail enthusiasm, and options market dynamics (gamma squeezes) can push the stock higher regardless of fundamentals. Position sizing must account for the possibility that the thesis is correct but the trade loses money for 6–12 months before working.

Why we're short anyway: The catalyst density over the next six months is unusually high. We are not betting on a slow grind toward fair value, we are betting on specific dated events that force the market to confront the gap between narrative and reality.

Risk 5: Trump Administration Intervention

Given Musk's role in DOGE and his relationship with the administration, direct government support for Tesla through regulatory fast-tracking, favorable trade policy, or subsidies is possible. This is a political risk that is difficult to handicap.

Why we're short anyway: Government support can change the regulatory timeline but cannot change the technology readiness. FSD still crashes 4x more often than human drivers. And political favor is a double-edged sword, Musk's DOGE involvement has already damaged Tesla's brand among its core environmental buyer demographic.

11. What Would Change Our Mind

The following conditions, if met, would cause us to reduce or exit the short position. These are falsifiable, verifiable triggers, not subjective judgments. We will review each quarterly.

1	Unsupervised FSD achieves a crash rate at or below human drivers (1 per 229,000 miles by Tesla's own benchmark) over a sustained sample of 5+ million miles. This would indicate the technology is genuinely approaching Level 4 safety.
2	Cybercab receives regulatory approval for unsupervised commercial deployment in any U.S. jurisdiction (not just testing). This would break the regulatory bottleneck argument.
3	Q1 or Q2 2026 free cash flow is positive despite the capex ramp, indicating that Tesla's operating cash flow is strong enough to self-fund the transformation without external financing.
4	FSD subscription take rate exceeds 25% of the global fleet (vs. 12.4% today), demonstrating genuine consumer demand for the software at scale.
5	Optimus demonstrates autonomous task completion in an unstructured environment (not a staged demo), verified by an independent third party, with a credible path to unit economics below \$50,000.
6	Energy storage gross margins remain at or above 28% through H1 2026 despite tariff headwinds, indicating Tesla has found mitigation strategies the CFO did not flag.

None of these conditions are currently met. If any are met, we will reassess the position size and thesis conviction in real time. This is not a permanent short, it is a thesis with a defined shelf life and specific exit conditions.

Source Notes

Primary Sources: Tesla Q4 2025 Earnings Call & Shareholder Update (January 28, 2026); Tesla FY2025 10-K Annual Report (SEC filing); Tesla quarterly 8-K filings (Q1–Q4 2025); NHTSA Standing General Order crash database; NHTSA PE25012 investigation filings.

Tracking & Crowdsourced Data: Robotaxi Tracker (robotaxitracker.com), independent crowdsourced fleet data; Electrek robotaxi and crash analysis (Fred Lambert); Ethan McKanna / Texas A&M reverse-engineering of Tesla robotaxi app.

Industry Data: Kelley Blue Book (U.S. EV sales); ACEA (European auto registrations); CPCA (China passenger car data); UBS consumer surveys; Goldman Sachs humanoid robot market sizing; Wolfe Research (robotaxi valuation model, Emmanuel Rosner); The Information (Optimus production report, July 2025).

Analytical Note: The TeraFab cost estimate of ~\$20-30 billion is extrapolated from comparable U.S. semiconductor fab costs (Intel, TSMC) and is not a Tesla disclosure. The xAI valuation comparison references the 2023 seed round (\$0.7B) against the January 2026 Series E (\$250B), while technically accurate, seed-to-late-stage markups of this magnitude are common in venture capital. The governance

concern is that Tesla shareholders voted against investing at the higher price, and Musk proceeded anyway.

END OF MEMO