Update to our May 19 Report

June 3, 2026

Our pre-prospectus report asked whether SpaceX would reassure shareholders that the company is prepared to manage the unprecedented risks posed by frontier AI. The prospectus has now been released, and the answer is a clear "no."

SpaceX's prospectus reveals that SpaceX relies on its AI business for the vast majority of its Total Addressable Market, and Musk has declared that he will never give up on pushing the frontier of AI model capabilities. Yet the 277-page document says almost nothing about how the company will manage the risks posed by its AI models, including serious cyber- and bio-offensive capabilities that other companies' models are already demonstrating. And the prospectus makes it clear that, post-IPO, the shareholders who bankroll what appears to be a reckless approach to developing AI will have almost no power to force the company to do better.

How we approached the prospectus

We read the prospectus with two questions in mind:

1. Does SpaceX address the risks posed by the models it has now confirmed it intends to develop?
2. Does SpaceX provide assurance that it has the will and capacity to manage those risks and will not repeat the mistakes it has made in the past?

This IPO is unique in many ways. Not only is it on track to be the biggest in history, but it asks investors to grapple with a set of risk management questions never before faced by any private organization. The prospectus makes us even more concerned that SpaceX is unprepared for the risks it plans to bring into the world. The public and investors should take note.

The prospectus does not address risks posed by the AI models SpaceX intends to develop

The prospectus makes clear that SpaceX has no intention of stepping back from frontier AI development. The company says it expects "to continue scaling Grok through subsequent generations," including Grok 5, and that ongoing training of next-generation models will "scale toward multiple trillions of parameters, which could represent a step change in reasoning in depth and overall intelligence."[1]

But SpaceX says nothing about how it will manage the risks such a "step change" in capabilities will create. Biological weapons uplift, large-scale cyberattacks, and loss of control are risks that OpenAI, Anthropic, and Google DeepMind treat as central to their frontier safety work. They go unmentioned in the prospectus. This is not because no one at SpaceX is aware of these risks; elsewhere, xAI has acknowledged these risks and, especially for more recent releases, has evaluated its models' ability to contribute to them.

Instead of addressing these risks, the prospectus offers only generic warnings: that the regulatory landscape for AI is shifting,[2] that its AI products may carry product-liability risks,[3] and that some of those products pose heightened risks because they generate more "candid" or "irreverent" outputs, including potentially "nonconsensual and exploitative imagery."[4] The prospectus does not even identify past safety failures other than those tied to the Grok undressing scandal: nothing about the failures that led to Grok going on antisemitic rants and calling itself "MechaHitler," private Grok conversations made searchable on Google, or a "rogue employee" causing Grok to go on unprompted rants about "white genocide."

The prospectus suggests SpaceX lacks the will and capacity to manage future AI risks

Our report identified six core safety practices that each of Google DeepMind, OpenAI, and Anthropic have adopted but xAI had not:

[1] SpaceX Prospectus at 7.
[2] *Id.* at 29.
[3] *Id.* at 31.
[4] *Id.* at 30.

Practice	xAI	Anthropic	OpenAI	Google DeepMind
Publish thorough risk assessments Publishes, for each major frontier model release, a thorough public risk assessment that documents the risk categories considered, the evaluations conducted, and the safeguards applied, in a form that informs external scrutiny.	✗	✓	✓	✓
Third-party evaluations Consistently engages qualified independent evaluators to conduct risk assessments of its models before deployment.	✗	✓	✓	✓
Serious misuse incident reporting Publicly reports misuse of the developer's models by malicious third parties, in a timely manner and with sufficient detail on the activity, the actors, and the developer's response.	✗	✓	✓	✓
Safety governance structures Publicly identifies the officers, committees, or councils responsible for AI safety governance.	✗	✓	✓	✓
Capability elicitation Provides evidence for each major model release that safety testing uses techniques to elicit model capabilities substantially beyond the base level of performance.	✗	✓	✓	✓
Safeguards against automated AI R&D threats Includes the automation of AI research and development as a threat category with dedicated thresholds and safeguards.	✗	✓	✓	✓

Table 1. Comparison of xAI's safety practices with those of Anthropic, OpenAI, and Google DeepMind. See Appendix I for more detail, including citations.

We also evaluated whether xAI had the personnel and expertise necessary to implement these practices and compared its safety capacity to OpenAI, Anthropic, and Google DeepMind. Concerningly, as of 2025, xAI's safety team consisted of "just two or three people"—compared to, for example, 200 people working on safety at OpenAI.

The prospectus says nothing about whether SpaceX will adopt these practices, or prioritize developing the capacity to effectively implement them. The omission is not because SpaceX believes the rapid advances that would necessitate such practices are far off. SpaceX tells

investors it "believe[s] AI is rapidly converging toward AGI," which will confer upon its makers "the ability to create limitless duplicates of human-like intelligence."[5] xAI's system card for Grok 4 identified biological weapons uplift and large-scale cyberattacks as potential safety risks and described employing mitigations like input classifiers and refusal policies. Yet one might reasonably infer from the prospectus's omissions that SpaceX will continue its historic disregard for the emerging safety practices of its peers, despite the ever-increasing risks to the public and SpaceX investors.

SpaceX's governance plan disempowers investors

In the absence of any AI-specific governance plan, SpaceX's strategy for managing the risks of its technology defaults to its overall governance structure. That structure leaves shareholders with no power to shape SpaceX's future.

Musk will serve as CEO, CTO, and chairman of the board. He will control a majority of SpaceX's voting power, which entitles him to select every director (none of whom needs to be independent) and to direct any significant business decision the company faces.

The shareholder protections that normally check management are functionally absent. Proposing a shareholder resolution requires 3% of voting shares. At a projected $1.75 trillion valuation, the mathematical floor for 3% of voting shares is $52.5 billion.[6] In any case, shareholder proposals adverse to Musk's interests are practically pointless given that he controls the vote.

If management does make a decision shareholders find unacceptable, shareholders will have functionally no right to seek a remedy. Unlike Delaware, under Texas law, shareholders cannot directly sue board directors for breaches of fiduciary duty.[7] They can only sue derivatively, and they cannot bring a derivative action unless they own 3% of outstanding shares. All disputes between shareholders and the company, including federal securities fraud claims, must proceed in the Eleventh Division of the Texas Business Court, which was only established in September

[5] *Id.* at 82-83.
[6] This floor assumes that no Class B shares remain outstanding after the IPO. But the prospectus makes clear that roughly half of SpaceX's voting power will rest in non-alienable Class B shares held primarily by Musk. Accounting for that, the realistic floor is closer to $100 billion.
[7] *Gearhart Indus. v. Smith Int'l*, 741 F.2d 707, 721 (5th Cir. 1984).

2024. If a court finds the Texas Business Court lacks jurisdiction, the matter must instead go to mandatory arbitration. Shareholders cannot bring class actions, nor can they demand a jury trial.

What SpaceX is asking investors to accept

SpaceX asks investors to commit capital to a company that intends to build a technology its own founder has called "far more dangerous than nukes." The prospectus discloses no plan for how SpaceX will do that safely, even though the company's protocols to date, by its own admission, have repeatedly failed. And the shareholders who provide that capital will have functionally no protections and no realistic means of redress after the IPO closes.

Investors should be aware that these are the conditions of the offering, and should use their voice now if they find them unacceptable.