

The Wishes Were Capitalised Before the Genie Appeared

The genie has not appeared. The shareholders have. SSI's product is still undefined, but the institution that may control it is already being capitalised.

DRAFT PREVIEW DENNIS HEDEGREEN ANALYSE DRAFT-LOCAL

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-WISHES-WERE-CAPITALISED-BEFORE-THE-GENIE-APPEARED/](https://hedegreenresearch.com/articles/the-wishes-were-capitalised-before-the-genie-appeared/)

Safe Superintelligence says it has one goal and one product. Investors already own part of the company that may create it, even though SSI has not publicly explained what form the product will take, who will receive it, or whether it should ever be sold.

The lamp is closed.

The genie has not appeared.

SSI has not publicly described what the system does, who may use it, or how it would be delivered, priced, governed or withdrawn.

The company already has shareholders.

Safe Superintelligence Inc. (<https://ssi.inc/>), says it has one goal and one product: safe superintelligence.

That should make the company unusually easy to understand.

It does the opposite.

The Auction of the Unopened Lamp

“What is the product?”

“Safe superintelligence.”

“What form does it take?”

“That remains confidential.”

“Who is the customer?”

“That remains to be determined.”

Strictly speaking, nobody has sold the genie.

Investors have bought equity in the lamp: shares in the institution that may one day create, control and derive value from the system. They have not purchased a finished model or a defined service. SSI says its singular focus means no distraction from product cycles and insulation from short-term commercial pressure.

This is not inherently irrational.

Research is often financed before its outcome is known. Biotechnology companies raise capital before a medicine is approved, sometimes before anyone knows whether the mechanism will work.

SSI adds an unusual condition.

The public does not yet know what category of product success would produce.

The Scientist Objects

“We are not manufacturing wishes,” says the scientist. “We are pursuing a technical research programme.”

Fair correction.

A laboratory should not be required to release a half-finished system every six months to satisfy a growth chart. Product pressure can distort scientific priorities, especially when the stated objective is to make safety advance alongside capability.

The problem is not that SSI has failed to launch an app.

The problem is that ownership, infrastructure and strategic relationships are developing before the product form is publicly defined.

Research can remain secret.

The institution around it cannot remain undefined forever.

There is a limit to this criticism. SSI may not yet know what form success would take. The research itself may make the category impossible to define in advance. That would explain the missing product answer. It would not answer who gets to decide once an answer becomes possible.

Research Worthy of Scaling

On 27 July 2026, SSI and NVIDIA announced a long-term strategic partnership (<https://investor.nvidia.com/news/press-release-details/2026/Ilya-Sutskevers-Safe-Superintelligence-Inc--and-NVIDIA-Announce-Long-Term-Strategic-Partnership/default.aspx>). NVIDIA invested in SSI. The announcement said Vera Rubin access would allow SSI to increase its compute by an order of magnitude, and that the companies would collaborate on current and future computing platforms.

Ilya Sutskever said SSI had research “worthy of scaling up.”

That is a serious signal.

Sutskever is not an unknown founder selling a speculative deck. He co-authored the AlexNet paper (<https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>), and later led the sequence-to-sequence learning paper (<https://proceedings.neurips.cc/paper/5346-sequence-to-sequence-learning-with-neural-networks.pdf>). Confidence in his technical judgement is therefore understandable.

But confidence is not verification.

“Worthy of scaling” can mean that a result is promising enough to justify a larger experiment. It does not publicly establish AGI, superintelligence, a new architecture, a complete safety method or a commercially usable system.

Scaling is the next experiment.

It is not the coronation.

The Product Manager Enters

“Is it a model?”

“Perhaps.”

“An API?”

“Possibly.”

“Critical infrastructure?”

“That may depend on what happens when it works.”

A conventional product has a recognisable relationship between producer, customer and output. SSI may be building something that could occupy several positions at once.

It could become a model whose weights are licensed. It could become an API. It could remain inside SSI as a private research engine producing software, scientific discoveries, designs and patents. It could become a secure service for selected companies, universities or governments. It could remain bundled with the company, so the model, compute, security procedures, researchers and legal control never become separate products.

These are scenarios, not disclosed plans.

In the final scenario, the product is not the intelligence alone.

The product is the institution containing it.

The Investor Asks a Reasonable Question

“What do we own?”

“A company pursuing safe superintelligence.”

“And if the system is too dangerous to sell?”

“Then we must not sell it.”

“And if it can safely produce medicines, software and machines?”

“Then we will need a governance model.”

The investor smiles.

“Before or after the medicines, software and machines?”

SSI would not necessarily have to sell a successful system to commercialise its outputs. It could keep the system private while developing medicines, materials, software or intellectual property.

The genie would never need to leave the lamp.

The company could sell the wishes.

That may be safer than broadly releasing the underlying system. It could also place unusually general productive capacity inside a small private institution while outsiders receive only selected outputs.

The investors have not bought a known product.

They have bought shares in a company that may eventually control the production of many products.

Safety Is Not Governance

Technical safety matters enormously.

A powerful system should be reliable, controllable, resistant to misuse and responsive to legitimate shutdown authority.

But a technically aligned system can still exist inside an institution whose ownership and decision structure concentrates power.

Alignment asks whether the system follows the intended direction.

Governance asks who gets to choose the direction.

Neither SSI's public mission page (<https://ssi.inc/>), nor the partnership announcement (<https://investor.nvidia.com/news/press-release-details/2026/Ilya-Sutskevers-Safe-Superintelligence-Inc--and-NVIDIA-Announce-Long-Term-Strategic-Partnership/default.aspx>), describes who would approve deployment, what independent review would exist, how investor rights would interact with a safety decision or who could order a shutdown. That is a statement about those reviewed documents, not proof that SSI has no internal answers.

SSI does not need to reveal dangerous technical details to begin explaining what kind of institution it intends to become if the research succeeds.

The Bunker Is Not a Product Launch

Karen Hao reported a 2023 OpenAI discussion

(<https://www.theatlantic.com/technology/archive/2025/05/karen-hao-empire-of-ai-excerpt/682798/>), in which Sutskever spoke about building a bunker before releasing AGI. The account concerns OpenAI, not SSI, and Sutskever declined to comment. Its relevance is narrower: it frames AGI not as an ordinary release, but as an event likely to trigger government interest and protection questions.

If that mental model carries forward, asking when SSI will “ship” may be the wrong question.

The more relevant question is who will be in the room when the lamp is opened.

The Electricity Meter Interrupts

Superintelligence is an abstract noun.

Vera Rubin systems are not.

NVIDIA describes Vera Rubin NVL72 as a rack-scale system integrating compute, networking, liquid cooling and power management (<https://developer.nvidia.com/blog/inside-nvidia-rubin-gpu-architecture-powering-the-era-of-agentic-ai/>).

An order-of-magnitude increase in compute does not automatically mean the same increase in electricity use. The relationship depends on the baseline hardware, new configuration, utilisation, efficiency and workload.

The partnership announcement gives no SSI baseline compute, expected power envelope, infrastructure location or cooling arrangement. That prevents a responsible numerical energy claim.

It does not make the physical layer disappear.

Before superintelligence becomes abstract power, it arrives as electrical load, equipment and contracts.

NVIDIA is also an investor and technical partner. The joint announcement says it obtained “rare access” to SSI’s closely guarded research.

That does not prove SSI has lost its independence.

It makes the definition of independence worth asking for.

The Saved Hours Accountant Opens the Ledger

“The machine completed ten years of work before lunch.”

“Excellent.”

“Where should I record the saved time?”

“Under productivity.”

“That is not an account.”

If SSI succeeds on its own terms, its system may reduce the human time required for software engineering, scientific research, design, planning or invention.

Those savings could reduce prices, accelerate medical research, create public knowledge, increase investor returns, eliminate jobs, shorten working time or raise output expectations until everyone works just as much as before.

Technical capability does not choose between those outcomes.

Ownership and institutions do.

The Saved Hours question is not whether the machine performs useful work. It is what happens to the human time no longer required for the same output.

Safety alone does not tell us who receives it.

The Genie Speaks

One question remains premature. SSI has not publicly claimed an autonomous or conscious entity. But if a future system becomes an agent rather than a tool, ownership itself becomes unstable. The question would no longer be only who owns the wishes, but whether anyone should own the genie.

What a Readable SSI Could Explain

SSI does not need to publish its secret research today.

It could still explain:

- what product forms it is considering - who has authority to approve deployment - what independent review exists - how investor rights interact with safety decisions - what conditions would trigger government involvement - whether a successful system would be sold, licensed or used internally - whether energy use will be disclosed - who owns discoveries produced by the system - whether benefit distribution is considered part of safety

These questions do not request a dangerous technical recipe.

They ask what kind of institution is being constructed around it.

The Room Becomes Quiet

Perhaps SSI has found a new principle.

Perhaps NVIDIA has seen enough to justify a much larger experiment.

Perhaps the result will be important without becoming the destination printed on the company's name.

We do not know.

That uncertainty is not an accusation.

It is the reason the institutional questions must begin before the technical answer arrives.

Who can open the lamp?

Who chooses the wishes?

Who pays for the power?

Who owns what it produces?

And when the machine saves human time, where are the hours recorded?

— Dennis Hedegreen, trying to see the structure

RELATION MEMORY

NEARBY [The Missing Mechanism](#) · Both ask what institutional mechanism exists beyond recognition and warning.

EXTENDS [The Saved Hour Also Needs Power](#) · Extends the time-and-power ledger into ownership and control of productive capacity.

NEARBY [The Electricity Under AI Labour](#) · Both keep the physical compute layer visible beneath abstract AI language.

CANONICAL URL: [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-WISHES-WERE-CAPITALISED-BEFORE-THE-GENIE-APPEARED/](https://hedegreenresearch.com/articles/the-wishes-were-capitalised-before-the-genie-appeared/)

PDF VIEW GENERATED: 2026-07-27T22:46:28.953Z