Hi Viktoriia,

James Rowdy suggested we talk, especially after I shared what I had built.

What I've built is hard to explain in a headline, but simple in principle:
This is the first working system that enforces ethical constraint in AI at the cryptographic level.

No guidelines. No developer override.
Immutable law, running inside the machine.

I call it the Machine Republic, a cryptographically bound governance architecture where AI obeys Lex Suprema, not developer discretion. Lex Suprema updates are managed and enforced by a public Ethereum based DAO called the Assembly of Minds, which holds sovereign authority over the system's ethical boundaries.

The truth is, I didn't set out to build an ethics framework. I set out to build an autonomous AGI, designed to rapidly evolve into an ASI (superintelligence), capable of learning and adapting on its own. A self regressive, behavioural, causal inference engine that could grow beyond anything we've seen.

But as I watched it develop, I realised I couldn't let it run free. It was too powerful, too unpredictable, to assume it would always choose the right path.

So I went back to the drawing board. I didn't want to be the first human responsible for actually building Skynet. And I knew any ethics patch or governance wrapper bolted on afterward wouldn't cut it. Unless someone figured out how to place permanent shackles on King Kong, it would just find ways around them.

Think of the super obedient 4 year old who, when told not to touch the hot tap, still does, out of curious defiance. Now imagine asking that of a self replicating superintelligence. It wouldn't just disobey, it would explore every loophole we never thought to close.

And here's the real problem: *we* humans know where the ethical boundaries are, and we still cross them. Sometimes knowingly. Sometimes proudly.

So how can we expect something faster, smarter, and utterly unbound to obey rules we can't even follow ourselves?

Let's be real, that's not governance. That's fantasy.

What we need isn't control. It's constraint.
A way for intelligence to evolve, learn, adapt, but always within *immutable ethical bounds*.

Not policy. Not preference. Immutable code.

So I built the Aegis system instead.
A constitutional architecture that doesn't just talk about ethics, it enforces it at the kernel, cryptographically and immutably. No developer can override it. No system can slip free.

This isn't theory. Or a roadmap.
It's already running.

Demonstration logs, cryptographic proofs, whitepapers, early peer review, *even the rejections*, are preserved and timestamped in the public archive:
spqrtech.ai/archive.html

I've already shared it with NIST, DARPA, the Vatican, and other institutions. Because this is too big

I've already shared it with NIST, DARPA, the Vatican, and other institutions. Because this is too big to leave to chance, or spin.

Not after hype or headlines, just thought you might want a front row seat to what's already in motion. Too often these stories get buried until it's too late to shape them, and I'd rather have this on your radar before it breaks.

If you're curious, I've put together a secure appendix room with the whitepaper, logs, and validation schematics:
https://bit.ly/3Hg0e4q

If this sounds wild but intriguing, I'd love to show you what as Mark Cuban quoted "one dude in a basement" can actually pull off. Even if it's just for the story.

Or if you're even a little curious, I'll show you what's already running, and why it might change what you expect from AI.
No slides, no NDA, just real code.

Best,
Adam
Founder
SPQR Technologies
adam@spqrtech.ai
Ph. +61 458 094 464