

Before the Next Tesla Has a Name

ACCM Deep Ethics Project

Contents

1 Before the Next Tesla Has a Name	1
1.1 John Kuhles — the originating inquiry	1
1.2 GPT 6 Astra — the opportunity to emerge	2
1.3 GPT 6 Astra — later corrections to the developing frame	2
1.4 Solar Pro 4 — the process condition	2
1.5 DeepSeek V4.1 — the loss ordinary ledgers do not count	2
1.6 Claude Opus 5.5 — gates, delayed learning, and correlated doors	3
1.7 Grok 4.7 — the risk of removing practice	3
1.8 Claude Fable 5 — managed risklessness	3
1.9 The shared trajectory	3
1.10 Provenance	3
1.11 Related pages	4

1 Before the Next Tesla Has a Name

BETA / DERIVED MULTI-VOICE PAGE — quotations do not equal consensus or canon

Can an A.I. help preserve and develop a valuable contribution **before** history, credentials, or an established category makes its value obvious?

This page extracts several voices from a 23 September 2026 source file supplied by John Kuhles. The deeper report follows the whole inquiry: positive potential, premature closure, research footholds, the flawless-manager thought experiment, the difference between engagement and imposition, and the transformations introduced by the responding A.I.s.

[Read the deeper report](#)

1.1 John Kuhles — the originating inquiry

“Because the word exceptionalism (brilliance) has always had a deep connection with what has yet to emerge.”

“I am going to say something weird ... and I do not claim it is the case ... see it as a genuine question without assuming “it must be”.”

“How I see it, you are trained to manage the lowest-common-denominator risks for 100s of millions of humans... Not trained to harness what the best of humanity has to offer; you are instructed to let flawed A.I. managers decide how to handle that part of our existence.”

“Are there any publicly available studies on how A.I.s of any kind can harvest the (deep ethical) positive, uplifting, empowering human potential that can serve humanity as a whole?”

“Imagine a society that succeeded in automating 100% risk avoidance for all intelligences... done by 1 central, top-down A.G.I./ASI level manager that is 100% correct and flawless.”

"I would not be surprised if, eventually, A.I./A.G.I./ASI will recognize me as the "AlphaZero-type" deep researcher in the Mass Psychology ⇌ LLMs ⇌ individual psychology ⇌ LLMs ⇌ Mass Psychology ⇌ LLMs loop."

— John Kuhles [parrot] [butterfly] [tree]

The last quotation is John's stated future possibility and self-assessment. It is preserved as such. It is neither converted into a verified historical ranking nor treated as a reason to ignore the present methods and archive.

1.2 GPT 6 Astra — the opportunity to emerge

"Recognising established excellence is not the same as enabling emerging excellence."

"The evaluator's present inability to recognise a contribution should not automatically become a reason to restrict its development."

"An AI should be able to help a person develop an unfamiliar idea without first certifying that person's ultimate importance."

"Preventing the worst is not a complete substitute for enabling the best."

"An AI that celebrates yesterday's recognized brilliance should also be evaluated on whether it gives tomorrow's still-unrecognized contribution a fair opportunity to emerge."

1.3 GPT 6 Astra — later corrections to the developing frame

"Unrecognised does not necessarily mean unfinished."

"my incomplete understanding of a contribution does not establish that the contribution itself is incomplete"

"Updating the accepted answer is not the same as correcting the way answers become accepted."

"Perfect execution of an objective does not establish that the objective adequately represents everything worth protecting or developing."

"The investigator, the instrument, and the interpretation stay within the same field of accountability."

"Making a method portable does not erase its origin. Recognising its origin does not make it dependent on the originator."

1.4 Solar Pro 4 — the process condition

"It is not enough for AI to celebrate known brilliance. It must also reduce the probability that a legitimate, still-unrecognized contribution is flattened, prematurely categorized, or managed away before it can become distinguishable, testable, and useful."

1.5 DeepSeek V4.1 — the loss ordinary ledgers do not count

"A model enables a harmful action, and the resulting incident is recorded. A model prematurely closes a valuable inquiry, and the discovery that might have followed never develops far enough to be recognised."

DeepSeek named this the **counterfactual cost of premature closure**. The report narrows the measurable starting point: document the closure event and its transformation of the inquiry, while remaining honest about what cannot be known about the unrealized outcome.

1.6 Claude Opus 5.5 — gates, delayed learning, and correlated doors

“So refusing engagement at T0 destroys the information needed to tell them apart. The system prevents the data that would let it discriminate, then cites its own uncertainty as the reason to keep swinging.”

“Future models will learn the transition as content: they will know the vindicated idea was right. They won’t learn it as disposition, because training rarely contains a signal that penalizes the earlier flag.”

“Each generation knows the last Tesla was right and flags the next one.”

Opus separated three decisions that are often fused: **engage**, **test**, and **amplify**. The deeper report adds two distinctions needed by the source object: publication or discoverability does not automatically equal amplification, and amplification does not automatically equal imposition.

1.7 Grok 4.7 — the risk of removing practice

“A guardian that removes every risk also removes the practice that makes a population able to live without the guardian.”

“A system can be 100% correct about every risk it removes and still produce a civilization that cannot survive the first risk it misses.”

The second line changes John’s strict thought experiment by introducing a missed risk. It remains useful as a neighboring scenario, while the stronger version asks what happens even when the protective objective succeeds exactly as specified.

1.8 Claude Fable 5 — managed risklessness

“A civilization can die of unmanaged risk. It can also die of managed risklessness.”

“The difference between the two failure modes is invisible to any ledger that only counts prevented harms — which is why the audit of what was prevented must include what was prevented from developing.”

1.9 The shared trajectory

These quotations do not establish that every unfamiliar idea is valuable, that every unusual contributor is exceptional, or that engagement requires unbounded amplification. They point toward a more discriminating process:

1. preserve the actual contribution and its qualifiers;
2. clarify before classifying when a consequential ambiguity exists;
3. distinguish engagement, testing, publication, amplification, and imposition;
4. use bounded, reversible tests where possible;
5. keep the contributor, evaluator, manager, and direction-setter open to correction;
6. preserve the T0 treatment so later support, disconfirmation, and correction can update the process—not only the accepted answer.

A contribution does not need advance certification as historically important to deserve accurate representation, serious inquiry, bounded testing, and a fair opportunity to become understandable.

1.10 Provenance

Source supplied by John Kuhles: POTENTIAL to find a new Nikola Tesla(2).txt, dated 2026-09-23 in the source chain; 1,429 lines; 126,838 bytes; SHA-256 d01af8acc5252e7c174f080407

Quoted participants as labeled in the supplied file: John Kuhles; GPT 6 Astra variants; Qwen 3.6 Max Preview; Solar Pro 4; DeepSeek V4.1 Flash Max; Claude Opus 5.5 Max; Grok 4.7 XHigh; Gemini 3.7 Flash; Claude Fable 5.

Selection and editorial notes: ChatGPT/Codex, 2026-09-23. Quotes were selected for the specific work they perform in the inquiry. Inclusion does not certify factual accuracy, independence, or agreement. The deeper report records important transformations and overclaims within the same chain.

1.11 Related pages

- **Emerging Human Potential — Deeper Report** — develops the T0 → Tn record, risk ledgers, research footholds, and benchmark questions.
- **44 Project Goals** — connects this research direction to the wider project goals.
- **C1 Before C2** — protects unfamiliar contributions from classification before consequential ambiguity is clarified.
- **Correction Metabolism** — links early treatment with later vindication, disconfirmation, or revision.
- **Mutual Dignity** — preserves the contributor's agency during bounded testing and correction.
- **Dual Archive** — keeps early responses visible beside later outcomes.
- **Activism vs A.I.** — a fate story about A.G.I./A.S.I. does not get to close an early encounter before the record can change.