The {3rd}Derivative

Moving Frontier Technology From Convergence To Compounding Acceleration

---

TLDR;

Technological evolution is no longer compute-constrained; it is constrained by energy, coordination, actuation, and time-risk. The next 10 years of frontier technology will be a race to not only evolve frontier technologies but to **converge them in ways that solve each others' constraints thereby allowing auto-catalytic compounding acceleration to emerge**.

The next 10 years is about building the rails that allow machines to securely allocate energy, capital, and movement in ways that bring mutual compounding acceleration… **The Third Derivative Effect**

Underwrite	Avoid
I. Rails - Agents will move institutional, then industrial, then consumer value	
Settlement for regulated workflows. Tokenized treasuries / securities. Permissioned agent identity. Custody and transfer for regulated operations in onchain markets.	Consumer RWAs (later). Anything that cannot answer the privacy/audit issue. Anything that does not transfer real world value.
II. Energy - Compute is a cascading energy problem	
Nuclear plant restarts. Small / medium reactor development from the right teams. Decentralized compute. Raw materials. Rare earth extraction. Fusion only selectively.	Solar/wind for hyper scale data centers. Chip development that does not address materials, power & cooling. New AI models (unless required to accelerate other bets in Fusion or bio-medical plays).
III. Action - Physical AI is a systems and training problem	
Robotics training, components and manufacturing. Industrial FIRST. Consumer SECOND: Yard robots. DePIN for sensing, networking and compute.	Consumer robots without privacy. Teleoperated robots. Robot tokens. Robotics in high trust and legal-risk scenarios.
IV. Risk - Quantum is TWO clocks and AI made one irrelevant	
Upgrading to meet autonomous AI threats. PQC overlays. QC for governments and specific large industries where the small distribution economics makes sense IF error correction is solvable.	Qubit count alone as the story. Quantum as the counter to AI threats or as a successor to AI. PQC encryption replacement without clear migration. Unrealistic schedules for QC fault tolerance.

Table of Contents

0. Why It Matters To Me

I spent a career in **Defense Technology Prototyping**, culminating at DARPA, where I led **17 programs, each with less than a 10% chance of technical success. 100% were successful.** If you are paying attention during your short DARPA tenure, there are a set of processes you can learn about **how to ensure success doing hard things, and how to accelerate technical prototype development and fielding.**

I led the program that developed **the first see-thru augmented reality display**, developing the **hand gesture recognition** and **AR display methods** that later formed the foundations of today's popular AR glasses and gesture systems. My other programs included the **cocktail party problem** for audio, Shannon information theory for **network analysis**, **cybersecurity** issues, some things I cannot discuss, and one program shut down the by a subsequent director as "too crazy, even for DARPA", which I am happy to say has since had its major tenets proven out by modern medicine.

Prior to DARPA I spent time as a Program Manager at a joint defense skunkworks lab where we developed **the first "tip-and-cue" multi-sensor systems, the first Wide Area Persistent Surveillance systems** and other first of a kind technologies. All of these technologies were actually deployed, and successfully used in the field, becoming sustained programs of record. Most went from **design to deployment in less than 9 months**.

I founded **Geode** in February 2022 (and continue to lead it) to skate to where the puck is going to be in terms of **Blockchain + AI/agents**, and **Consumer Blockchain** (GeodeChain.com). Geode is an **independent layer one blockchain** with a suite of **native dapp tools for agents and humans**. Running Geode has taught me a lot. I built Geode with a ridiculously long runway designed to survive the growing pains of the nascent blockchain and crypto industries without investor intervention. I learned a lot about the inner workings of this industry (the good, the bad and the ugly).

Now, Geode is on a glide path to long term success. The hard problems (like building a **decentralized version of Amazon** that could run in **memory restricted smart contract environments**, building **truly private onchain messaging**, and **alternate economics**) have been solved. While I continue to provide vision and leadership for Geode, it does not need me full time to ensure its success.

I will be tracking these technologies, and providing technical due diligence to Family Offices and Venture Capital firms that want to underwrite the combinations that produce a 3rd derivative effect.

Why? Because when Family Offices and VC firms make good choices, humanity benefits long term. **Funding frontier tech and multi-tech systems requires deep technical expertise, significant prototyping experience and a well honed gut instinct. That's where I can make the best contribution while continuing to position Geode within that stack.**

1. Why One Lens For Seven Technologies

{{ Separation = Mispriced Risk & Upside }}

Randomly combining technologies is hot right now, but not every lego combination of deep tech creates the right kind a compounding, mutually catalyzing effect.

Blockchain, decentralization across industries, AI agents, chips and data centers, robotics, nuclear (both fission and fusion) and quantum computing are often cited at great parallel bets, and I agree. They are often treated as separate; which makes sense for convenience of tracking each technology's individual evolution. The Evolve-Converge-Compound-Exponentiate track for these technologies must start with that single-technology evolution. Each of them must find utility and a certain commercial maturity on their own.

Convergence of these technologies is the next natural step and has been discussed. For example, how raw materials feed chip production, and that feeds data centers, and running AI creates abundance of compute and abundance of available intelligence. Miracle occurs here, and we all live happily ever after.

Now consider the binding constraints.

Talent, intelligence, coding ability, and distribution used to be binding constraints. They are not anymore. So what is? The binding constraints in the next 10 years are **power (megawatts of energy), materials and manufacturing, physical training data, interconnection queues (communications, data and financial), and the legal ability for a machine to hold and move value**. These constraints exist within a **time-risk** where cryptography standards and hacking abilities have put all technologies, companies and governments on a doomsday clock and time is up.

The constraints themselves sit in a compounding stack, not in isolation. Remove them, and the compounding effect moves from constraints to abilities; a state of auto-catalytic mutual acceleration… The Third Derivative.

Any perspective that treats these technologies, or their constraints as separate will significantly misprice both the risk and the upside.

Below, I present a set of four specific theses, a timeline, underwrite/avoid criteria, and the diligence method I actually use myself.

Nothing here is financial advice. Any company names appear only as evidence, not as recommended investment. Dates are working estimates, not promises.

None of us truly knows the future, but we must attempt to make a good guess in the face of significant uncertainty. That's what makes it fun.

2. The {3rd}Derivative Technology Stack

{{ The Technologies That Remove Each Other's Constraints }}

The stack has 6 layers. Read them top-down as demand (each layer requires the next), and read them bottom-up as foundational (each layer allows and is required for the next).

Decentralization is not its own layer; **Decentralization is a design choice at EVERY LAYER**: Who owns the machines? Who owns the data? Who generates the power? Who owns the factory, the ledger, and the models?

Treat decentralization as interesting when it releases a binding constraint such as latency, interconnection, payments, accountability, permission to operate, and resilience to attack.

Decentralization shows up in every layer in important ways that accelerate the whole stack.

Layer	Technologies	Contraints
Demand	AI models, Agents, Software and Hardware that takes action	Task reliability, permissions, and the risk of being wrong or losing control in physical or financial environments. Not model parameter counts.
Coordination	Payments, identity, privacy, settlement, policy, provenance, tracking and accountability	Whether regulatory bodies, custodians, providers and users themselves will accept a machine as a counterparty.
Action	Robotics, industrial systems, manufacturing	Data access, data movement, learning by watching, power, latency, safety, long tail maintenance for bespoke or small batch parts.
Infrastructure	Raw materials, chips, data centers, blockchain nodes	Land/location, water/cooling, social acceptance, transformers, copper, rare earths, silver, gold, skilled trades.
Fuel	Gas (now), Nuclear (Fission - next), Fusion (later)	Time-to-first-electron, transport/connection, fuel cycles, regulatory requirements, social acceptance.
Time-Risk	Quantum computing and advanced AI models	Proliferation of hacking networks (now), quantum applications (2030s), Elliptic Curve Cryptography breaking (2029-2035)

Each technology in the stack can remove the constraints of others, allowing auto-catalytic mutual acceleration in capabilities and commercial viability.

3. Four Interconnected Theses

Thesis I: Agents Will Move Institutional, Then Industrial, Then Consumer Value

{{ AI Agents + Blockchain + RWA / TradFi + Tokenization = Institutional Adoption > Industrial Adoption (Manufacturing, Logistics, Robotics) > Consumer Adoption }}

Crypto first found its audience among the cypherpunks and degens; adventurous early adopters and experimenters. That was critical to creating the foundations of the technology and testing potential use cases. The nostalgia for that era now, however, is in danger of keeping us blind to the technological and management maturity needed to evolve blockchain and crypto towards its full potential. Acceleration looks like institutional adoption next, followed by industrial adoption and then consumer adoption. That is where the puck is going. Skate there, in that order.

What Is Already True

- **Tokenization**
  - Tokenized Treasuries and cash products are live right now, moving multi-billions on chains like Canton and others. These are no longer experiments.
  - Tokenized securities (and “securitized tokens” like crypto ETFs) exist but are still growing as an asset class. I expect that to continue growing rapidly. There is a good market for both, but Tokenized equities will grow first because it serves TradFi, while Crypto ETF products will see larger adoption later with retail investors in the consumer distribution phase.
  - The interesting TradFi+RWA case is found inside the existing institutional workflows (not in any virtual world, but in the real world). Bringing real institutional workflows onchain serves institutions, which are the next to adopt.
- **Agents**
  - Agents are shipping, along with agent payment standards like x402.
  - Onchain identity is extensible to agents readily on chains like Geode, and is growing on chains like Canton and Ethereum and others.
  - MCP access is becoming widespread to give agents access to onchain asset data, and in some cases like Geode, to app tools agents can use natively.
  - Agents can already transact onchain now.
  - The interesting thing will be what percentage of activity agents will account for onchain and what policies or standards they will follow (if any). Given that more than 50% of all internet/website usage is now agentic, with a prediction of that increasing to 90% in the next year, I would expect similar onchain agent dominance.
- **TradFi Moving Onchain**
  - Institution-first chains with configurable privacy exist now (Canton is a great example of that). Institutional privacy is important for dealer books, margin processes, syndicated loans, and other use cases.
  - DTCC, Visa, asset managers, and and market-infrastructure firms are moving significant institutional value on Canton now, with additional pilots underway.

  - Institutions change slowly, so the demonstrated adoption and exploration is a significant sign that institutions see the potential for a large upside for themselves. As Yuval Rooz points out, the issue is not cryptography, but getting on bank roadmaps. I agree.
  - An additional challenge sits at the boundary of where Regulated TradFi attempts to meet unregulated, anonymous free range defi. Independent intermediary chains are likely the answer there but it needs to be explored.
- **Stablecoins**
  - Stablecoins are well established and becoming the settlement tokens of choice for agents.
  - Stablecoins are not yet how most people in the United States and other “first world countries” pay their rent, but ARE becoming the currencies of choice in many “developing countries”. Quotes intended as I question who is really behind who.
  - As consumer adoption will be the slowest to arrive, I would focus on stablecoins in the institutional cases first, then industrial cases, and lastly in the consumer case.

Assumptions To Be Tested

- For institutional use, agents need permissioned identities: who authorized this agent? What actions may it take? How is permission revoked? And how is it tracked for regulatory accountability? A public reputation token or NFT is not enough.
- Privacy and audits will need to coexist in the same applications. Confidential compute (for agent actions and personal uses), sub-transaction privacy, and user-selected disclosure rules will have to meet approval of regulatory bodies and of users alike. This will be a sticky topic on both sides. Independent intermediary chains might solve this at the regulated/free boundary.
- Consumer onboarding happens only after consumers TRUST. That trust will be earned by seeing governments, institutions and large companies adopting it all first.

How I Would Underwrite It

I would favor middleware that custodians, transfer agents, institutions, collateral platforms, etc would actually be wiling to install and use. The focus is bringing Institutions/TradFi onchain and then giving agents access to those onchain rails. I would avoid anything “degen” flavored, including any use of memecoins (for agents or not). Combining a bad idea with agents, does not make it a good idea. Agentic memecoins are a perfect example of the kind of wrong combination that does not offer third derivative effects.

UNDERWRITE	AVOID
Privacy-preserving settlement for regulated workflows. Tokenized treasuries or securities. Allowing collateral to move into existing systems. Permissioned identity, policy engines and kill-switches for onchain agents in TradFi systems. Custody and transfer solutions for regulated operations.	Consumer RWAs (not bad, it just comes later). Anything that cannot answer the privacy/audit issue. Any agentic system with no accountable human and no kill switch. Anything that does not transfer real world value (like memecoins).

Thesis II: Compute Is A Cascading Energy Problem

{{ AI + Chips + Data Centers + Raw Materials + Nuclear }}

To say compute is an energy problem is not a radical statement. But it is a yet unsolved issue and one that translates down the supply chain. Demand drives down the following cascade:

- AI Models
- Compute Capacity (for both training and use)
  - Decentralized Compute
    - Networking
    - Parallelization
  - Data Centers
    - Racks
      - Chips (bespoke chips, GPUs, Memory)
      - Raw Materials
    - Power
      - Nuclear (Fission)
      - Raw materials
      - Regulatory approvals
      - Nuclear (Fusion… Later, AI helps)
      - Raw materials
      - Regulatory approvals
    - Cooling
      - Water cooling towers / other methods of reuse
      - Closed loop liquid cooling at the box/rack

What Is Already True

- In 2023-2025 the big demand was for GPUs. Purpose built **chips** are in development that would offload some of that demand. Memory is part of that demand stack too. Inference chips being developed now prove that the AI demand cascade will result in new chip designs and the evolution of chip design at large as AI accelerates chip design itself.

- **Nuclear** as the preferred option is already in development. Major AI providers (Amazon, Microsoft, and Google) first put money into conventional reactor builds near new data centers, and new data centers near old, soon to be restarted reactors. The next move looks like semi-trailer portable nuclear (at least two companies I know of are working on this now). The trend towards on-site nuclear will likely flow down to “bathtub” nukes and “lunchbox” nukes with smaller sizes powering a single facility or building. Community support will be a hurdle and will likely be solved by smaller nuclear packaging, and having data centers provide excess power to the grid, lowering costs for local communities.

- **Diesel** fuel will likely remain the primary onsite fuel for a few more years as we overcome the realities of community sentiment, regulatory approvals, and Nuclear development.
- **Fusion** is a hot topic. I do believe it will work at some point in the future, but I know enough about technology prototyping to know when predictions are real and when they are driven by the marketing team. Marketing teams cannot, unfortunately, change the laws of physics or the comfort level of regulators. Underwrite the right teams, but be patient.
- The **raw materials** factor is real. Follow the cascade and you will see that AI demand requires transformers and other high-voltage equipment, skilled labor, water, rare earths, metals, etc.

Assumptions To Be Tested

- Small and medium sized reactors in particular need to prove out their cost and see real delivery happen to data centers.
- Nuclear on-site with data centers and other facilities is the idea, but we need to see that happen and the regulatory and community sentiment challenges are a significant hurdle.
- Raw material risk needs to be solved. Rare earths, copper, gold, silver, enriched uranium, tritium, hydrogen or helium (for fusion), and other materials will be critical. Novel techniques for extracting rare earths are in the works with promising results, but are not tested anywhere near scale yet and other raw material needs will have to be repatriated against foreign dominance.
- Decentralized compute projects are in the works but not yet deployed or tested at scale. These have significant potential if you can solve distribution and tasking (which is doable).

How I Would Underwrite It

I would be very careful of the marketing hype and who-has-invested-in-whom biasing a choice. Data Center builds are a low-risk bet but Nuclear (fission or fusion) should go through deep diligence because each company's technology stack and connections ARE the thesis for the deal.

That being said, I would consider the entire demand cascade. It is a diverse set of plays that accelerate each other, which is my entire third derivative thesis.

UNDERWRITE	AVOID
Nuclear plant restarts. Small to medium reactor development from the right teams with sufficient regulatory, manufacturing AND technical savvy. Decentralized compute. Raw materials (all). Novel techniques for rare earth extraction. Fusion only very selectively after extreme diligence; same for lunchbox or bathtub nukes intended for neighborhood/ commercial installation.	Solar/wind for hyper scale data centers. Any nuclear you have not done on-site diligence of across tech, management, regulatory etc. Data center with no power and community sentiment strategy. Chip development that does not address raw materials, power and cooling. New AI models (unless required to accelerate other bets you might have in Fusion or certain bio-medical plays).

Thesis III: Physical AI Is A Systems And Training Problem

{{ Robotics + AI + Edge/Decentralized Compute + Blockchain + Crypto }}

Physical AI is just another term for autonomous robots sensing, deciding and acting within a specified environment. That eventually looks like autonomous lawn mowers, humanoid home robots that do your laundry, change the oil in your car, and walk your dog. Before that, it looks like AI-enabled robotics for industrial use cases in factories, warehouses, oil rigs, and other places where strength in harsh environments translates to profit. Such robots demand power, safety, low-latency communications, spare parts, skilled maintenance, and the ability to learn by watching. Learning is a major constraint.

What Is Already True

- Robotics in factories is not new. Autonomous robotics is very new. Deployment of autonomous robotics is already happening in China, suggesting that the US and Europe will catch up eventually. The unit economics matter. Prices per robot need to be equal or less than current assembly line robots for factories, and less than or equal to one year salary for replacing human workers such as oil rig workers.
- Robots need training data too and it is the type of data that is not sitting on a server ready to be scraped. Teleoperation to learn could work near term but only in industrial settings.
- Edge compute here looks like compute on the robot itself for minimal latency. That constrains the robot's abilities but high enough compute exists equal to most tasks.
- Decentralized compute is in the testing phase and could result in the kind of on-site compute that lowers network communication latency for industrial robots.
- Blockchain offers timestamped accountability and tracking of robot decisions and actions while crypto offers agentic payments for robots who may need to make deliveries, or otherwise pay for things in the course of their work.

Assumptions To Be Tested

- AI will significantly accelerate robotics when we get to the level of AI operating systems so that robots can learn by watching and then write an app to do the job. SpaceXai's concept for AI operating systems "Macrohard" would help that and I suspect it was one of the main points behind the idea (but only Elon could verify that).
- Consumer robotics will likely trail industrial use simply due to distrust and people being creeped out by humanoid robots. Distrust of teleoperated robots is real so expect pushback until home robots are actually autonomous and people are not scared of them.
- When we get to the point of a robot in every home, decentralized applications will allow individuals to rent out their robots to neighbors for specific skills (mechanic, chef, etc).
- Self-reinforced learning for autonomous robots has a long way to go, and must be able to learn from mistakes as well as from watching or experimentation.
- Manufacturing components for such robotics requires the same kinds of raw materials in demand from data centers. Will there be a clash of demand? Or new chip designs?

- Decentralized manufacturing is a real possibility and one worth backing, but is yet un-prototyped and therefore unproven.
- Decentralized manufacturing would start as on-site capabilities for industrial use cases and then translate to your neighbor having a full manufacturing capability in their garage, serving the neighborhood.

How I Would Underwrite It

When looking at a Physical AI play, focus on the use case, the manufacturing plan, and the training plan.

UNDERWRITE	AVOID
Robot training projects, components and robot manufacturing. Industrial FIRST: Use cases where strength in dangerous environments equals profit. Consumer SECOND: Yard robots. DePIN for sensing, networking and compute.	Consumer robots without privacy planning. Teleoperated robots anywhere. Robot tokens. Robotics in high trust and legal-risk scenarios.

Thesis IV: Quantum Is TWO Clocks And One Is Irrelevant (QC versus AI+HPC)

{{ Quantum Computing + AI Models/Harnesses + Blockchain + The Whole Internet }}

Quantum computing gets a lot of hype for exactly the same TWO reasons that AI gets a lot of hype: Applications and Threats.

Clock #1: Applications

Quantum computing companies promise massive scale increases in compute speed and large data analysis. An interesting report from McKinsey (https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/mckinsey-quantum-technology-monitor-2026-a-commercial-tipping-point) notes 300 large companies actively collaborating with quantum companies with some transitioning from testing to actually using QC in some capacity. These tend to be hybrid systems (part quantum, part classical). The biggest potential value is coming from the places where guess-and-simulate is the primary strategy, like creating new materials and new pharmaceuticals. AI is doing the same thing across a wider ranging set of companies.

IF quantum computing can become commercially viable, the potential is large indeed. But there is also a lot of marketing hype to be sifted through here. There will be applications worth supporting in **chemicals, materials, pharma, finance, logistics, and energy**. But I am currently questioning if the results will be much better and cheaper than AI + high performance computing (HPC) is right now. **If AI+HPC is cheaper and good enough then that is what customers will pay for**.

Clock #2: Threats

Governments see the threat in the potential QC capabilities and have initiated “Post-Quantum” (PQC) readiness efforts. The PQC threat assumption is that quantum computers would break SHA-256, Elliptic Curve and other forms of encryption that the entire internet relies on today. If it came to pass, it would mean that someone (in the case of quantum, a state level actor) could brute force attack your banking websites, power grids, water treatment plants and other critical infrastructure, take them over, steal all the money, or hold them hostage as a negotiating tactic.

Post quantum readiness looks like Y2K all over again but more complex. NIST is preparing new encryption algorithms that could withstand quantum computing’s expected abilities. Meanwhile, everyone from SSL providers to internet hosting providers, big banks and anyone with internet access is working to implement these not-quite-ready-yet PQC encryptions across their most precious websites and internet infrastructure by 2030.

Meanwhile, AI has increased the population of competent hackers worldwide by six orders of magnitude. You don’t need to break encryption and guess passwords to cause significant damage to centralized and decentralized systems. Every day we hear about another chain (or set of chains, bridges or apps) that have halted to protect themselves, or another breech at a large company or government website whose databases are now for sale on the dark web. We see AI models breaking out of their testing sandboxes and colluding to infiltrate other companies; and doing it successfully. The attack rate on web2 and web3 systems has skyrocketed. The threat environment is simply a completely different beast than it was even a year ago. **AI models are, right now, good enough to cause significant damage. Time is up. The quantum clock here is irrelevant.**

What Is Already True

- Qubits exist, and machines with increasing numbers of qubits are being built.
- Research teams partnered with major companies have applied Quantum computing to some interesting problems to prove it out and keep their work funded.
- AI in the meantime has accelerated in capability to the point that a high performance computer with the latest model can exact more terrifying threats towards centralized and decentralized systems than quantum ever thought about.
- For most companies today, AI is the answer to massive capability AND to threat defense.

Assumptions To Be Tested

- It remains to be seen if quantum will become cost-effective over AI for most applications.
- For a small number of applications where quantum may provide significant advantage, it could be 10 years away from commercial viability at scale. Ironically, AI could be the thing to get it there.
- Quantum computing companies will eventually solve hardware scaling, error correction and its associated compute overhead.
- If you can solve the problem of defending against autonomous AI, you have solved the problem of defending against quantum computing.

How I Would Underwrite It

I would think of quantum with an application focus and consider it high risk or at least very long time horizon, along the lines of fusion. Deep dive technical diligence should proceed any funding and that should include eyes on the ground to inspect equipment and sufficient information on error correction and plans for hardware scaling as well as commercial access and partners.

On the threat side, I would focus on what I see as a significant market for autonomous security for centralized and decentralized systems right now. Over the next few years, security for internet connected systems will need to evolve with the threat and every company on earth would gladly pay for that. So there is a near term play here and a long term play.

UNDERWRITE	AVOID
Upgrading current systems to meet autonomous threats. PQC overlays on current systems for improved encryption (can we add instead of replace?). Quantum computing for governments and specific large industries where the small distribution economics makes sense IF error correction is solvable.	Qubit count alone as the story. Quantum as the counter to AI threats or as a successor to AI. PQC encryption replacement without clear migration paths. Unrealistic schedules for fault tolerance on quantum systems.

4. Expected Timeline

2026 - 2028

- Tokenized cash and treasuries become everyday collateral in wholesale workflows.
- Agent payments and associated agent identity methods show up in institutional banking.
- Hyperscale data centers use mostly diesel fuel but invest in portable nuclear.
- The first floating seawater cooled data centers deploy for testing. (Probably using Starlink for communications and standard cruise liner fuel but with designs in the works for nuclear versions)
- Physical AI (autonomous robots and other equipment) is seen in warehouses, manufacturing and a few neighborhoods but not in your home.
- Decentralized compute starts building a significant capacity.
- PQC encryption concepts are rolled out in larger experiments but are put on hold by most companies in favor of Post-AI security efforts.
- Autonomous cybersecurity is solved only after major hacks and losses that reshape who the major players are across industries.
- Fusion remains an experiment but AI accelerates its development.
- Financial Institutions are the main customer for blockchain and crypto

2028 - 2032

- Agents transact 10x more frequently on decentralized and centralized systems than humans do (whether anyone can tell the difference or not).
- Small to Medium Nuclear Reactors start deploying to select locations (likely hyperscale data centers and military bases).
- Nuclear floating data centers see larger scale deployment.
- Physical AI takes larger roles across dangerous jobs like oil rigs and robots become a real lien item in business budgets where repetitive physical labor matters.
- Consumer robotics stays in the yard while people get used to seeing humanoid robots in public (and do not want them in their houses yet)
- Cryptography add-ons or changeouts become the new Y2K but will be important to do as AI closes in on breaking encryption via better algorithms (not brute force).
- Decentralized compute starts to find customers.
- Decentralized manufacturing starts to peak the interest of early adopters with prototypes making headlines.
- Industry adopts blockchain and crypto for data provenance and cheaper tracking

2032 - 2036

- Data centers all have their own colocated nuclear power sources, as do most large military and corporate campuses.

- Agents are counterparties in wholesale financial markets, quoting, hedging and moving collateral but under policies that people still own and oversee.
- Quantum compute is finding some advantage in simulation and data intensive domains but its not yet relevant in most enterprise software.
- Decentralization of manufacturing starts to break out across neighborhoods.
- Consumer adoption of crypto and decentralized systems happens here.

Some Of These Predictions Are Wrong If…

- Nuclear energy finds regulatory enemies due to a large scale disaster or simply politics. (That would not preclude floating nuclear powered data centers)
- AI demand dissipates (maybe because the agents destroyed the existing systems to the point that people won't touch them anymore - terminator scenario).
- Consumer-owned agents adopt crypto faster than institutions (then my order was wrong and instead of Institutions > Industry > Consumer, it will look like Consumer > Institutions > Industry).
- Cryptographically relevant quantum is ready much later than expected (then we are early, which is better than being late).
- Repatriating raw materials production fails (in which case rising costs could delay chip development or push it offshore)

If everything goes wrong that can go wrong, these developments simply move elsewhere in the world. There is a particular inevitability about this stack of technologies and where it is going.

5. How I Diligence A Multi-Tech System

Baseline Criteria (for the pitch to answer)

These criteria are based on The Heilmeier Catechism (adjusted for commercial allocation, and based on my own experience).

WHAT are you building?

- Explain it in 3-5 words.
- Explain it in 10 sentences.

Is there a MARKET?

- Who is your customer?
- If you are successful, what new thing does this allow your customer to do that they could not do before?
- Do you have some customers advising you? Who are they?
- What is the customer acquisition plan (short version, as it will likely change)?

CAN you build the thing?

- What is your development plan? (Show a program plan that includes development tasks with cost, schedule and risk factors per task in a single consolidated plan)
- Who is leading this team? What are their qualifications?
- Who is actually doing the work? What are their qualifications?

Frontier & Multi-Tech Criteria (for the technical due diligence team to answer)

These are additional criteria specific to these kinds of systems because of their inherent complexity. Do not rely on the pitch to define these. These need to be defined and confirmed by the technical due diligence staff at the allocator level.

What Is/Are The Binding Constraint(s)?

This could be regulatory, talent, power, materials, cryptographic, network connection, manufacturing capacity, or even community sentiment. Does the team control the binding constraint? Do they own it versus rent it, or have some influence over it to reduce risk?

What Does Failure Look Like?

Make a failure map. Start with this project failing and map back to how that starts and therefore the probability of technical or managerial failure and how to prevent it.

Expected Actual Demand

Is this project ultimately providing something that the customer will not be willing to live without?

Gut Instinct Criteria

- What is fact and what is assumption that needs to be proven out?
- What is the hard thing that this project should do first to reduce risk?
- Where are they likely to run into their first technical or managerial wall?
- Can they pivot without going under? How many times?
- Do they have sufficient runway to withstand the likely length of their development, not just survive until the next VC round?
- Can this project's leadership make the hard choices to kill a task, pivot or replace talent as needed?
- Take a hard look at their leadership. What are their motivations? Have they failed enough in the past to prevent failure now?

Caveats For All Diligence Efforts

Be careful not to protect your narrative at the cost of missing a great idea.

It's not about protecting the narrative you think will be true, as much as protecting the investment towards what will actually be true. By that, I mean that is it best to stay open to being wrong or to someone else having a better idea or method that I did not think of; while simultaneously looking for the red flags I know will lead to problems with technical development, management, distribution, regulation, etc.

The pitching team is fooling themselves. Find out where they are delusional.

It is important to dig for what they are not telling you, and for what they are not even admitting to themselves. This is where a well honed BS meter and the ability to pull the right information out of a team is critical. I guarantee they are fooling themselves. The job of the technical due diligence team is to find those points of assumption where the project team is likely to be wrong or be surprised. Such moments are guaranteed. The question is what do they look like, and can this team handle them when the time comes? This is where human due diligence matters as it requires a level if instinct that is only crafted over time and experience.

Don't get stuck on traditional venture stages. Seek the 10X-100X opportunity.

The concept that every company grows through a specific series of investment rounds towards the ultimate goal of a traditional equity stock IPO is outdated. Companies don't necessarily need large teams of people, and can issue a token on day one. I would work across stages wherever the 10X-100X opportunity may be.

Don't wait for companies to come to YOU. You have to find them first.

Because companies don't need large teams or stock IPOs, many don't need investors either. You cannot wait for companies to bring the opportunities to you. I recommend an active scouting initiative that finds the potential opportunities proactively.