CAPRIOLE RESEARCH

Capriole
INVESTMENTS

The Abundance Paradox

Investing when intelligence becomes cheap

AI rewrites economics.
We invest in the new scarcity.

CHARLES EDWARDS
Founder & CIO, Capriole Investments

OCTOBER 2026

Contents

Technology does not eliminate scarcity. It moves it.

A research framework for identifying scarce assets at the frontier of massive demand, assessing who captures the value, and deciding what the market has mispriced.

About this paper . . . 3

Executive summary . . . 3

1. Intelligence becomes infrastructure . . . 5

2. The New Scarcity . . . 7

3. The four laws of abundance . . . 10

4. Tokens are made of atoms . . . 12

5. The second engine: quantum computing . . . 16

6. The synthetic ocean and the human premium . . . 18

7. Finance becomes software . . . 21

8. The Speculation Reflex . . . 24

9. Will power laws get steeper? . . . 27

10. The Scarcity Map . . . 28

11. What could break the thesis . . . 34

Conclusion — Moving Scarcity . . . 37

References and research notes . . . 38

Research scope: eleven chapters, source-linked evidence and explicit tests that could disprove the thesis. Conceptual diagrams are not forecasts or portfolio weights.

About this paper

This paper is an investment thesis about a specific transition: useful machine cognition becoming cheaper, more capable and easier to deploy. Abundance means broad availability at a falling cost for a defined quality of work. It does not mean infinite intelligence, zero resource costs or equal access for everyone.

The central claim is that technology does not eliminate scarcity. It moves it. When one productive input becomes cheaper and more abundant, the relative value of other inputs and outcomes can rise. Cheaper intelligence benefits users. Exceptional investment returns require something more: an asset that remains difficult to reproduce, a durable claim on the economics, and a price that understates how long that advantage can last.

The method is to follow the cost curve, measure adoption, map where scarcity is moving and identify who can retain the resulting cash flow before capital fully reprices it. Historical evidence, current observations and company announcements are cited. Forward-looking interpretations and named Capriole frameworks are analytical hypotheses, not established economic laws.

> The central question is not only how intelligent machines become. It is what their success makes scarce—and who can profit from owning it.

This differs from a forecast of the date of AGI. The thesis can matter well before full automation: lower costs for particular tasks can alter bargaining power without replacing every occupation. We distinguish a technology's usefulness, the profits a business captures and the return available at today's valuation.

Executive summary

A world can become richer while the things people most want become harder to own.

Humanity is industrializing parts of intelligence. Skills that once had to be recruited and trained can increasingly be delivered as software, replicated across workflows and improved centrally. The significant change is not a chatbot that answers questions. It is a falling cost of getting useful work done.

The direction is measurable. Stanford recorded a more than 280-fold decline in the inference cost of GPT-3.5-level capability between November 2022 and October 2024.[2] That is a benchmark-specific historical result, not a claim that an entire employee or business process became 280 times cheaper. The distance between those two statements is where the next phase of investment will occur.

The obvious question is what AI replaces. The more valuable question is what its success makes scarce.

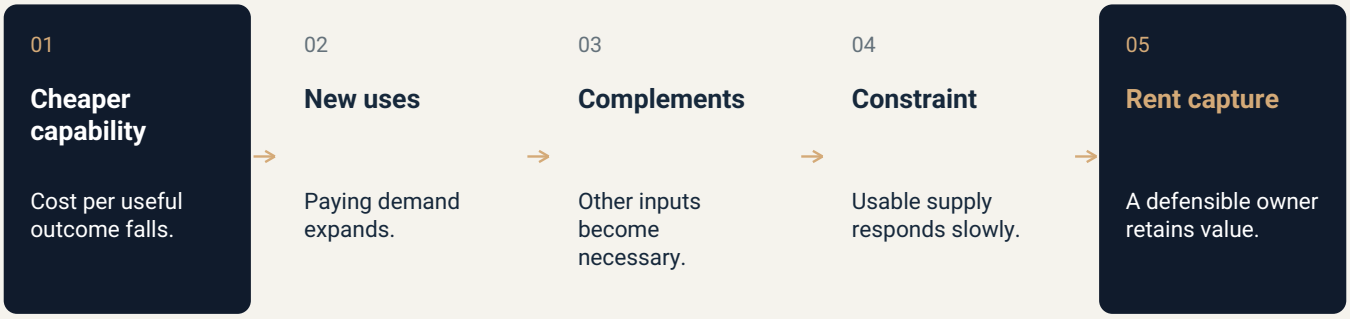
The greatest investment regime of the next decade may not be AI itself. It may be the scarcity caused by AI.

Technology does not eliminate scarcity. It moves it. Intelligence can become abundant while power, physical capacity, trusted execution, capital, attention, ownership and status become scarcer relative to it. The investment opportunity is not a fixed sector. It is the movement of relative scarcity.

Printing expanded access to words. Electrification expanded access to work. The internet expanded access to information. These transitions enlarged the economy without making every participant profitable. They shifted advantage toward new complements, distribution systems and institutions. Intelligence may now do the same across much more of production.

Our thesis is a two-sided repricing: machines demand scarce capabilities, while people continue to compete for scarce outcomes. Between them sits a growing supply of competent, reproducible digital output.

> THE ABUNDANCE PARADOX When useful intelligence becomes abundant, value can migrate to the scarce complements it consumes and the scarce human outcomes it helps people pursue.



Demand must be funded. The constraint must bind. Its owner must capture the economics.

Figure 1. The Abundance Loop: a conditional mechanism, not an inevitable sequence. Demand must expand, the constraint must bind, and its owner must capture the economics.

On the machine side, tokens require atoms: chips, memory, networks, power and physical execution. Quantum computing may become a specialized second engine for selected scientific problems, while also forcing changes to digital security. Neither application is a reason to buy every supplier.

On the human side, cheap function need not eliminate the premium for history, trust, status, privacy or control. A Rolex is not principally a solution to inaccurate time. A private jet is not principally a solution to the absence of commercial flight. AI can reproduce a design more readily than the shared recognition that gives ownership its meaning.

Finance connects the two. Stablecoins, tokenized claims and machine-controlled accounts make ownership and payment more programmable. The wrapper becomes less informative; the underlying rights become more important.

The social transition may be less comfortable. Greater productivity can coexist with weaker labor bargaining power and anxiety about access to ownership. Our Speculation Reflex describes one possible response—not a proven causal chain from AI to unemployment to gambling.

Three distinctions govern the investment case: value created is not value captured; a temporary shortage is not a permanent franchise; and a correct thesis is not necessarily a cheap security. The task is to identify the new scarcity, determine whether its owner can keep the economics, estimate the Scarcity Clock, and move before those economics are fully reflected in price.

> The future is a power law. Own the few.

1. Intelligence becomes infrastructure

Intelligence is becoming infrastructure.

Technologies become infrastructure when their extraordinary capabilities become ordinary inputs. Light, mechanical power, communication and computation followed this path, although access remains unequal. We no longer buy them as miracles. We build businesses around their availability.

Intelligence has unusually wide reach because it helps improve other inputs. Better reasoning can support chip design, drug discovery, logistics, software and coordination. Electricity expanded the work machines could perform. Cheap, reliable cognition could expand the decisions they can make. The distinction is not that earlier technologies were narrow; it is that more of the work of invention itself becomes reproducible.

> Energy converts scarcity into work. Intelligence converts work into possibility. Trust converts possibility into value.

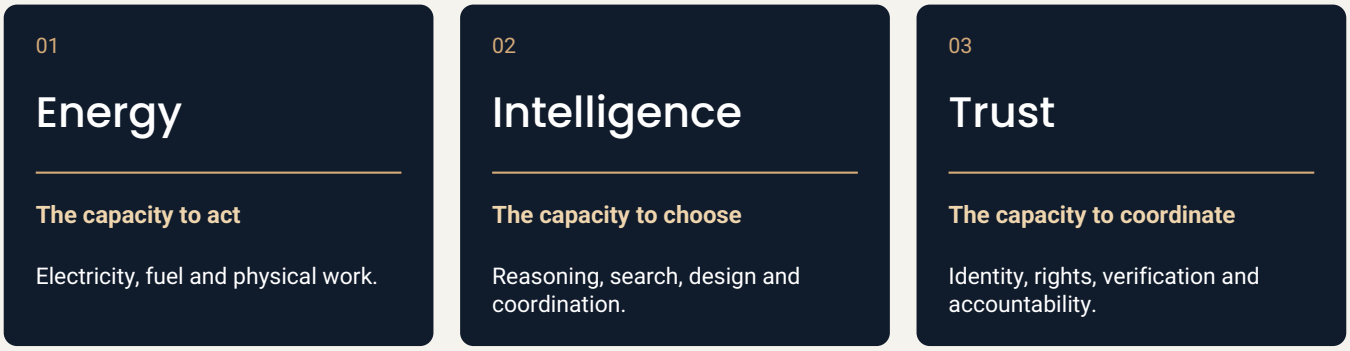


Figure 2. The Three Currencies are a metaphor for productive capacities: action, choice and coordination. They are complements, not interchangeable units or a valuation identity.

From salary to token

Useful knowledge has long been stored in books, software and institutions. But much adaptable judgment has still been purchased through people: recruit a specialist, give them context, and receive a finite number of working hours. Their knowledge can be shared; the person cannot be copied.

Machine cognition changes the scaling model. Capability can be purchased per token, task or outcome, then deployed across many instances. One improvement can reach all of them. Replicating the model is relatively easy; replicating the permissions, clean data, reliability and operating context needed for productive work is not.

This is the capitalization of cognition in an economic sense. Some work previously supplied by labor is embodied in models, compute, proprietary data and operating systems that can be financed and scaled. It is not an accounting instruction: API payments remain an operating expense for many customers, even when the provider owns the productive infrastructure.

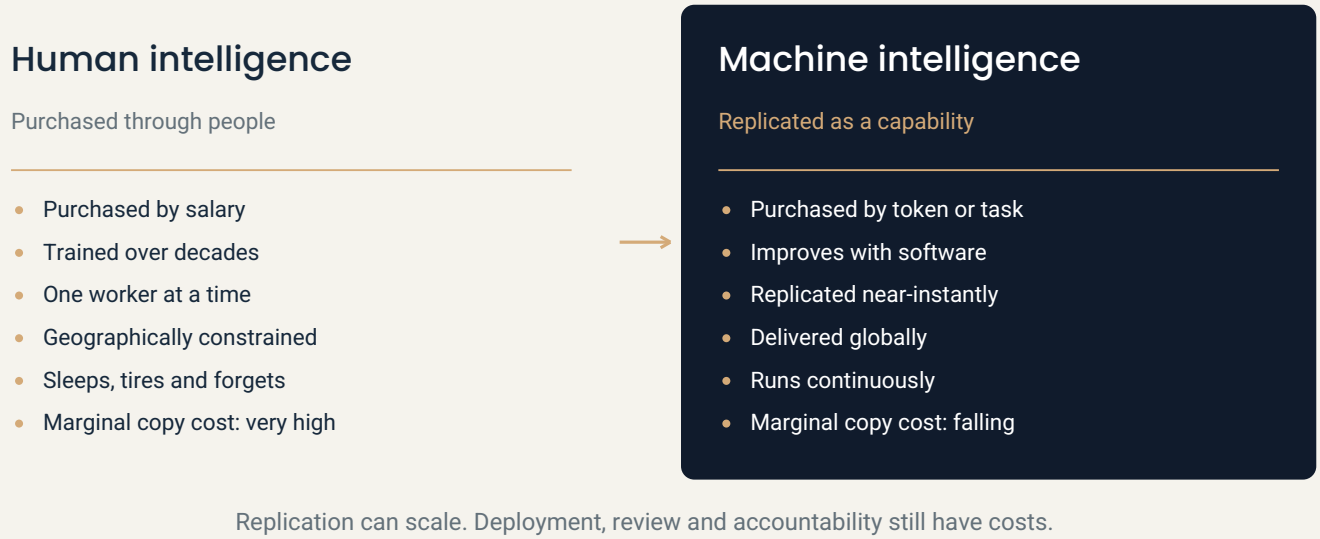


Figure 3. The Salary-to-Token Transition concerns the economics of replication. Deployment, review and accountability still have costs; the diagram is not an accounting classification.

The analogy is not perfect. Intelligence will not become literally free. The token has an energy cost, a chip cost, a latency cost and often a human-verification cost. Many tasks require accountability, context, physical presence or trust. But the direction of travel matters more than the endpoint. When the marginal cost of a useful cognitive task falls by one or two orders of magnitude, markets reorganize around the new price.

The relevant price is trusted work

A cheap answer is not necessarily a cheap outcome. The useful denominator is an accepted unit of work at a stated quality and reliability: a resolved case, a tested code change, a verified design. Its cost includes inference, integration, supervision, correction and expected loss from errors. Tokens are an input to that result, not a substitute for measuring it.

Consider a workflow in which 80% of the original time is spent producing an answer and 20% reviewing it. Make production ten times faster while review stays unchanged. Total time falls to 28% of the original, a 3.6-times improvement—not ten times. This is an illustrative calculation: 0.8 / 10 + 0.2 = 0.28. Extra review can reduce the gain further.

Early evidence illustrates this dependence on context. A 2023 study of customer-support work estimated roughly 14% higher productivity with generative assistance.[40] In contrast, METR's randomized early-2025 study found experienced open-source developers took 19% longer with the tested AI tools. METR's February 2026 follow-up cautioned that selection effects made its newer speed estimates unreliable.[34][35] These are different tasks and populations, not a scorecard for or against AI.

The next valuable business may not generate the answer. It may make the answer safe to use.

The first-order mistake

The first-order question is which tasks can be automated. The second is how businesses reorganize when those tasks become much cheaper. Substitution is one source of value. Expansion into work that was previously uneconomic may be another.

The second-order analysis asks what happens when every company, consumer and machine can afford vastly more intelligence than before. The answer is not merely fewer salaries. It is more experiments, more software, more personalized media, more security monitoring, more scientific search, more financial optimization and more machines acting continuously.

A cheaper unit can support a much larger system. But greater usage is not the same as greater spending, and neither is the same as greater electricity demand. All three depend on how quickly demand responds relative to improvements in efficiency. The investment case needs to measure that response rather than invoke a slogan about it.

> **THE SALARY-TO-TOKEN TRANSITION**
> The decisive shift is not that a machine can perform a task. It is that the same capability can be replicated, metered and compounded like capital.

2. The New Scarcity

Technology does not eliminate scarcity. It moves it.

History does not repeat, but value migration does. Civilization advances by turning expensive capabilities into widely available inputs. The important question is not simply who creates the abundance. It is what becomes scarce relative to it, who controls that scarcity, and whether the market has already paid for the answer.

Scarcity is relative

Scarcity is not an absolute count of units. It is a relationship: demand relative to supply, one capability relative to its complements, one asset relative to alternatives, and one person's time, wealth or status relative to everyone else's.

If machine intelligence scales faster than electricity supply, power becomes scarcer relative to cognition. If synthetic output grows faster than human attention, attention becomes scarcer relative to content. If technological opportunity expands while financing costs rise, capital itself can become the binding constraint. In that regime, the attractive claim may sit in a bank, insurer, commodity producer or cash-rich incumbent rather than in the most obvious technology beneficiary.

Everything is relative. Wealth is measured against purchasing power and other owners. Status is positional. Time cannot be accumulated beyond the hours available to a person. Financial assets trade against the price of capital and every competing use of capital. Excelling in one dimension does not abolish constraint; it increases the value of the next unsolved constraint.

THE NEW SCARCITY

Every solved bottleneck changes the relative price of what remains unsolved. There is no final state of abundance. There is a handoff. Scarcity moves.

The investment universe

AI rewrites economics. We invest in the new scarcity. The Capriole Fund seeks what the world needs next, but cannot produce fast enough, and which investments stand to win. The opportunity set spans three complementary forms of scarcity.



Figure 4. The New Scarcity opportunity set. Digital, competitive and physical scarcity are research categories, not fixed portfolio allocations or assurances of returns.

The historical record is useful because this handoff appears repeatedly. The table below is an interpretive map of where value migrated after major cost collapses; it is not evidence that every named sector earned superior returns.

Technology	What became abundant	Where value migrated
Printing press	Words and books	Literacy, authorship, distribution, reputation
Steam engine	Mechanical force	Coal, rail, factories, capital and coordination
Electricity	Power at the point of use	Grid access, copper, generators and appliances
Internet	Information and communication	Attention, search, networks, trust and identity
Smartphone	Cameras, maps and access	Audience, authenticity, convenience and status
Machine intelligence	Cognition and digital production	Energy, compute, data, trust, physical agency and meaning

Table 1. Historical analogies identify questions to investigate. They do not establish automatic rent capture or predict shareholder returns.

Printing: words became cheap

Movable type reduced the labor needed to reproduce a book and expanded the reach of ideas. Dittmar's research links early adoption of printing with subsequent European city growth.[37] The investment interpretation is broader than the finding: once reproduction becomes easier, selection, distribution and reputation become more important. But the same technology can destroy the income of existing copyists and publishers.

Steam and electricity: power became available everywhere

Steam extended productive power beyond the immediate limits of muscle, wind and water. Electricity made it practical to distribute power differently within a factory. The economic prize was not just a cheaper version of the old task; it included redesigned production. Nordhaus's history of lighting illustrates why the price of a useful service, rather than the price of a particular device, is the relevant unit of progress.[38]

AI has an analogous reorganization problem. Adding a model to an unchanged process may save minutes. Redesigning the process around continuous, machine-assisted work may change its economics. These

benefits need investment in complementary systems and time to diffuse. A capability breakthrough and a productivity breakthrough need not arrive together.

The internet: information became universal

The internet made information easier to publish and distribute. Search, marketplaces and social networks illustrate the value of organizing that abundance. The analogy is not that every intermediary survives. When a new interface reduces discovery costs, it can destroy the incumbent gatekeeper as readily as create the next one.[7]

The smartphone: function became invisible

The smartphone makes the distinction between function and meaning visible. Maps, photographs and accurate time became available through a device people already carried. The marginal price of another reading of the time is effectively zero. That says little about the value of a particular watch as an object, experience or social signal.

Rolex, Porsche, Ferrari and private aviation illustrate purchases that contain more than basic function: provenance, design, privacy, access and control. Their continued appeal is not proof that AI will raise luxury share prices. It is a reason to distinguish the service a product performs from the identity its ownership communicates.[8]

> Your phone made accurate time abundant. It made the Rolex more purely a Rolex.

AI: cognition becomes universal

Our expectation is that the cost of competent writing, coding, design, translation and analysis will continue to fall. Access can broaden even if the leading systems remain expensive. This is valuable for users and uncomfortable for producers whose only defense was that competent work required scarce hours.

That does not imply every industry concentrates. Entry may become easier at the same time that the strongest platforms become larger. The investable hypothesis is narrower: where output is easy to reproduce but a complement remains difficult to replicate, bargaining power migrates toward that complement.

Value is not a conserved pool

There is no rule that a dollar lost by an incumbent reappears as profit for a new bottleneck owner. Competition can pass it to consumers. Better health, cheaper education and more leisure may create enormous welfare without producing a proportionate stream of listed-company earnings. Abundance can destroy scarcity rents as well as create them.

This is the essential difference between optimism about civilization and a portfolio. We can be optimistic about the first while refusing to pay an excessive price for the second.

3. The four laws of abundance

Four propositions organize the thesis. We call them laws as a shorthand for the framework, not as claims of universal or empirically estimated relationships.

Law 1: cost collapse creates demand

A lower price opens uses that were previously uneconomic. Continuous auditing, repeated simulation and personalized assistance may create demand beyond today's human task inventory. But demand must be both useful and funded. Generating more output is not, by itself, evidence of a larger paying market.

Counting only the tasks humans already perform misses some of the opportunity. Counting every technically possible task overstates it. The important frontier lies between them: work that becomes worth purchasing when its complete cost falls below the value of the outcome. That boundary should be measured customer by customer.

A simple distinction disciplines the demand story. Spending equals usage multiplied by price per unit. If price falls 90% and usage rises fivefold, spending falls 50%. If usage rises twentyfold, spending doubles. These are illustrative identities, not forecasts. Resource consumption requires a second calculation using energy or hardware per unit. Greater adoption is not sufficient evidence for a larger supplier profit pool.

Law 2: scarcity moves to the bottleneck

Production is a chain of complements. When one link becomes abundant, bargaining power can migrate toward the link that expands more slowly. If machine researchers generate ten times more promising experiments while laboratory capacity only doubles, experiment access becomes scarcer relative to ideas. If customers can obtain more compute than grid capacity, a power connection can become more valuable than another accelerator. Scarcity moves; the right to profit from it still does not follow automatically.

The original Scarcity Equation is better understood as three gates:

> THE SCARCITY EQUATION: THREE GATES
> 1. VALUE CREATED: does cheaper capability produce a useful surplus?
> 2. VALUE CAPTURED: which owner can retain it, and for how long?
> 3. RETURN PURCHASED: how much of that future is already in the price?
> A decision framework—not an equation with estimated coefficients.

A bottleneck can create value without capturing it. A price-regulated utility may be essential yet unable to charge a scarcity price. A supplier may surrender gains to a concentrated customer. A freely available standard may unlock an industry while earning little itself. We need demand, an effective constraint and a durable claim on the resulting cash flows—not merely strategic importance.

Law 3: abundance raises the value of selection

When production becomes cheap, selecting the correct output can become more valuable. But selection is also a cognitive task. AI will automate parts of search, evaluation and ranking. "Human judgment" is not automatically a permanent moat.

The more durable opportunity is responsibility for a verified outcome: a workflow embedded in a business, a trusted institution willing to stand behind a decision, or a distribution relationship that survives better

competing models. The scarce product may become permission to act, rather than another recommendation about what to do.

> When answers become abundant, permission to act can become the product.

The Delegation Premium

Recommendations are becoming abundant. Authority remains scarce. A business can test ten agents; it may authorize only one system to pay invoices, change production code or move capital. The authorized system occupies a different position because the customer has connected its identity, data, permissions and risk appetite to it.

We call the potential value of that position the Delegation Premium. It depends on earned trust, embedded workflow and credible responsibility for mistakes. Open standards may make it easy to replace; regulation may limit its economics. But it is a more defensible research target than the claim that humans will always be better at choosing.

Law 4: utility saturates; status does not

Technology can give everyone better tools. It cannot give everyone a higher relative rank. Machines can make abundance. They cannot make everyone above average.

Human desire is not a software bug waiting to be fixed. We compare, compete, signal, collect, belong, hoard and seek control. Material abundance changes the objects of competition, not the existence of competition. Once basic function is satisfied, spending migrates toward time, privacy, access, authenticity, identity and positional goods whose value depends partly on who else can possess them.[7][8]

Our expectation is that some digital functions will become cheaper while spending pressure grows on selected scarce experiences, access and ownership. That is a relative-price thesis, not a forecast that all luxury goods appreciate or that aggregate inflation must rise. Interest rates, household income, new supply and changing preferences still matter.

4. Tokens are made of atoms

AI looks like software. Economically, it is an industrial system.

Measure	Evidence	Interpretation
Frontier training compute	About 5× annual growth since 2020 [1]	Historical estimate; not a forecast of economic output.
Pre-training compute efficiency	Compute requirements halved about every eight months in a study covering 2012–2023 [45]	More capability per unit does not automatically increase spend.
GPT-3.5-level inference price	More than 280× lower, Nov 2022– Oct 2024 [2]	A fixed capability benchmark, not the total cost of a job.
Global data-center electricity	485 TWh in 2025 to 950 TWh in 2030 [3]	IEA estimate and central projection; includes non-AI workloads.

Figure 5. Different measures, different clocks. Historical compute and inference-cost trends are not the same thing as future adoption or energy demand. Sources: Epoch AI, Stanford HAI and IEA.[1][2][3]

A token appears on a screen as if it came from nowhere. It did not. It came from a chain of mines, fabs, packaging plants, memory stacks, optical links, substations, generators, cooling systems, data centers and software. Tokens are made of atoms.

Epoch estimates roughly fivefold annual growth in frontier language-model training compute since 2020.[1] Separately, a 2024 study of language-model evaluations from 2012–2023 estimated that the compute needed for a fixed performance level halved about every eight months, with a 95% confidence interval of five to fourteen months.[45] These are estimated historical trends, not a forecast that a current efficiency rate persists. They should not be multiplied mechanically into an economic forecast: architectures, training, inference and deployment are different parts of the system.

The IEA's April 2026 central projection puts global data-center electricity use at about 950 TWh in 2030, up from an estimated 485 TWh in 2025. This covers all data centers, not only AI.[3] Separately, the BIS describes more than US$1 trillion of AI-related capital expenditure planned by the five largest hyperscalers across 2025 and 2026.[4] One number is projected electricity use; the other is a spending plan. Neither is proof of supplier profits.

The financing is part of the technology story. A July 2026 BIS working paper models how competition for a dominant AI position can encourage excessive capacity investment and amplify losses through debt and interconnected financing.[5] Its results are a model, not a measurement of how much today's industry has overbuilt. The warning is useful: rational behavior for a firm trying not to lose a strategic race can still produce poor returns for the industry.

> A technology can be inevitable and an investment can still be terrible. A correct decade does not guarantee a correct entry price.

Moving scarcity: constraint rotation

Multiple constraints can bind at once: accelerators, high-bandwidth memory, packaging, optical links, power, cooling, robotics, deployment capacity or capital. They need not appear in a neat sequence. Training and inference have different requirements; one geography can have spare power while another has an unusable queue of proposed projects. Moving scarcity is a map of changes in relative bargaining power, not a fixed trade itinerary.

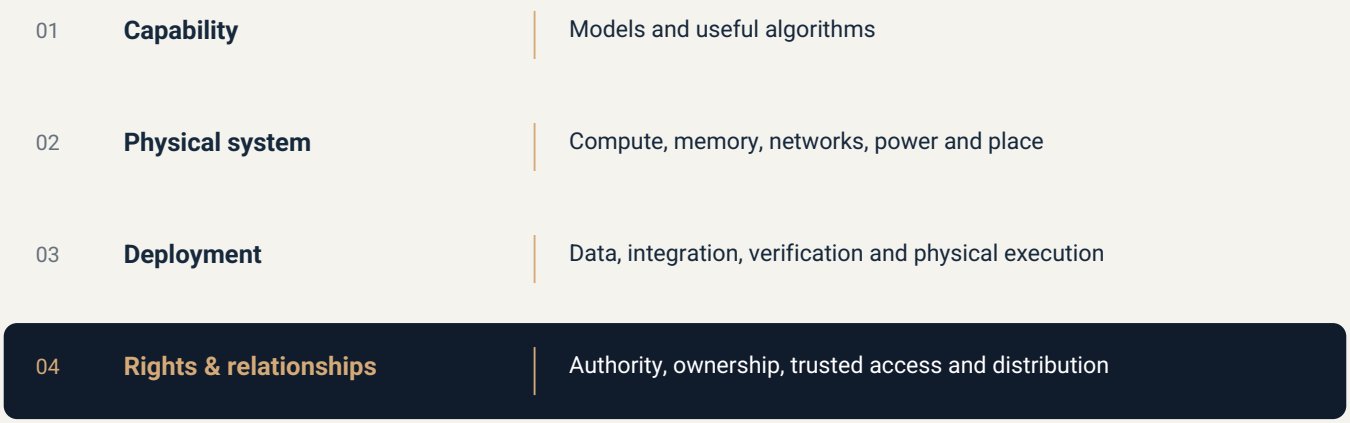


Figure 6. Potential constraint rotation. The sequence is illustrative, not a forecast; several constraints can bind simultaneously and rent capture depends on contracts, competition and regulation.

The investment mistake is to recognize a secular theme and assume its first bottleneck lasts forever. Shortages attract investment and substitution. The same AI that increases demand may help redesign the product around the constraint. A theme can strengthen while a supplier's pricing power weakens.

Stage	Question	Likely constraint owners
Breakthrough	Can it work?	Researchers, architectures, unique data
Buildout	Can it scale?	Chips, memory, networks, power, land
Distribution	Can users adopt it?	Platforms, workflow, integration, trust
Commoditization	Who keeps margin?	Lowest-cost operators and differentiated ecosystems
Embodiment	Can it act in the world?	Robotics, factories, logistics, permits
Meaning	What remains worth paying for?	Brands, status, authenticity, access, scarce assets

Table 2. Stages overlap. A company may control one constraint and be exposed to the resolution of another.

The Scarcity Clock

A shortage has a clock. The clock ends when new capacity, a substitute, customer bargaining or regulation removes the excess return. We call this the Scarcity Clock: the expected duration of a constraint's profitable life. It is a research estimate, not a measured law.

Start with four times: when customer demand becomes real, when usable supply can arrive, when substitution can work and when invested capital must be recovered. The attractive interval is not simply the longest construction delay. It is the period in which demand is funded, the asset is usable, pricing is defensible and the owner can earn more than its cost of capital.

An energized data-center site and a proposed site with an impressive power queue are not the same asset. A supplier with a firm order and a supplier whose customer can cancel have different clocks. A specialized machine may be difficult to build yet lose value rapidly when its next generation arrives. Lead time without customer commitment is not a moat.

For an illustrative project costing 100, five annual cash receipts of 30 have a present value of about 114 at a 10% discount rate, assuming no residual value. Cut the profitable window to three years and the value falls to about 75. The demand theme has not changed; the investment has. Duration belongs beside growth in every valuation.

Own the constraint for longer than the price assumes—not longer than the constraint survives.

AI is both demand shock and supply response

This is what makes the cycle different from a static “picks and shovels” thesis. AI can increase demand for computation and help design cheaper computation. It can create demand for scientific experiments and improve which experiments are worth running. It can raise the need for security while automating part of security itself.

The supplier most visibly exposed to AI growth may therefore face AI-enabled substitution. Our job is to estimate the race between demand expansion and the erosion of the constraint. The investment opportunity lies in that gap, not in assuming a permanent shortage.

Bearings: old components, new demand curves

Intelligence can be copied. Motion has to be manufactured.

A robot’s software can be deployed across a fleet. Its moving parts still have to be made for every machine. If cheaper, more capable intelligence makes robotics useful in more settings, the demand shock can reach a mature industrial component rather than a new software product. Bearings illustrate the transition from digital capability to physical demand.

P Equity Research’s summary of Morgan Stanley research frames bearings as an architecture-agnostic robotics opportunity: exposure to the components used across competing robot designs, rather than a bet on one winning manufacturer. The post attributes an approximately 300-fold expansion in the robot-bearing market by 2050 to Morgan Stanley, and illustrates how bearing content rises with machine complexity.[43]

Measure	Reported in the supplied post	Interpretation
Robot-bearing market	~300× growth through 2050	Long-horizon third-party forecast; not a Capriole base case.
Small quadcopter	8–12 bearings	Illustrative component content; specifications vary.
Humanoid robot	70 or more bearings	More complex movement can increase content per unit.
Individual bearing price	Under US$1 to about US$100	Wide application-dependent range; not a blended robot selling price.

Source summary: P Equity Research, “Morgan Stanley: Bearings,” supplied X-post text.[43] The underlying Morgan Stanley report and forecast assumptions have not been independently reviewed. The 300× figure refers to the reported robotics opportunity, not the entire bearings industry or a forecast of stock returns.

There is also primary-source evidence of the component’s relevance. OpenAI’s January 2026 U.S. hardware manufacturing RFP lists precision bearings—ball, roller and harmonic—among six robotics categories, alongside actuators, drives, gearboxes and motors, magnets, and power electronics. This confirms procurement interest; it is not an endorsement of any bearing producer, a purchase commitment, or validation of the 2050 forecast.[44]

The investment translation

For new robot production, the revenue bridge is straightforward: robot shipments multiplied by bearings per robot, multiplied by the realized average price. A supplier then needs to win a share of that content and retain a margin after manufacturing costs and capacity investment. More robots and more components can increase demand even while the price of each component falls. The net result has to be modeled, not assumed.

The post argues that substitution, insourcing and obsolescence risks are low.[43] Our interpretation is narrower: a common component can reduce dependence on one robot architecture, but it does not remove supplier competition or changes in component specifications. A necessary part is not automatically a scarce part. A scarce part is not automatically a profitable franchise.

The research target is therefore the qualified supply that cannot scale as quickly as funded demand. We would test manufacturing precision, customer qualification, usable production capacity, realized pricing and margins. We would also measure how much of a diversified producer’s earnings actually comes from robotics. A large industry forecast matters little to shareholders if the company has limited exposure, customers capture the savings, or the valuation already assumes the opportunity.

This is the Abundance Paradox in miniature. The breakthrough makes intelligence cheaper; deployment creates demand for movement; the investment question becomes who can supply that movement reliably and retain the economics. The winning technology may be new. The valuable component need not be.

The three currencies

The Three Currencies are a way of asking whether a system has the capacity to act, choose and coordinate. They are not currencies in the monetary sense.

- Energy is the ability to act. It powers computation, transport and physical production.
- Intelligence is the ability to choose. It turns energy and matter into useful outcomes.
- Trust is the ability to coordinate. It establishes identity, ownership, authenticity and settlement.

Cheaper intelligence may increase the volume of valuable actions and the need to verify them. Energy and trust can then become more important complements. AI can also make both more efficient. The balance —not an assumption that all three rise together—determines the economic outcome.

5. The second engine: quantum computing

AI makes thought scalable. Quantum may make parts of reality computable.

Quantum computing is often described badly. It is not a universally faster computer. It will not make email, spreadsheets or ordinary databases magically faster. Its promise is narrower and potentially more important: for selected problems, quantum systems represent and manipulate information in a way classical machines cannot efficiently reproduce.

The scientific case begins with the structure of the problem. Exact representation of a general quantum state can require resources that grow exponentially with system size. Many practical problems nevertheless admit useful classical approximations. Quantum advantage therefore has to be demonstrated for a particular task against the best relevant classical alternative; it cannot be inferred from the word “quantum”.[10]

> AI lowers the cost of asking questions. Quantum may lower the cost of answering some of the hardest ones.

Why the combination matters

Imagine a future laboratory with one million AI researchers. They can read every paper, generate hypotheses continuously, write code and design experiments. Their limiting resource is no longer researcher-hours. It is access to high-value experiments and accurate simulation.

A fault-tolerant quantum computer could become a specialized accelerator within that discovery system. Simulating molecules and materials is a natural candidate because the underlying physics is quantum. Some mathematical problems have known quantum speedups; broad commercial superiority in optimization or machine learning should not be presumed. The opportunity is to expand a useful frontier, not replace ordinary computing.

Evidence already exists in the opposite direction: AI helping build quantum systems. AlphaTensor-Quantum uses reinforcement learning to reduce costly operations in selected quantum circuits.[11] A July 2026 Nature experiment on Willow reported a 3.5-fold improvement in logical-error-rate stability against injected drift using reinforcement-learning control and decoder steering.[36] That is a result for a defined error-correction experiment—not a 3.5-times-faster general-purpose quantum computer.

DIRECTION	WHAT IS SUPPORTED	WHAT STILL NEEDS PROOF
AI to quantum	Circuit optimization and experiment-specific control improvements.	Generalization, scale and sustained system reliability.
Quantum to discovery	Selected algorithms and scientific applications motivate development.	End-to-end useful advantage over the best classical alternative.
Quantum to trust	Post-quantum migration is already a planning priority.	Timing of the threat and reliable implementation across real systems.

Figure 7. The Second Engine separates demonstrated AI-for-quantum work from the conditional opportunity for quantum-enabled discovery. It is not a claim that quantum hardware is necessary for AGI.

The scarce unit is not the qubit

Raw qubit counts make good headlines and poor economics. The valuable unit is a useful, error-corrected computation: logical qubits, low logical error rates, sufficient circuit depth, reliable control and an algorithm that produces an advantage after the full system cost is included.

Google's below-threshold surface-code experiment demonstrated improved logical error suppression as the code grew.[9] IBM's roadmap published in June 2025 targets a 2029 Starling system with 200 logical qubits and capacity for 100 million gates.[12] The first is an experimental result; the second is a vendor target. They answer different questions and should not be presented as equivalent evidence of commercial readiness.

From a laboratory milestone to a scarce service

DARPA's Quantum Benchmarking Initiative makes the economic threshold explicit: utility-scale operation means computational value exceeds cost.[42] That is a better organizing principle than a leaderboard of device sizes. A commercially valuable quantum service must cross a chain of technical and economic gates.

A useful quantum investment case crosses several gates. First, the physical system must preserve information well enough. Second, error correction must support the depth and scale of a relevant algorithm. Third, the complete workflow must beat its best classical alternative on an economically important measure. Finally, the provider must be able to supply the service repeatedly and retain attractive economics.

This changes the unit investors should follow: not the physical qubit, but the cost and throughput of a verified useful result. Preparation, error correction, classical processing, queue time and readout belong in the comparison. A special benchmark advantage can be scientifically important without clearing the commercial gate.

The potentially scarce complement may be the processor, but it could also be control electronics, photonics, fabrication expertise, error-correction software or access to an integrated system. Different architectures distribute the value differently. Picking the quantum sector is not the same as identifying its eventual bottleneck.

There are two distinct investment clocks. Commercial quantum services depend on useful application advantage. Cryptographic migration can require spending before such services are widespread, because

institutions must protect long-lived systems in advance.[13] Security preparation and quantum-computing revenue need not mature in the same year.

Quantum is also a trust reset

Quantum computing is not only a production technology. It is also a demolition technology for cryptographic assumptions. A sufficiently capable fault-tolerant quantum computer would threaten widely used public-key systems. NIST finalized its first three post-quantum cryptography standards in 2024 and now encourages organizations to begin migration.[13]

That migration creates a second scarcity: crypto-agility. The valuable institutions will not merely own stronger algorithms. They will know where vulnerable cryptography exists, control the upgrade path, preserve identity through migration and earn trust while the old foundations are replaced.

Digital assets face the same need for cryptographic migration. A threat to signing keys is not the same as changing the monetary issuance rule, but compromised ownership would still undermine confidence. The relevant investment questions are technical capability, exposed cryptography, governance and the time needed for a safe migration. A date attached to “Q-Day” is not a substitute for those questions.[13]

The error bars are large

Quantum may remain commercially narrow for longer than enthusiasts expect. Engineering requirements differ by architecture; noise, error correction, fabrication, control and system integration all matter.[10] A 2025 systematic review found weak evidence for useful quantum-machine-learning advantage specifically in digital health.[14] That result cautions against premature application claims. It is not a verdict on every quantum algorithm.

That uncertainty does not make the field irrelevant. It makes selectivity essential. The important question is not whether every quantum company succeeds. It is whether fault-tolerant quantum capability becomes a scarce complement to a world with abundant machine researchers. If it does, the strategic value of useful quantum time could be extraordinary.

> **THE SECOND ENGINE**
> AI may multiply the number of researchers. Quantum may multiply the hardest experiments they can run.

6. The synthetic ocean and the human premium

When media becomes infinite, proof becomes valuable.

The internet made publishing abundant. Generative AI makes production abundant. Text, images, video, music, voices, characters, software and advertising can now be generated at industrial scale and personalized to the individual.

The supply of competent digital output can expand far faster than the time people have to consume it. That is valuable to a customer who needs an illustration or a first draft. It is threatening to a producer whose pricing

depends only on the labor needed to make that output. End-to-end quality, rights clearance and responsibility still cost more than the raw generation.

Attention is scarce per person, not a literally fixed global total. Population, screen time and preferences can change. But adding a million plausible pieces of media does not add a million hours to a customer's day. As synthetic supply grows, discovery, credibility and a reason to care become harder parts of the business.

> The value of a photograph was once in taking it. Tomorrow, part of the value will be in proving it happened.

Brand is compressed trust

A brand is not a logo or a font. AI can generate those instantly. A brand is a compressed expectation built from history: what the product will do, how the company behaves, what ownership says about the buyer, and how other people will interpret the signal.

In a low-information world, brand reduced search costs. In a synthetic world, brand also reduces authenticity costs. It becomes a shortcut through an ocean of plausible claims. The strongest brands can therefore become more valuable even as the cost of creative production collapses.

Two buyers: the optimizer and the human

A procurement agent instructed to minimize the verified cost of a commodity task may be less impressed by advertising than a human buyer. It can compare, switch and negotiate repeatedly. A consumer choosing a watch or a holiday may be purchasing precisely what cannot be reduced to the cheapest equivalent function.

This is a useful division, not a claim that machines are immune to brands. An agent may use a reputable supplier because failure is expensive or because its owner values that supplier. The question is why the preference survives better information. Brands based on confusion are vulnerable. Brands based on reliable outcomes or socially recognized identity have a different defense.

When the buyer is a machine, prove the outcome. When the buyer is a human, understand what the outcome means.

Provenance becomes infrastructure

Provenance and truth must be separated. C2PA credentials can bind claims about a file's origin and editing history to the file in a tamper-evident way. They do not certify that its content is true, that the signer is honest, or that a depicted event happened as described.[15] Provenance is evidence to assess—not a machine-issued guarantee of reality.

That distinction creates two complementary needs: authenticate where information came from, then test what it means. Contracts, research data, insurance claims, identity and financial communications all depend on both. A verifiable source can be valuable even when its conclusions are disputed. A beautifully documented falsehood remains a falsehood.

The Rolex Principle

A Rolex is useful here precisely because it is not simply a timekeeping device. Its appeal can include craftsmanship, recognition, provenance and belonging. A competing watch may reproduce the function. It cannot automatically reproduce the social agreement about what the original signifies.

A Porsche or Ferrari offers transport, but also a particular experience and identity. A private jet offers mobility, but also scheduling control, privacy and time. The relevant scarcity is not “cars” or “flight” in general. It is a differentiated outcome for which a particular group of customers is willing and able to pay.

We call this the Rolex Principle: as basic function becomes easier to reproduce, a larger part of willingness-to-pay can depend on identity, relationships and authenticated history. It is an interpretive principle, not a claim that technology caused historical luxury returns. It directs research toward pricing power, repeat demand and disciplined supply—not an indiscriminate luxury allocation.

> Machines can make abundance. They cannot make everyone above average.

Humans remain human

Technology evolves faster than human desire. We will have better tools, more entertainment and cheaper advice. We will still want more status, more autonomy, more beauty, more belonging, more security and more power. AI changes production faster than evolution changes preference.

That does not mean every luxury company wins. Fashion changes, brands decay and supply can be mismanaged. It means the category of positional demand does not disappear simply because functional goods become abundant. In some cases, it strengthens because more people gain the income and tools to compete for goods whose supply cannot expand without destroying their meaning.

Status itself is a relative position; any particular status object is a fragile expression of it. A brand can damage its scarcity by selling too much, lose relevance to a new community, or be displaced by a different signal. The investor should examine price realization, discounting, repeat purchase, distribution discipline and the cost of preserving desirability.

The distinction also explains why abundant synthetic content might benefit live sport or a trusted creator without benefiting all entertainment businesses. The scarce asset is not “human content” in general. It is a relationship or shared occasion that customers cannot obtain from interchangeable output.

Proof of human

A live performance, an authenticated object or direct access to a respected person can retain a premium when synthetic alternatives multiply. AI cannot create additional places in the original event or additional hours in one person’s day. But audiences must still care about those distinctions. A synthetic experience that customers prefer can win; “human-made” is not sufficient by itself.

The Human Premium will not replace the machine economy. It will sit on top of it. The winning brands and institutions will use abundant intelligence behind the scenes while making the customer experience feel more trusted, personal, scarce and meaningful.

7. Finance becomes software

Crypto will not defeat traditional finance. It will disappear into it.

Crypto and conventional finance are often described as competing destinations. A more useful model is convergence in the infrastructure. Bitcoin is a digitally native asset. Stablecoins can carry private fiat-denominated claims. Securities can be issued or represented on a ledger, while derivatives can provide price exposure without ownership. The interfaces are converging; the economic rights are not becoming identical.[18][20][22]

Our five-year hypothesis is that by 2031, many users will choose financial services by price, access and protection rather than by whether the backend is called crypto. The distinction may become less informative at the interface even while law, custody and settlement remain fragmented. This is a forecast to test, not a claim that all markets will merge on a deadline.

The token is not the asset

Tokenization does not create economic value by itself. A token can represent a share, a Treasury bill, a dollar, a concert ticket, a loyalty point—or nothing at all. A token cannot rescue a bad asset. It can make a good asset programmable.

The durable innovation is not that every object needs a coin. It is that ownership instructions and financial claims can become easier for software to use. Transferability, division, auditability and conditional settlement can improve, where the legal structure permits them. The token is an interface to a claim; that interface cannot create rights that the claim does not contain.

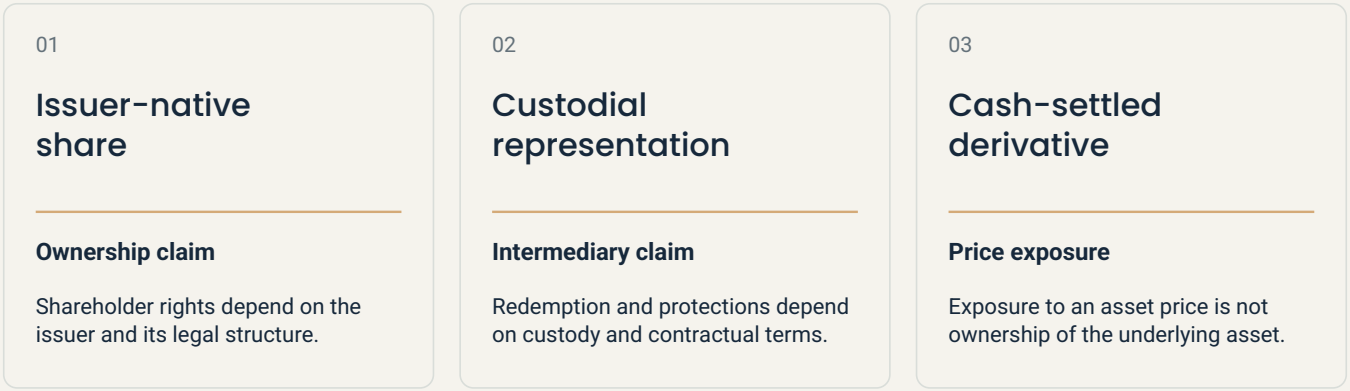


Figure 8. Finance moves from documents and closed ledgers toward programmable claims, while law, rights and trust remain essential.

One screen, three different claims

An issuer-native tokenized share, a custodian-backed representation of a share and a cash-settled derivative can display the same familiar ticker. They are not the same investment. The first may carry shareholder rights through the issuer's legal structure. The second depends on an intermediary and its contractual terms. The third supplies price exposure without ownership of the company.

The diligence must look through the interface: who issued the claim, who holds the underlying asset, what can be redeemed, how dividends and votes are treated, and what survives an insolvency. A token can improve transfer instructions while weakening other protections. "Backed" is the start of the questions, not the end.

This is why tokenization belongs in the paper's trust thesis. The winning infrastructure must make assets easier to use without making ownership harder to understand.

The convergence is already visible

In February 2026, Kraken reported more than US$25 billion in cumulative xStocks activity within eight months of launch. Its definition combines centralized trading, decentralized trading, minting and redemption. The same announcement reported roughly US$225 million of aggregate assets under management.[18] Turnover and the outstanding asset base differ by orders of magnitude. Neither should be mistaken for the other.

The London Stock Exchange announced plans in September 2026 for rights-preserving tokenized UK equities and a proposed xStocks listing on LSE 24 in 2027, subject to approvals.[21] NYSE's January 2026 announcement described a proposed tokenized-securities platform, including continuous trading and stablecoin funding, for which it would seek regulatory approval.[22] These are significant commitments of intent—not evidence that those venues are already operating at scale.

Alongside those plans, Crypto.com announced tokenized-stock access in the EEA in August 2026, while Bybit offers TradFi-linked perpetual contracts.[19][20] These products are not interchangeable with directly owning a listed share. The important observation is that conventional exposures are entering crypto interfaces as traditional infrastructure providers explore tokenization. The evidence does not establish that equities and commodities dominate all crypto trading.

Stablecoins are machine cash

Bitcoin and stablecoins illustrate two different functions: a digitally native asset without an issuer's promise to redeem at par, and a privately issued claim designed to track a currency. For reserve-backed stablecoins, reserve quality, redemption terms and the identity of the issuer matter. Neither the word "stable" nor an onchain balance guarantees that a holder has the same protections as a bank depositor.[39]

Visa reported a US$7 billion annualized stablecoin-settlement run rate for its pilot in April 2026.[23] Circle announced policy-controlled agent wallets and USDC transfers as small as US$0.000001 in May.[24] One is a reported rate of activity; the other is a product capability. Together they show development of the infrastructure, not the current size of autonomous commerce.

An agent that recommends a purchase is different from a system authorized to spend. Useful machine-to-machine finance needs a budget, identity, permitted counterparties and a record of who bears responsibility. The machine economy requires a Machine Balance Sheet—but not necessarily a new token for every machine.

The Machine Balance Sheet

A company has accounts, ownership records, credit arrangements and expense policies. A software agent acting for it needs delegated versions of the same functions: limits, identities, permissions, working capital and audit records. The economic actor remains the accountable person or institution unless the applicable law provides otherwise. Autonomy of execution is not the absence of an owner.

Function	Conventional arrangements	More programmable arrangements
Cash	Deposits and existing payment APIs	Stablecoins, tokenized deposits or other compliant payment rails
Ownership	Legal records plus broker/custodian ledgers	Claims connected to machine-readable transfer instructions
Trading	Electronic and human orders; venue-specific hours	API routing across a wider set of continuous markets
Settlement	Netting, batch or real-time processes	Atomic or near-instant settlement where suitable
Identity	Account-level identification and authority	Agent identity linked to delegated authority
Risk	Policy, software controls and human approval	Executable limits with auditability and revocation

Table 3. Finance becomes machine-readable when the asset, authority, payment and risk controls can all be expressed in software.

Programmability is not the same as liquidity

A market can accept orders at any hour without offering dependable depth at every hour. A stablecoin can transfer while its bank reserve cannot. A tokenized stock can trade while the underlying market is closed. The result may be useful access, but also wider spreads, basis risk and a need for market-making capital.

Faster settlement has its own trade-off. Settling each transaction gross and immediately can reduce the benefit of netting and increase intraday funding needs.[32] Atomic settlement may reduce one risk while making collateral management more demanding. The objective is better total market structure, not speed in isolation.

Agent payments are a competitive opportunity for stablecoins, tokenized deposits and conventional APIs. Users will care about cost, reach, finality, privacy, recourse and reliable execution. The thesis does not require every transaction to use a public blockchain. It requires financial systems to become easier for authorized software to operate.

The end state

A user may eventually see one interface across a bank, broker, exchange and custodian. That does not make those functions economically identical or legally interchangeable. Balance sheets, insolvency rights, settlement assets and accountability remain essential. The convergence thesis is about programmable finance, not the disappearance of financial institutions.

Value may accrue to networks and institutions that combine deep liquidity, enforceable rights, dependable custody, identity and distribution. Open chains, permissioned systems and conventional payment APIs can compete or interoperate. We should follow the best economics for each use, rather than assume one technical architecture must capture every market.[39]

This also changes the investment map. A scarce economic exposure may appear as an equity, commodity, digital asset, derivative, debt instrument or private security. A serious Intelligence Age strategy should be instrument-agnostic and claim-obsessed. Artificial asset-class boundaries will increasingly hide the opportunity.

8. The Speculation Reflex

When the ladder slows, more people reach for the elevator.

Abundant intelligence can make society richer while making many individuals feel less economically secure. Productivity can rise, services can become cheaper and corporate profits can grow while the market price of routine cognitive labor falls. Aggregate abundance and personal displacement are not opposites.

This matters because work is not only a source of consumption. It is the default ladder through which most people expect to gain ownership, status and control. If that ladder lengthens at the same time that power-law fortunes become more visible, speculation can begin to feel less like entertainment and more like an alternative career path.

The Ownership Gap

Cheap services and accessible ownership are different forms of abundance. A household may get better software, entertainment and advice while remaining unable to acquire the assets generating the new income. Falling consumption costs do not automatically close the distance between wages and ownership.

We call that possible divergence the Ownership Gap. It connects the machine and human sides of the thesis: abundant capability raises what the economy can produce, but access to the claims on that production may remain scarce. Its size depends on competition, wages, entrepreneurship, savings, taxes and how widely capital is held. It is not predetermined by the technology.

The labor shock starts on the screen

The first large labor shock does not require a humanoid robot. It begins wherever work enters as digital information and leaves as digital information. If a worker spends most of the day receiving instructions from a machine, transforming symbols and returning the result to another machine, an agent can increasingly insert itself into the loop.

The ILO's 2025 task-based assessment places roughly one in four workers in occupations with some generative-AI exposure, but only 3.3% in its highest-exposure category. It judges transformation more likely than complete replacement.[25] The OECD similarly measures task exposure rather than a count of jobs certain to disappear.[26] These are maps of where work could change, not unemployment forecasts.

Anthropic's March 2026 study makes the distinction particularly clear. It finds a large gap between theoretical task exposure and observed use of Claude, no systematic increase in unemployment in highly exposed occupations since late 2022, and only suggestive evidence of slower hiring of younger workers.[27] The study uses a particular platform's activity and early labor data; neither optimism nor catastrophe is established by it.

The Wealth Jump Economy

Industrial transitions can create new work, but people experience the transition through their own income, obligations and time. The relevant risk is not only unemployment. It includes weaker entry-level hiring, slower wage progression, underemployment and a loss of bargaining power in tasks that once supported a career.

If gains accrue faster to owners of productive assets than to workers selling routine skills, access to ownership can feel further away even while some services become cheaper. Many will adapt by learning, building or saving. Some may seek a discontinuous jump through speculation. We call that possible response the Wealth Jump Economy. It is a behavioral hypothesis, not a judgment about every displaced worker.

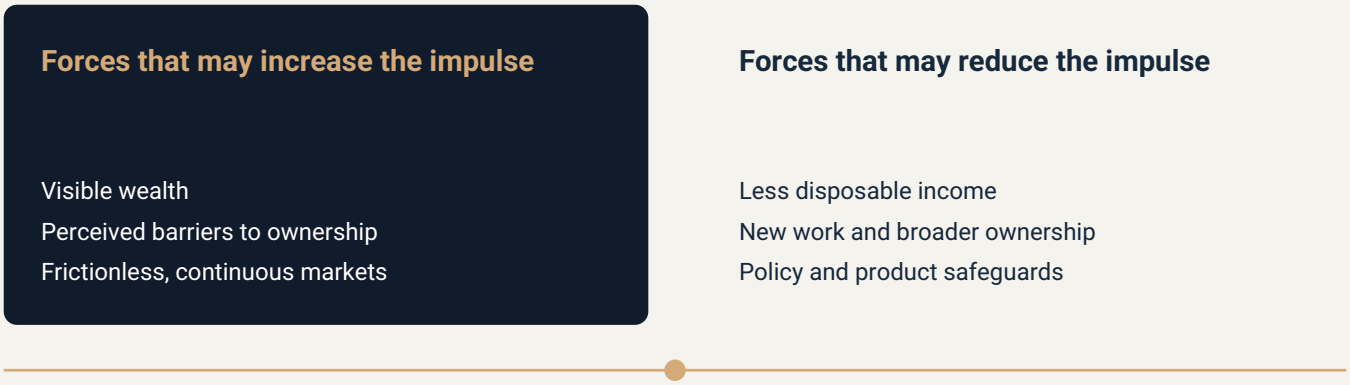


Figure 9. The Speculation Reflex is a conditional social thesis, not a prediction that every displaced worker will gamble.

The mechanism has opposing forces. Insecurity and greater leisure can raise the attraction of a lottery-like payoff. Lower income can reduce the money available to risk. Broader ownership and new work can restore a credible path to advancement. Public policy and product design can change both sides. The net effect is an empirical question, not a one-way consequence of automation.

The casino is becoming an API

Digital distribution makes betting and short-horizon trading easier to access. A phone can support continuous markets, small stakes and rapid feedback. AI could personalize these products further. Accessibility can expand participation, but advertised ease should not be confused with better odds for the customer or sustainable profit for the operator.

The American Gaming Association reports US$166.94 billion of legal U.S. sports-betting handle in 2025, up 11%, and US$16.96 billion of sportsbook revenue, up 22.8%. Its reported iGaming revenue rose 27.6%.[28] Handle is the amount wagered, including money staked repeatedly; it is not operator revenue or unique capital invested. These figures cover legal commercial markets, not all gambling.

OCC separately cleared 15.3 billion contracts in 2025, versus 12 billion in 2024.[29] That is evidence of expanding derivatives activity, not a measure of gambling or proof of an AI-driven change in behavior. Legalization, mobile distribution, market volatility and financial-product design are competing explanations. The causal leap from growing activity to AI displacement must not be smuggled into the data.

Hyperliquid, Pump and the claim on the fees

Hyperliquid illustrates the continuous-market side of the thesis. Its published fee documentation describes allocations to community mechanisms, including the assistance fund, HLP and deployers.[30] Pump illustrates rapid asset creation and the monetization of attention. Its fee schedule distinguishes creator, liquidity-provider and protocol fees.[31] Those categories matter: gross trading fees are not automatically the income of the protocol or of a tokenholder.

Pump's token page also describes a buyback-and-burn program.[41] A program can transmit economic value without making a token equivalent to equity. Its durability, governance, supply unlocks and relationship to holder rights must be assessed. The same test applies to any protocol whose activity is used to justify a token valuation.

The research bridge is therefore explicit: credible volume, less rebates and incentives, becomes retained revenue; retained revenue, less operating and capital needs, becomes distributable value; only a defensible

mechanism connects that value to the instrument being purchased. Growth in the first number does not guarantee growth in the last.

The Attention-to-Asset Pipeline

The time between an online narrative and an instrument that trades on it can be short. Token-launch infrastructure and continuous derivatives make it possible to financialize attention without the traditional sequence of a corporate listing.[20][31] The mechanism is not unique to crypto, but programmable issuance makes it easier to repeat.

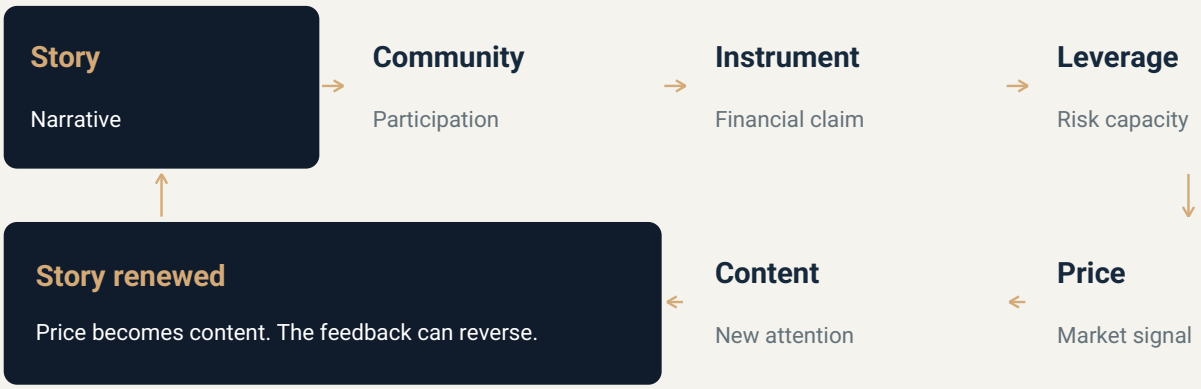


Figure 10. The Attention-to-Asset Pipeline turns narrative velocity into order flow, and order flow into more narrative.

Price becomes content. Content recruits participants. Participants may create liquidity, which attracts more attention. The feedback can run in reverse just as quickly. A liquid market is not evidence of a productive asset; it may only show that participants agree to trade a narrative for now.

Layer	What becomes abundant	What remains scarce
Narrative	AI-generated stories, analysis and predictions	Attention, credibility and cultural timing
Issuance	Tokens and contracts created on demand	Trusted listing, legal rights and distribution
Access	Global 24/7 interfaces	Liquidity and reliable market making
Leverage	Fractional and convex exposure	Collateral, risk capacity and survival
Settlement	Stablecoin and onchain execution	Identity, custody, compliance and finality
Retention	An infinite product catalog	Habit, trust, community and brand

Table 4. The modern speculation stack converts abundant narrative into financial turnover, but the rent still belongs to scarce trust and liquidity.

Who captures the economics?

The investment conclusion is not to own every casino chip. It is to identify who owns the scarce rails: liquidity, settlement, identity, custody, compliance, risk management, distribution and trusted access. A high-volume venue can still be hacked, regulated, copied or displaced. A platform token can rise without conveying an enforceable claim on cash flow. The house itself can be a poor investment when the rent is not legally or economically captured.

The strongest businesses will make financial participation easier while keeping the system credible. In a world where every agent can trade and every story can become an asset, trust becomes more valuable—not less.

The social cost is part of the valuation

A Finnish register-linked study found an association between unemployment, social-benefit receipt and problem or at-risk gambling.[33] It does not establish that unemployment caused gambling, still less that AI will cause it. The harms are nevertheless part of the economic assessment: vulnerable participation can produce unsustainable revenue and invite product restrictions, litigation, taxation and reputational damage.

A profitable demand curve can be socially destructive. It can also be temporary. The task is to understand the mechanism without romanticizing it, and to price the legal, ethical and reputational risk with the same seriousness as revenue growth.

9. Will power laws get steeper?

Abundance raises the floor and stretches the ceiling.

Rosen's superstar framework explains how scale and imperfect substitution can concentrate rewards.[7] The Intelligence Age could reinforce that mechanism where the best product can reach a global market and users strongly prefer it to alternatives. This is a conditional economic argument; we are not claiming that every sector's return distribution has already become more unequal.

The feedback can be powerful: more customers finance better systems, which attract more customers. Yet a small model improvement need not create a lasting advantage. If buyers can switch instantly or use several providers, the gain may accrue to customers instead. Distribution, integration, rights and the cost of switching determine which feedback dominates.

Abundance can therefore lower the barrier to starting a company while increasing the scale of the most successful networks. More entrants and more concentration at the top are not contradictory. The investment question is whether competition dissipates the surplus or whether a firm retains it after the technology spreads.

> Abundance can widen participation and concentrate the rewards at the same time.

Power law does not mean permanent winner

Agents may weaken advantages based on customers failing to compare prices or suppliers. They may also strengthen a trusted system authorized to manage the comparison and transaction. Our hypothesis is that some value moves from being the product customers browse to being the system customers permit to act. That is the Delegation Premium—not a guarantee that any existing platform owns it.

Power laws are therefore not a license to own the largest companies at any price. They are an argument for selectivity and continuous re-underwriting. The winner is the firm that controls the relevant constraint now and can defend that position as the constraint moves.

The politics of concentrated abundance

If a small number of firms control essential intelligence, physical capacity and delegated authority, economics and politics become difficult to separate. Customers may demand interoperability. Governments may regulate prices, tax rents or require local infrastructure. National security can favor a domestic supplier while increasing the capital it must commit.

Concentration is therefore not only a potential source of returns. It creates a response that may constrain them. An investor should underwrite who is allowed to earn the rent, not merely who has the technology to charge it.

The middle is the danger zone

Businesses most exposed are those whose output, customer access and credibility are all easy to reproduce. A product can become much better while its pricing power deteriorates. Conversely, an application can thrive when it owns a workflow, scarce proprietary information or a customer outcome that a generic model cannot deliver alone.

The strongest businesses are likely to cluster at two ends: the scarce inputs machines must consume and the scarce outputs humans continue to value. This is the Abundance Barbell.

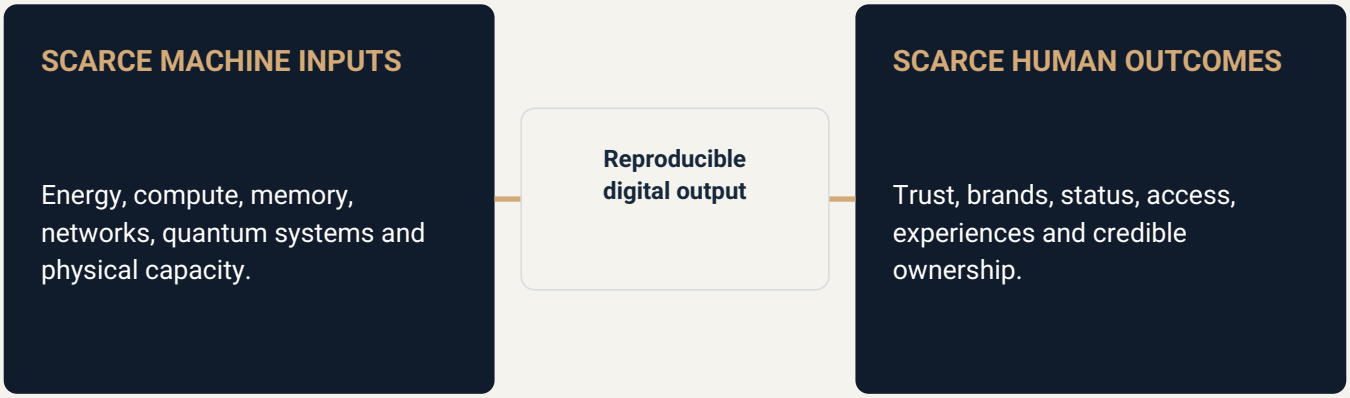


Figure 11. The Abundance Barbell is a research lens, not an asset-allocation rule. Applications in the middle can win when they retain defensible rights, distribution or workflow ownership.

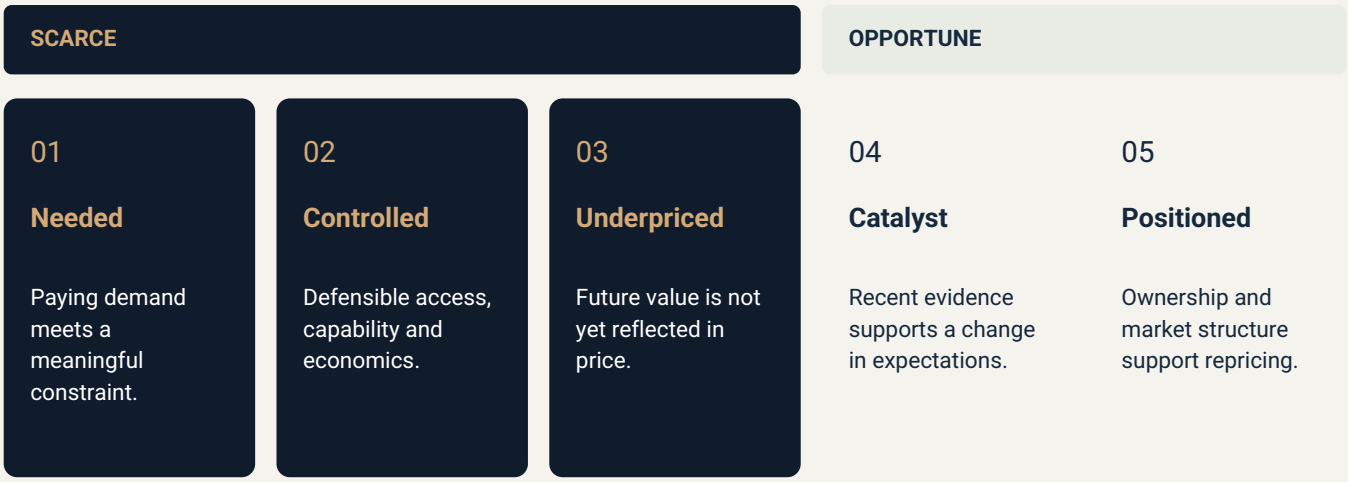
10. The Scarcity Map

Find where scarcity moves before capital does.

The Abundance Paradox is not an AI index, a small-cap screen, a crypto mandate, a value style or a momentum rule. Those are possible expressions. The underlying regime is Moving Scarcity: technology, policy, capital and human behavior repeatedly change which constraint has bargaining power. A useful framework must tell us where to look, when the scarcity is strengthening, and when it is being competed away.

Finding scarce assets at the opportune time

Proprietary systematic strategies and fundamental research identify opportunities. Valuation, catalysts and market evidence determine allocation. We look for new scarcities the market has mispriced. Each opportunity is assessed through five complementary lenses:



Five complementary lenses. One decision: scarce assets at the opportune time.

Figure 12. Five complementary lenses. Needed, Controlled and Underpriced assess scarcity; Catalyst and Positioned assess whether the opportunity is timely. These investment lenses complement the value-created, value-captured and return-purchased distinctions developed earlier.

A living Scarcity Map

The Scarcity Map is a research framework for following four states: scarcity emerging, scarcity accelerating, scarcity peaking and scarcity easing. It is updated as evidence changes. The map is not a portfolio sheet. It is a hypothesis sheet showing where relative bargaining power appears to be moving and what evidence could reverse that view.

For each candidate scarcity, track four things: funded demand, supply elasticity, rent capture and expectations. Then add the Scarcity Clock: how long the advantage can plausibly survive before capacity, substitution, regulation or a change in the price of capital resolves it.

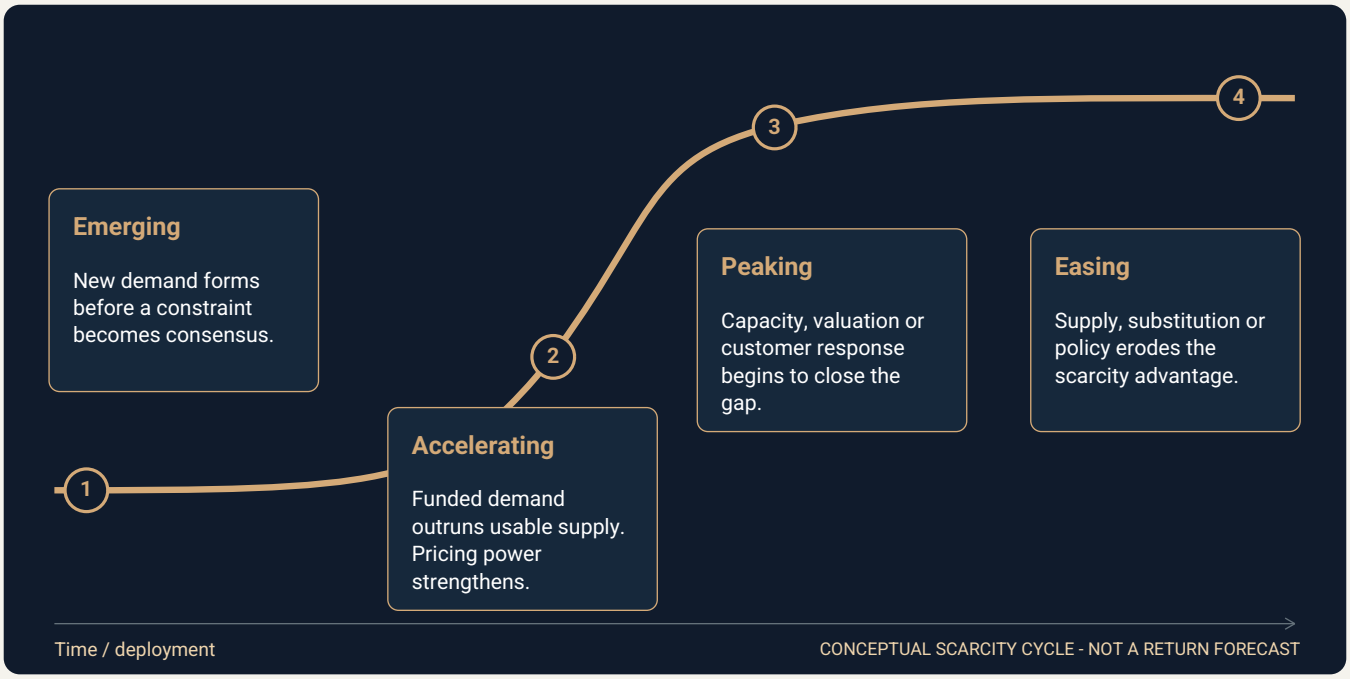


Figure 13. The Scarcity Map. Emerging, Accelerating, Peaking and Easing describe changing opportunity states. The S-curve is conceptual: it is not a return forecast, measured adoption series or a claim that scarcity rents rise monotonically. The Scarcity Clock asks how long the advantage can last before capacity, substitution or regulation changes the economics.

The map should be allowed to rotate radically. A period dominated by AI infrastructure can give way to energy, robotics, biotech, banks, commodities, digital finance, premium brands or cash. The consistency is not the sector. The consistency is the search for the binding constraint and the claim that captures it.

From research to allocation

Proprietary AI, systematic strategies and fundamental research identify opportunities. Valuation, catalysts and portfolio impact determine allocation. The process connects research to implementation and continuous review, rather than treating a thematic label as an investment decision.

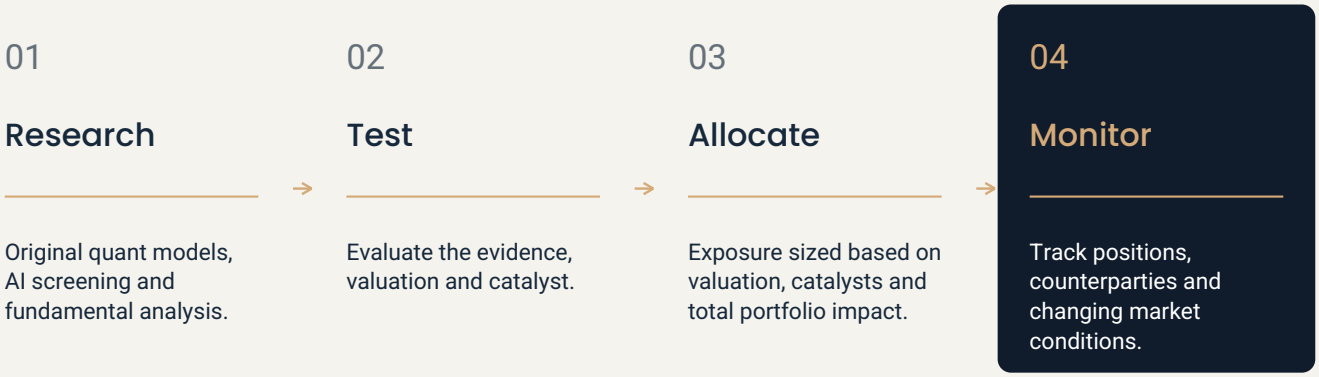


Figure 14. Research, Test, Allocate, Monitor. Models support opportunity selection and sizing; human oversight and ongoing risk assessment remain integral.

The detailed research checklist

Step	Research task	Required evidence
1	Define the abundance	Specify the task, quality level, full cost and measured change.
2	Trace the demand	Identify paying users and the complements they will consume.
3	Locate the constraint	Measure usable capacity, qualification, location and supply response.
4	Establish the claim	Identify rights, pricing power, competition and who retains cash flow.
5	Set the Scarcity Clock	Estimate when capacity, substitution or policy erodes excess returns.
6	Price expectations	Model duration, capital needs, dilution and downside—not just sales growth.
7	Seek confirmation	Track fundamentals, liquidity and market evidence without treating price as proof.
8	Size the opportunity	Require a survivable downside and a better use of risk capital than alternatives.

Three worked research tests

An AI infrastructure supplier. Establish the funded demand for its product, the time to deliver qualified capacity, and the extent to which contracts protect pricing. Then model customers switching, a competing architecture arriving and growth normalizing. The thesis fails if returns on new capital fall below the required rate even while revenue keeps expanding.

A premium human brand. Identify the outcome people purchase beyond function. Test pricing against unit volumes, discounting, retention and distribution. Do not infer desirability from a high price alone. The thesis fails if additional supply, changing tastes or synthetic alternatives erode willingness-to-pay faster than operating efficiency improves.

A programmable financial market. Distinguish underlying assets from wrappers and usage from ownership value. Examine customer retention, market depth, fee rebates, security, legal rights and token dilution. The thesis fails if volume depends on subsidies or the instrument lacks a durable claim on the economics.

These are templates, not security recommendations. Their common structure is the contribution of this paper: the same test of scarcity, duration and value capture can be applied across assets without pretending those assets carry the same rights or risks.

The Copy Cost Test

Every business in the digital economy should now face a brutal question: what happens when a competent agent can copy the product?

01	Can an agent replicate the product at near-zero cost?
02	Can the distribution and customer relationship be replicated?
03	Does the business own a physical, legal, network or data constraint?
04	Does trust, provenance, brand or status materially affect willingness to pay?
05	Does greater intelligence increase demand for the product - or destroy it?
06	After all of that, what expectations are already in the price?

A strategic constraint can still be a poor investment at the wrong price.

Figure 15. The Copy Cost Test is an analytical screen, not an empirical scoring model. Low replication cost matters only alongside distribution, rights, competition and valuation.

A low copy cost does not automatically destroy a business. Distribution, proprietary data, regulation, network effects, integration, trust and brand can still protect it. But a product whose output, distribution and credibility are all reproducible is not a moat. It is a temporary feature.

Two ends of one opportunity set

The barbell describes two kinds of scarcity, not two mandatory portfolios. Capriole's investment universe spans global equities, Bitcoin, digital assets and critical commodities. The broader research thesis also considers how economic claims can be expressed through different instruments. The underlying exposure, legal rights and implementation risks come first; the paper does not expand the Fund's mandate or prescribe asset-class weights.

- Machine-side scarcity: energy, compute, memory, networking, semiconductor equipment, power infrastructure, quantum systems, robotics, digital scarcity, programmable settlement and other scarce physical or financial inputs.
- Human-side scarcity: trusted platforms, provenance, security, premium brands, scarce experiences, prime access, distinctive intellectual property and credible ownership systems.

The middle is not uninvestable. Some application businesses will create extraordinary value. But they must prove that abundant intelligence increases their demand faster than it commoditizes their product, and that the company—not its upstream model supplier or downstream customer—captures the gain.

The portfolio implication

A skewed distribution of realized returns does not prove that investors can identify the winners in advance. It supports both the case for broad diversification and the case for selective concentration when there is genuine evidence of an edge. Capriole's approach must earn the latter through research, sizing and ongoing challenge—not through the slogan alone.

Concentration should be measured through shared economic risks, not a count of ticker symbols or strategy names. Several models and assets may all depend on the same liquidity cycle, capital-spending boom or funding regime. Independent quantitative return streams, fundamental conviction, cash and hedges can coexist. No fixed position count, asset allocation or rebalancing calendar follows from this paper because scarcity itself does not respect those boundaries.

> A position should not enter the portfolio because it is interesting. It should enter because it can matter.

Sizing must also reflect financing and liquidity. A thesis with a five-year payoff can fail as an investment if it depends on collateral that can be liquidated next week. Multiple strategies should be assessed on their combined exposures, including common customers, funding sources, counterparties and downside scenarios. Staying invested by choice is different from being unable to exit.

Our process combines proprietary models, AI, systematic strategies and fundamental research. Valuation, catalysts and portfolio impact determine allocation. Experienced judgment is layered over quantitative systems. The framework informs which opportunities deserve investigation; it does not insist that each trade be discretionary or that a fixed number of holdings expresses conviction. Cash is available when price, evidence or implementation is unattractive.

Our systems continuously monitor portfolio positions, market conditions and trading counterparties. Automated alerts, calls and risk controls work alongside human oversight, with exposure assessed across strategies and asset classes. Guardian monitors on-chain and counterparty risks in real time. Independent administration, annual auditing and controlled asset-transfer processes support the Fund's operating framework.

Research as a living system

In an agentic world, generic research becomes abundant too. The advantage does not come from producing the most summaries. It comes from maintaining a live Scarcity Map, asking better questions, and recording which proposed scarcities strengthened, disappeared or moved somewhere else.

The research system should connect dependencies, supply, pricing power, customer behavior, scientific milestones, capital costs, valuation and market structure. Agents can help maintain the map, but the map must be redrawn when the evidence changes rather than defended because a previous thesis became part of the brand.

> HUMAN CONVICTION. MACHINE DISCIPLINE. Use quantitative strategies and human judgment where each has an edge. Keep responsibility, risk limits and evidence explicit.

11. What could break the thesis

A useful thesis must explain what would prove it wrong.

The Abundance Paradox is a framework, not a promise. Several outcomes could materially weaken its investable conclusions.

AI capability or adoption could slow

Capability progress could slow, or useful deployments could require more context, oversight and reorganization than expected. The risk is not only that models stop improving. It is that benchmark gains do not translate into sufficiently cheaper, reliable outcomes. Current infrastructure valuations may assume a faster conversion than businesses achieve.[34][35]

Efficiency could outrun demand

The Jevons-like response is not a law of nature. Better hardware and algorithms may reduce aggregate resource demand in some workloads. Customers may satisfy their needs with smaller models. Local power bottlenecks may ease. The IEA's own scenarios show meaningful sensitivity to efficiency and adoption assumptions.[3]

The infrastructure boom could bust

A real technological revolution can contain an investment bust. Debt, aggressive capacity commitments and dependence on a few customers magnify the effect of disappointments. The sign to watch is not just slower AI growth; it is an unfavorable gap between delivered capacity, paying usage and cash conversion.[5]

The profit pool could bypass the obvious suppliers

More compute demand does not guarantee durable semiconductor margins. More AI use does not guarantee model-provider profits. Customers can multi-source, open systems can commoditize capability, and regulators can intervene. The economic winner may be a less visible provider of power, workflow, security, distribution or end-customer value.

Quantum could take much longer

Useful fault-tolerant systems may take longer, or the best classical alternative may improve enough to postpone commercial advantage. Attractive device metrics can coexist with poor application economics. The appropriate response is milestone-based exposure, not treating a vendor's calendar as a certainty.[10][12][14]

Human preferences could move differently

Some premium brands will be diluted by synthetic imitation rather than strengthened. Younger consumers may value access over ownership. Virtual status goods may substitute for physical ones. Provenance systems may fail to gain adoption. The Human Premium is a category thesis, not a claim that every incumbent luxury brand is permanent.

Finance could remain fragmented

Tokenized assets may remain trapped in jurisdictional, custodial and technological silos. Some wrappers provide weaker rights than the securities they reference. Bridges, issuers and custodians can fail. Regulation may force closed networks and preserve market-hour or settlement frictions. Finance becoming programmable does not guarantee that open crypto rails capture the value.

The Speculation Reflex could fail—or be regulated away

AI may broaden entrepreneurship and ownership, while cheaper services and new occupations offset disruption. Insecurity may reduce disposable money rather than increase betting. Product restrictions and consumer fatigue can break the link between attention and platform revenue. An expanding speculative market does not validate the labor mechanism.

Valuation can overwhelm truth

The most dangerous version of this thesis is one that is directionally correct and universally believed. If future growth is already capitalized at an extreme rate, a great business can deliver poor returns. The discipline is to separate technological inevitability, corporate value creation and shareholder return. They are three different questions.

Dimension	Measures	Question
Cost of intelligence	Cost per accepted outcome; review time; task success and latency	Is complete, reliable work becoming cheaper?
Demand response	Paid usage; retention; utilization; revenue net of incentives	Is efficiency expanding demand or reducing spend?
Physical constraints	Power queues; HBM/network supply; data-center utilization	Where is scarcity actually binding now?
Rent capture	Contract pricing; free cash flow; incremental returns on capital; dilution	Who keeps the economics?
Quantum progress	Logical errors; useful depth; full-workflow advantage; repeat paid use	Is fault-tolerant utility moving closer?
Trust migration	PQC adoption; provenance usage; fraud losses	Are trusted systems adopted and failures reduced—not just credentials issued?
Human premium	Luxury pricing; live experience demand; creator concentration	Is status and authenticity gaining share of wallet?
Financial convergence	Real assets outstanding; adjusted settlement; redemption rights and use	Are financial claims moving to programmable rails without weakening ownership?
Speculation reflex	Cohort-level betting and trading; income; access; labor outcomes	Does the proposed labor mechanism survive alternative explanations?

Table 5. A monitoring framework, not a backtested predictor. Use consistent definitions, fixed comparison periods and the full set of positive and negative evidence.

> The thesis should be bold enough to act on — and measurable enough to abandon.

Predictions that can fail

A thesis about “moving scarcity” can explain almost anything after the event unless its predictions are specified in advance. The following research commitments turn the framework into propositions that can disappoint. The dates are review horizons, not promised milestones or a backtested forecast.

The forecast register below is original Capriole analysis. Published indicators provide context; they do not validate the predictions.

Hypothesis / review horizon	Evidence that would support it	What would weaken it
Trusted work, 2027–28	Fixed-workflow cost per accepted outcome falls after review, integration and errors.	Benchmarks improve but complete workflow costs and reliability do not.
Constraint rotation, 2027–29	Contract prices, utilization and cash returns shift toward newly limiting complements.	Broad supply expansion removes excess returns before owners recover capital.
Delegation Premium, 2027–29	Retention and pricing are stronger for authorized, embedded workflows than interchangeable output.	Open standards and easy switching prevent any provider retaining the surplus.
Human Premium, 2027–30	Selected authenticated brands and live experiences sustain pricing after controlling for demand conditions.	Synthetic alternatives or changing preferences erode willingness-to-pay.
Quantum, annual to 2030	Independent full-workflow benchmarks demonstrate useful advantage and repeat paid use.	Only device milestones or narrow tests improve; classical methods remain superior in practice.
Financial convergence, to 2031	Repeat settlement, real assets outstanding and enforceable rights grow across compatible systems.	Activity is mainly churn, weak wrappers or isolated pilots without durable adoption.
Speculation Reflex, 2027–30	Comparable exposed cohorts show persistent behavioral change after access, income and market effects are considered.	No independent relationship; lower income or broader opportunity offsets insecurity.

Conclusion — Moving Scarcity

The important question is no longer only what machines can do. It is where the success of those machines moves scarcity next.

Not AI. Not small caps. Not crypto. Not value. Not momentum. Those are sectors, wrappers and styles. The deeper investment regime is Moving Scarcity.

Technology does not eliminate scarcity. It moves it. Every breakthrough changes the relative value of its complements. More intelligence can make power scarce. More content can make attention scarce. More automation can make ownership and trusted authority scarce. A boom can make capital scarce when rates rise. Solving one constraint raises the shadow price of another.

This is not an edge case of economics. It is much of the world order. Value is relative to alternatives. Wealth is relative to purchasing power and ownership. Time is relative to opportunity. Status is relative to other people. Financial assets are priced relative to capital. Abundance in one dimension therefore does not end competition; it changes where competition concentrates.

Past technologies expanded the supply of useful capabilities while reorganizing the economy around new constraints. Intelligence could extend that process into invention, coordination and decisions. The scale may be extraordinary. The distribution of the gains is not predetermined.

Cheaper cognition can enlarge the world's possibilities and erode the income of particular producers at the same time. We should expect both new businesses and stranded assumptions. A rising technological ceiling is not a rising floor under every asset price.

But every token will still require atoms. Every autonomous system will still require energy. Every synthetic claim will still require trust. Every scientific hypothesis will still confront the difficulty of nature. Every person will still have only one life, one reputation and twenty-four hours in a day.

Quantum belongs in this map as a potential specialized complement to abundant machine research, and as a reason to modernize cryptographic trust. Useful computational advantage, access to it and the price of that access remain questions to resolve. Research progress is real; the investment payoff is not automatic.

People will still seek time, autonomy, security, beauty, belonging and status. AI can reproduce a style, but not instantly recreate a shared history. It can offer more possibilities, but cannot give everyone the same exclusive place. The Human Premium depends on what people continue to recognize and value—not simply on what a machine cannot manufacture.

> THE FINAL IMPLICATION The most attractive opportunities combine expanding demand, a defensible claim on scarce value, sufficient duration and a price that leaves room for error.

In finance, the distinction between a conventional asset and its digital interface should become less decisive. The future may combine tokenized deposits, stablecoins, securities, digital stores of value and ordinary financial institutions in unfamiliar arrangements. Instrument-agnostic investing will matter. So will obsessive attention to the claim underneath.

The transition also has a political and social balance sheet. Broad access to powerful tools can coexist with narrow ownership of the assets those tools depend on. If the gains are shared poorly, insecurity and

speculation may grow. That possibility is a risk to understand, not an outcome to celebrate or an excuse to confuse turnover with durable value.

Our conclusion is neither “buy AI” nor “buy whatever is scarce”. It is to find where scarcity is moving before capital fully follows: identify the cost collapse, trace the new demand, locate the slowest response, determine who owns a durable claim on the economics, estimate how long the advantage lasts, and compare that answer with the price.

Technology does not eliminate scarcity. It moves it. Abundance creates new possibilities; relative scarcity determines where bargaining power migrates; price determines whether the investor participates. That is the Abundance Paradox.

> The future is a power law. Own the few.

References and research notes

[1] Epoch AI, Trends in Artificial Intelligence. Historical trend compilation; page reviewed 7 September 2026. Source

[2] Stanford HAI, The 2025 AI Index Report. The >280× comparison is Nov 2022–Oct 2024 at GPT-3.5-level performance. Source

[3] International Energy Agency, Key Questions on Energy and AI, executive summary, 16 April 2026. 2025 estimate and 2030 central projection. Source

[4] Bank for International Settlements, Progress and peril, Annual Economic Report 2026. Hyperscaler expenditure refers to 2025–26 plans, not realized supplier profits. Source

[5] Phurichai Rungcharoenkitkul, The AI investment race, BIS Working Paper 1367, July 2026. A theoretical investment-race model. Source

[6] Leopold Aschenbrenner, Situational Awareness: The Decade Ahead, June 2024. Author-supplied reference PDF; historical influence, not a current forecast adopted here. Source

[7] Sherwin Rosen, The Economics of Superstars, American Economic Review 71(5), 1981, pp. 845–858. Source

[8] Fred Hirsch, Social Limits to Growth, 1976. See also J. J. Tyssedal, What is a positional good? Recovering Hirsch’s insights, Economics and Philosophy, 2026. Source

[9] Google Quantum AI and Collaborators, Quantum error correction below the surface code threshold, Nature, published online December 2024. Source

[10] Yuri Alexeev et al., Artificial intelligence for quantum computing, Nature Communications, 2 December 2025. Review of AI applied to quantum systems. Source

[11] Francisco J. R. Ruiz et al., Quantum circuit optimization with AlphaTensor, Nature Machine Intelligence 7, 374–385, 2025. Source

[12] IBM Quantum, How IBM will build the world’s first large-scale, fault-tolerant quantum computer, 10 June 2025. The 2029 Starling specification is a vendor roadmap target. Source

[13] NIST, Post-Quantum Cryptography program and migration guidance. The first three final standards were released in August 2024. Source

[14] Riddhi S. Gupta et al., A systematic review of quantum machine learning for digital health, npj Digital Medicine 8, 237, 2 May 2025. Scope: digital health, not all quantum algorithms. Source

[15] C2PA, Explainer, specification 2.4. Provenance records assertions and edit history; it does not establish the truth of the content. Source

[16] Charles Edwards, Bitcoin Energy-Value Equivalence, Capriole Research, December 2019. Source

[17] Charles Edwards, Capriole Research and newsletter archives. Stylistic and research lineage; not independent validation of this thesis. Source

[18] Kraken, xStocks remain largest provider of tokenized equities, surpass $25 billion in total transaction volume, 19 February 2026. Issuer-platform report; activity includes minting and redemption as well as trading. Source

[19] Crypto.com, Tokenized Stocks launch in the EEA, 12 August 2026. Product announcement, eligibility and terms apply. Source

[20] Bybit, Introduction to TradFi Perpetual Contracts. Product documentation reviewed 7 September 2026; derivatives do not constitute underlying share ownership. Source

[21] LSEG, London Stock Exchange to launch UK tokenised equity structures and announces partnership with Payward, 1 September 2026. Plans and intended 2027 xStocks listing are subject to relevant approvals. Source

[22] ICE / NYSE, The New York Stock Exchange Develops Tokenized Securities Platform, 19 January 2026. Proposed platform for which regulatory approvals would be sought. Source

[23] Visa, Visa Accelerates Stablecoin Momentum: Adding Five Blockchains for Settlement, 29 April 2026. US$7B is the pilot's annualized settlement run rate. Source

[24] Circle, Circle Launches AI Infrastructure to Power the Agentic Economy, 11 May 2026. Product availability is not a measure of realized agent-commerce volume. Source

[25] International Labour Organization, Generative AI and Jobs: A Refined Global Index of Occupational Exposure, Working Paper 140, May 2025. Exposure is distinct from displacement. Source

[26] OECD, The OECD AI exposure measure, 26 May 2026. Source

[27] Anthropic, Labor market impacts of AI: A new measure and early evidence, 5 March 2026. Includes null unemployment findings and tentative younger-worker hiring evidence. Source

[28] American Gaming Association, Commercial Gaming Revenue Hits $78.7 Billion in 2025, 26 February 2026. Industry association compilation of legal U.S. commercial markets. Source

[29] OCC, 2025 Annual Report: A Message from Leadership. Total cleared contracts; not a count of gamblers or a measure of AI effects. Source

[30] Hyperliquid, Trading fees, official documentation reviewed 7 September 2026. Mechanism description, not audited revenue or a guarantee of tokenholder returns. Source

[31] Pump.fun, Fees, official documentation reviewed 7 September 2026. Distinguishes creator, protocol and liquidity-provider fees. Source

[32] Deutsche Bank, How tokenised assets transform liquidity management, 17 April 2026. Discussion of gross settlement, netting and intraday funding trade-offs. Source

[33] Latvala et al., Social disadvantage and gambling severity: a population-based study with register-linkage, European Journal of Public Health 31(6), 2021. Finnish observational data; association does not establish direction of causality. Source

[34] METR, Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity, 10 July 2025. Randomized study in a specific developer setting. Source

[35] METR, We are Changing our Developer Productivity Experiment Design, February 24, 2026. Explains selection effects and limitations of newer speed estimates. Source

[36] Volodymyr Sivak et al., Reinforcement learning control of quantum error correction, Nature, 8 July 2026. The 3.5× result concerns logical-error-rate stability under injected drift, with combined control and decoder steering. Source

[37] Jeremiah Dittmar, Information technology and economic change: The impact of the printing press, Quarterly Journal of Economics, 2011; author's summary at CEPR. Source

[38] William D. Nordhaus, Do Real-Output and Real-Wage Measures Capture Reality? The History of Lighting Suggests Not, in The Economics of New Goods, NBER, 1996. Source

[39] Bank for International Settlements, Anchoring trust in money: innovation beyond stablecoins, Annual Economic Report 2026. Assessment of stablecoins, tokenized deposits and monetary architecture. Source

[40] Erik Brynjolfsson, Danielle Li and Lindsey R. Raymond, Generative AI at Work, NBER Working Paper 31161, 2023. The roughly 14% figure refers to the working-paper estimate, not a general productivity forecast. Source

[41] Pump.fun, PUMP token page, reviewed 7 September 2026. Buyback-and-burn program description; changeable platform terms are not equivalent to shareholder rights. Source

[42] DARPA, Quantum Benchmarking Initiative. Utility-scale operation is assessed by whether computational value exceeds cost; program page reviewed 7 September 2026. Source

[43] P Equity Research (@pequityresearch), "Morgan Stanley: Bearings," X post dated 13 July 2026 in the supplied transcript. Secondary summary used as supplied; the original Morgan Stanley report, baseline, methodology and chart were not available for review. Forecast and component illustrations remain attributed to this source. Source

[44] OpenAI, Request for Proposals: U.S. Hardware Manufacturing, January 2026, p. 1, Robotics scope. Primary PDF reviewed 7 September 2026; precision bearings are one of six listed robotics categories. The RFP establishes a request for manufacturing proposals, not orders or investment endorsement. Source

[45] Anson Ho et al., Algorithmic Progress in Language Models, March 2024. Historical evaluations from 2012–2023; compute-requirement halving estimate approximately eight months, with a 95% confidence interval of five to fourteen months. Source

Research lineage

The framework's intellectual roots include the economics of complements, superstar markets and positional goods.[7][8] Aschenbrenner's June 2024 Situational Awareness is an influence on the treatment of AI scaling and industrial constraints; its timing and geopolitical forecasts remain its author's.[6] Capriole's prior research contributes a focus on measurable systems, energy and digital assets.[16][17] The present framework is a synthesis and investment hypothesis, not a claim to have discovered new universal economic laws.

Disclosure

This is general research containing opinions, illustrative examples and forward-looking hypotheses, not an offering document or a recommendation to buy any fund, security or digital asset. Company announcements establish what the company reports or intends, not independent proof of execution. Examples do not identify current holdings. No portfolio count, asset-class cap, allocation, fee or liquidity promise is created by this paper. Commercial fund terms remain governed by the applicable offering documents. Forecasts can fail; technologies, platforms and financial claims can become obsolete or lose all their value.