

GA-COURTENAY SPECIAL SITUATIONS FUND

GA-Courtenay Special Situations Fund is managed within Green Ash Partners Investment Management Ltd
11 Albemarle Street, London, W1S 4HH, UK. Green Ash Partners is regulated by the FCA.



A UNIFIED FRAMEWORK FOR MODELLING THE FUTURE

March 18th, 2026

Introduction

At the pinnacle of advantaged capital allocation sits the ability to make a series of accurate predictions about the future. This is not hyperbole but tautology. Equity valuation is optimised when accuracy regarding the consistency, duration, and growth of future cashflows is correctly estimated.

The contention that an investor can achieve *A Unified Framework For Modelling the Future* may seem like overreach, or at worst, delusion. Nevertheless, as we investigate the premise of this white paper, we will put forward a path to orientate toward this ambitious goal.

We begin by identifying the general conditions under which determinism occurs. The future path of an object through space can be accurately predicted when an existing orientation possesses momentum, and interference in its path is absent. Furthermore, predictive accuracy is heightened when the passage of the object to its predicted goal is also subject to reinforcement. Under these conditions, what is predicted to happen becomes what will happen.

In systems governed entirely by the laws of physics, this is self-evident. We can predict with certainty that Jupiter will continue to revolve around the Sun—its existing orientation of momentum, absent interference, is reinforced by gravitational pull.

Under rare circumstances, a comparable form of absolute determination applies to human trajectories as well. Consider a person jumping from a surfboard onto a beach, wearing only swimming trunks. It is impossible to predict their circumstances ten minutes later. There is no momentum following their initial orientation, and their subsequent behaviour will be subject to interference from new inputs. No reinforcement toward a particular end goal is present either.

Now consider the same person jumping from a plane mid-flight, again wearing only swimming trunks—absent a parachute. The initial orientation is followed by momentum, and the absence of a parachute eliminates plausible interference. And until terminal velocity is reached, the force of gravity provides reinforcement.

The outcome – sadly for its protagonist – becomes deterministic.

Figure 1: Reductively “similar” human trajectories are characterised by different forecast accuracy. Prediction requires all of – an initial orientation subject to momentum, an absence of interference, and reinforcement toward a particular end goal.



A person jumps off of a surfboard wearing just swimming trunks. It is impossible to predict what they will then be doing in ten minutes time.



The same person jumps out of a plane, mid-flight, again, wearing just swimming trunks and absent a parachute. The future outcome is deterministic, and no scenario of interference is plausible.

In our 2024 white paper, *Macro Protection Within A Unified Framework For Capital Allocation*¹, we identified instances in financial markets where initial orientation subject to momentum may continue unperturbed by interference at the same time as being subject to reinforcement—yielding an analogous forecast accuracy. However, in most human circumstances these conditions are absent. A framework for absolute prediction of the future will not apply.

The implication is that predicting the future path of technological development must embrace a more realistic ambition: whilst remaining within the framework of initial orientation and momentum, interference absence, and reinforcement, to master the assignment of probabilities to various scenarios to a far greater degree than is widely recognised as possible. It is the route to this mastery that this white paper explores.

“I look at the future from the standpoint of probabilities. It’s like a branching stream of probabilities.

There are actions that we can take that affect those probabilities or that accelerate one thing or slow down another thing.. And I may introduce something new to the probability stream.”

Elon Musk, *The Future We’re Building – and Boring*, April 2017²

Evolution has programmed humans to prioritise what we *see*. Our visual cortex, linked to *Daniel Kahneman’s System One*³ processing, absorbs by far the greatest volume of information—information that, in simpler times, did not benefit from the reasoning advances humans can now apply, including through artificial intelligence. The result is a knowledge acquisition route innately prioritised as observation of change, conjecture about governing rules, and stipulation of predictive ability. In other words: scientific method.

Yet data is only a subset of the inputs that an advanced methodology for future prediction must consider. Monotheistic dependence on data is a mistake, as Taleb’s turkey before Christmas discovered⁴.

To put it another way, whilst humans will always be interested in the side of the Moon we can see, predicting the future requires equal interest in the side we cannot see. The implication: despite *System One* feeling efficient and obvious, future prediction demands the determined and more time consuming deployment of abstract reasoning through *System Two*—a choice that feels hard precisely because it is.

“First principles is kind of a physics way of looking at the world. And what that really means is you boil things down to the most fundamental truths and then reason up from there. That takes a lot more mental energy.”

Elon Musk, speaking in interview with Kevin Rose, 2012⁵

The unified framework presented in this white paper represents a synthesis of knowledge which, whilst a number of its pillars focus on the principles that have been argued for practitioners including Elon Musk, Charlie Munger and Leopold Aschenbrenner, it also contains conjectures wholly unique, including the *Superorganisation* concept as a reinforcement mechanism, and we are aware of no other literature which aggregates all of our principles into *A Unified Framework For Modelling the Future* that this white paper presents.

In October 2015, Elon Musk appeared at Stanford University’s Future Fest, in conversation with venture capitalist Steve Jurvetson. He did not merely speculate about coming technologies or offer predictions about specific industries. Instead, he articulated a systematic reasoning process for predicting the future and assigning probabilities to scenarios well beyond what is entirely deterministic. What emerged was a coherent set of principles that treat prediction not as prophecy but as engineering.

For Musk, forecasting the future and constructing it are not separate activities—they are the same discipline viewed from different temporal vantage points. Musk argued that the most reliable way to predict what will happen is to identify what must eventually happen according to fundamental physical, economic and human incentive constraints, then work backward to determine the sequence of enabling conditions required to bring it about.

In modelling the future, however, we argue the predictor can further raise their accuracy by also incorporating the mode of thought of Leopold Aschenbrenner—formerly of OpenAI, author of the June 2024 treatise *Situational Awareness: The Decade Ahead*⁶, co-author of the December 2024 Global Priorities Institute working paper *Existential Risk and Growth*⁷, and one of the most formidable young analytical minds to have emerged from the intersection of economics, existential risk research, and artificial intelligence.

“Right now, there’s perhaps a few hundred people in the world who realize what’s about to hit us, who understand just how crazy things are about to get, who have situational awareness.

I probably either personally know or am one degree of separation from everyone who could plausibly run The Project. The few folks behind the scenes who are desperately trying to keep things from falling apart are you and your buddies and their buddies. That’s it. That’s all there is.”

Leopold Aschenbrenner, *Situational Awareness* white paper, June 2024⁶

Where Musk’s framework begins with the laws of physics and works upward through engineering to civilisational impact, Aschenbrenner’s begins with empirical trendlines and works outward through probability distributions to geopolitical consequence. Where Musk asks “what does the universe permit?”, Aschenbrenner asks “what do the trendlines compel?” Both arrive at predictions of extraordinary ambition, but through fundamentally different analytical routes.

Aschenbrenner's intellectual formation is unusual. He entered Columbia University at fifteen, graduated as valedictorian, and by his early twenties had produced academic work on the relationship between long-run economic growth and existential risk that drew the attention of Tyler Cowen, the economist who awarded him an Emergent Ventures grant. He subsequently worked at the Global Priorities Institute at Oxford, contributed to biosecurity and AI safety philanthropy, and joined OpenAI before departing to launch his own hedge fund, which has emerged amongst the best performing entities across global markets in recent years.

Unique to this white paper, we put forward our own forecast method additions, drawing on principles from natural world systems and the recognition of *Superorganisations*, as groups of humans, aggregated in the corporate form and in this context as a reinforcement mechanism driving future outcomes to group goals. We contend *Superorganisation* recognition as essential to be used alongside the approaches of Musk and Aschenbrenner for *A Unified Framework For Modelling the Future* to achieve predictive accuracy at the leading edge.

Musk's advocacy of testing the viability of predictions through the lens of the laws of physics may sound seductive, however, human knowledge of the laws of physics is incomplete. Certain natural world observations, such as abiogenesis, the defiance of entropy, and consciousness, *must* be explained by the laws of physics yet our knowledge of the exact laws required is absent. We also add therefore, that analysis of natural world phenomena can reveal what the laws of physics must also include, even in the absence of our identification today of those laws which remain unidentified themselves.

Our analysis further expands an aspect explicit in Musk's framework and implicit in Aschenbrenner's: that the toolkit is optimised not merely as one for anticipating technological change but for the predictor to participate in its creation, an aspect described by Musk as prediction being a "*contact sport*" and by Aschenbrenner as the predictor requiring "*situational awareness*."

Musk adds: "*I look at the future from the standpoint of probabilities. It's like a branching stream of probabilities.. and I may introduce something new to the probability stream.*" His explicit recognition is that it is humans driving technological change. It is the business leaders, and their workforces whose efforts change the probability stream of future outcomes.

"People are mistaken when they think that technology just automatically improves. It does not automatically improve. It only improves if a lot of people work very hard to make it better. And actually it will I think—by itself—degrade.

If you look at the progress in space, in 1969 we were able to send somebody to the moon. Then we had the Space Shuttle. The Space Shuttle could only take people to low-earth orbit. Then the Space Shuttle retired and the United States could take no one to orbit.

You look at great civilizations like ancient Egypt, and they were able to make the pyramids, and they forgot how to do that. And the Romans, they built these incredible aqueducts. They forgot how to do it."

Elon Musk, speaking at TED, May 2017⁸

The implication is that accurate detection of the path of future technological change should not be merely abstract, but grounded in the explicit recognition that highly productive groups of humans are collaborating in a reinforcement manner to drive each transition forward.

In our 2025 white paper, *Evaluating the Vector Sum Theory of Corporate Output*⁹, we identified that both Elon Musk, and Charlie Munger late in his career, have publicly recognised that the output of a group of humans can be understood as the vector sum of individual contributions, where each worker represents a vector combining productivity (magnitude) and alignment with corporate goals (direction). This framework, expanded in my 2025 book *The Super Organisation Secret*¹⁰, argues that decisive human productivity – including in driving technological shifts – occurs only when variables beyond traditional metrics are prioritised, specifically, the variables which empower the building and maintenance of high-performing, aligned teams.

If prediction truly requires the predictor to engage as a *contact sport*, gaining *situational awareness* — if forecasting the future and engaging in its building are inseparable—then the predictor appears bound to the same constraints as the builder: the investment of time, the burden of execution, the building of personal network, the limitation to those few ventures one can individually engage in. This is the paradox that the outside minority investor faces, they are *outside*, not engaged in the direct building of the future – and as such they would appear to face a handicap with regard to the predictive contentions they can reasonably expect to make.

However, this white paper argues that the full set of principles underpinning *A Unified Framework For Modelling the Future*, carefully examined, can deliver otherwise. On a combined basis, we are aware of no other literature which aggregates all of our principles into *A Unified Framework For Modelling the Future* that this white paper presents. As such, we present a cure for the paradox that the outside minority investor otherwise faces. We conclude the paper by orientating our analysis to reveal that our framework, fully understood and deployed, permits a position available to the investor who masters these principles that may prove superior even to that of the participant's own position—a path to amplified participation in the future's construction, a path allowing de-risking through hedging, a path which benefits from the benefits the instancy advantage of being able to switch between investments, and a path freed from the time allocation handicap of direct involvement. That resolution awaits in our conclusion.

Our principles do not guarantee success—Musk himself estimated his probability of success at SpaceX and Tesla at roughly ten percent when he began. But they do offer a disciplined approach to raising the accuracy of predictions, timing entry appropriately, and maintaining commitment through the inevitable failures that accompany any attempt to reshape how the world works.

The principles are sequential but not strictly linear. For coherency of argument, the structure of the white paper is organised along the lines of initial orientation and momentum, interference absence, and reinforcement. Each aspect builds upon one another while operating in parallel. Together, they constitute what might be called *a biophysics of foresight*—an approach to the future grounded not in intuition or optimism but in systematic analysis of what the laws of physics and of nature permit, and what human organisational dynamics propel.


Identifying high gradient, initial orientation with momentum

High gradient initial orientation and momentum is concurrent with *phase transition*

Our analysis begins with one of its most powerful elements for modelling the future – which is the correct recognition of the timing of high gradient, initial orientation with momentum.

Technologies do not develop along similar gradients. Some improve gradually along well-understood curves; others undergo accelerated *phase transitions* where the fundamental state of affairs changes rapidly and irreversibly. Recognising which trajectory a technology follows—and where it currently sits on that trajectory—is essential for determining when to engage.

Figure 2: Technologies do not all develop along similar gradients. Some improve gradually along well-understood curves; others undergo accelerated phase transitions where the fundamental state of affairs changes rapidly and irreversibly¹¹.



The metaphor Musk employs is water freezing into ice. The transition from liquid to solid is not gradual; it occurs very rapidly yet only when conditions cross a precise threshold. Before the threshold, entirely liquid. After, entirely solid. The transition point represents a qualitative shift, not merely a quantitative one. As Figure 2 illustrates, some technological developments exhibit similar dynamics.

"The internet was happening in '94, '95...

And then I thought, if I watch the internet get built while I'm doing this, that would be really frustrating."

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015¹²

In 1995, Musk arrived at Stanford to pursue a PhD in energy storage for electric vehicles—a field directly connected to the sustainable energy domain he had identified as civilisationally important. He left after two days to start an internet company. The decision was not an abandonment of sustainable energy; it was a recognition that the internet was undergoing a phase transition while energy storage was not.

The internet in 1995 was passing through a critical threshold. Foundational architectures were being established. Business models were being defined. Competitive positions were being staked out. Decisions made during that period would shape the internet's structure for decades. Participation during the transition window would have consequences that participation afterward could not replicate.

Energy storage, by contrast, remained on a path of slow, steady progress. Battery energy density was improving incrementally along predictable curves. No phase transition was imminent—no moment when fundamental relationships would shift irreversibly. The window for meaningful participation was wide. One could return later without having missed a critical inflection point.

The analysis implies a taxonomy of technological trajectories. Phase-transition technologies reward early positioning disproportionately; entrepreneurs and investors present during the transition capture durable advantages. Gradual-improvement technologies offer more forgiving timing; entry can occur across a wider range of dates without dramatically affecting outcomes.

Musk eventually returned to sustainable energy, as he had anticipated. But the sequencing was deliberate. The capital and credibility accumulated through Zip2 and PayPal¹³ enabled subsequent engagement with capital-intensive hardware ventures that would otherwise have been impossible.

Identifying phase transitions requires high resolution knowledge and alertness

The opacity challenge: exponential or accelerating processes appear linear when observed locally

Recognising that phase transitions matter is one thing; identifying them is another. The challenge with exponential or accelerating processes is that they appear linear when observed locally. The curvature only becomes apparent over extended timeframes, by which point the transition may have already occurred.

"The tricky thing about predicting things when there's an exponential is that an exponential looks linear close up. But it's actually not linear."

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

Step one: obtain high resolution knowledge by achieving proximity to the frontier where history is being made

For the first step, we adopt Leopold Aschenbrenner's prerequisite: the predictor must achieve what Aschenbrenner calls *situational awareness*. This is not mystical insight; it is the epistemic advantage that accrues to those that achieve proximity to where the trendlines are generated, the trillion-dollar contracts are signed, and the engineering breakthroughs are first observed.

“You can see the future first in San Francisco. Over the past year, the talk of the town has shifted from \$10 billion compute clusters to \$100 billion clusters to trillion-dollar clusters. Every six months another zero is added to the boardroom plans.”

Leopold Aschenbrenner, Situational Awareness white paper, June 2024

Aschenbrenner contrasts this form of tiny insider group—“perhaps a few hundred people”—with the vast majority of observers. The latter and larger group lack the information to form accurate predictions—Nvidia analysts who believed 2024 might represent peak demand; mainstream commentators anchored to the dismissive frame that AI is “just predicting the next word”; policymakers operating on timelines calibrated to decades rather than years.

“There are perhaps a few hundred people, most of them in San Francisco and the AI labs, that have situational awareness. Through whatever peculiar forces of fate, I have found myself amongst them.

A few years ago, these people were derided as crazy—but they correctly predicted the AI advances of the past few years.”

Leopold Aschenbrenner, Situational Awareness white paper, June 2024

The principle carries a direct implication for the investor. Information asymmetry in AI is not narrowing; it is widening. Capital allocators who achieve even partial situational awareness—through proximity to the labs, through careful reading of technical publications, through relationships with the people building these systems or by studying the public disclosures from these same people including their capital allocation decisions—operate with a structural informational advantage over those who do not.

This is not an argument for insider trading or privileged access. What distinguishes the few hundred people who have situational awareness is not access to secrets but the willingness to take niche specific, focused publicly available information seriously.

Step two: commence prediction not with theory but by observation – of initial orientation with momentum

The second step orientates the predictor toward efficiency – prioritise observation before theory. Ordering the analysis in this manner, the input as the observation and the output as the analysis, lowers the cognitive load required for initial detection at the same time as calibrating initial detection around reality rather than theory.

Aschenbrenner advocates that the observation process should target the extrapolation of the existing trendlines of technological process, a reality-targeting rather than theory orientation. By example, Aschenbrenner does not argue for AGI by 2027 through philosophical reasoning about the nature of intelligence, nor through physics-based analysis of computational limits. He argues for it by what he calls “counting the OOMs”—orders of magnitude—tracing the empirical growth in compute, algorithmic efficiency, and capability gains, and extending those slopes forward.

“I will simply ‘count the OOMs’: look at the trends in 1) compute, 2) algorithmic efficiencies (algorithmic progress that we can think of as growing ‘effective compute’), and 3) ‘unhobbling’ gains...”

“We trace the growth in each over four years before GPT-4, and what we should expect in the four years after, through the end of 2027.”

Leopold Aschenbrenner, Situational Awareness white paper, June 2024

The method – prioritising observation before theory – is sufficiently far from orthodoxy to hand to its practitioners an advantage differentiation. For the AI case, compute scales at roughly 0.5 orders of magnitude per year. Algorithmic efficiency gains contribute another 0.5 OOMs per year. “Unhobbling” gains—fixing obvious limitations and giving models tools—deliver step-changes on top. GPT-2 to GPT-4 represented a leap from approximately preschooler-level to smart high-schooler-level capability in four years. Another equivalent jump, which the trendlines indicate, would land somewhere around expert-level AI by 2027.

“AGI by 2027 is strikingly plausible. GPT-2 to GPT-4 took us from ~preschooler to ~smart high-schooler abilities in 4 years...”

“We are on course for AGI by 2027. That doesn’t require believing in sci-fi; it just requires believing in straight lines on a graph.”

Leopold Aschenbrenner, Situational Awareness white paper, June 2024

Critically, the AI-relevant trendlines have held despite persistent scepticism. Aschenbrenner notes that, concurrently, theory-based predictions of a deep learning “wall” have been made every year for the past decade and have been consistently wrong. Aschenbrenner explicitly contrasts ‘productive empirical work’ on trendlines with ‘unproductive philosophical debate’—what he calls “shape-rotating” versus “word-selling”:

“The trendlines have been astonishingly consistent, despite naysayers at every turn.”

Leopold Aschenbrenner, Situational Awareness white paper, June 2024

“I was kind of stuck in these philosophical arguments with people... and the thing that sort of just like made it visceral to me was: I’m gonna look at the research work that is actually happening.”

Leopold Aschenbrenner, interview with The Center for the Study of Partisanship and Ideology, 2024¹⁴

Aschenbrenner extends the same empirical scepticism to expert surveys and superforecaster aggregates, which he calls “pretty generally low signal”—noting that two different sets of reputable superforecasters he commissioned on the same question gave “wildly divergent numbers.” His contention – the trendlines themselves are more reliable than the opinions of those who observe them.

Step three: augment trendline extrapolation by tracking expert predictions

Aschenbrenner begins not with what the predictor believes is theoretically possible but with what the empirical data shows is actually happening. The framework is inherently self-correcting: if the trendlines bend, the prediction updates accordingly, because it is anchored to observable reality rather than to assumptions about what should be possible.

It is at this point Elon Musk’s approach argues that expert opinion and conclusion should be brought into the analysis in order to raise forecast accuracy further. However, Musk still does not advocate that direct technological knowledge is necessary.

Instead, rather than measuring technological progress against some absolute scale—which requires defining and quantifying capability in domain-specific terms— Musk argues that expert predictions about arrival dates should be monitored to determine if they are shifting. If predictions to end point consistently contend sooner arrival than previously forecast, the underlying process is accelerating. If predictions remain static or recede, progress is linear or decelerating.

Musk illustrates this with autonomous driving. *“Three years ago I thought it was 10 years away. Two years ago I thought it was 5 years away. Now I think it’s 3 years away or less.”* In three years of elapsed time, Musk argued seven years of predictive time had vanished. The gap between elapsed time and predicted time is the signature of acceleration.

“If at any point in time your predictions are going further out or coming closer in, that actually would help define whether it’s accelerating or not... If a given technology is always 20 years in the future, it’s more logarithmic. But if predictions keep arriving sooner than expected, it’s exponential.”

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

The method has practical advantages. It requires neither deep technical expertise nor measurement of capability on any absolute scale—only honest tracking of predictions over time and attention to whether those predictions are converging toward the present faster than time itself is passing.

Exponential tracking also identifies domains where hype exceeds reality. Technologies that have been “20 years away” for the past 30 years—nuclear fusion is a canonical example—reveal through their static prediction horizons that no genuine acceleration is occurring. The prediction is not converging; the technology is not commencing a phase transition.

Figure 3: Elon Musk’s framework for exponential tracking advocates monitoring whether expert predictions are for arrival dates shifting to nearer in the future. If predictions are consistently revised to sooner than previously forecast, the underlying process is accelerating.

Shifting expert predictions about AGI arrival dates reveal, using Musk’s framework, exponential growth. Microsoft President Brad Smith, in 2023, estimated a time period up to “many decades”; by 2025, Musk estimated AGI arrival in 2026.

“There’s absolutely no probability that you’re going to see this so-called AGI, where computers are more powerful than people, in the next 12 months. It’s going to take years, if not many decades.”

Microsoft President Brad Smith,
speaking to reporters in November 2023¹⁵

“I think we’ll hit AGI next year – in 2026.

And then by 2029 or 2028, AI will exceed the intelligence of all humans combined.”

Elon Musk, speaking with
Peter Diamandis in December 2025¹⁶

Determining an absence of interference

Establish that the prediction end point is permitted by the laws of physics

In assessing the next stage of prediction, that the forecast path can continue unperturbed by interference, we first bring in Musk’s advocacy that the end point prediction of the phase transition should be consistent with both what the laws of physics permit and continue until the point to which the laws of physics deny.

The laws of physics guiderails are applied to address the difference between correctly recognising a phase transition and correctly predicting its end point. For example, one might have correctly recognised the internet phase transition yet incorrectly forecast its terminus as a personal, internet-connected computer in every home—rather than a personal, internet-connected smartphone carried by every human regardless of their location. End point prediction is therefore as critical as recognition of the phase transition itself.

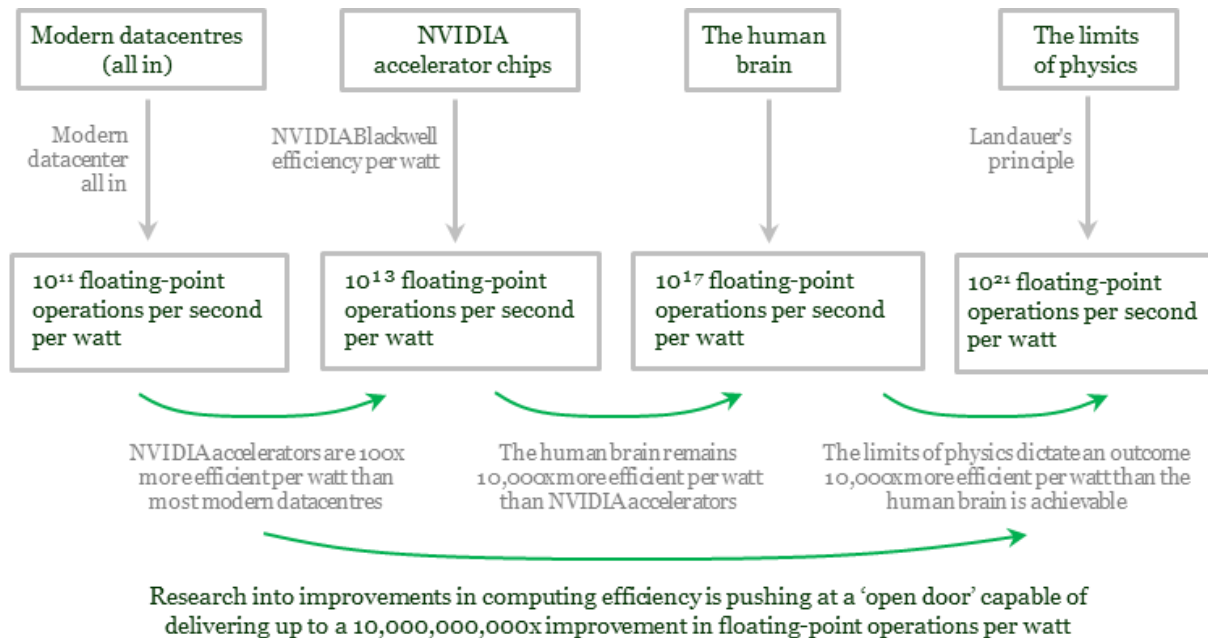
Here, instead of analysing markets, competitors, or business models, Musk advocates that the predictor must establish that the predicted destination is permitted by the laws of physics.

His flagship example – multi-planetary life. Humans – with the right support infrastructure – can survive on another planet. For other Musk ventures: vehicles can be propelled by electricity stored in batteries, information can be transmitted globally at the speed of light, the laws of physics also dictate that orders of magnitude improvements in computing efficiency from current levels remain achievable. These are not aspirations; they represent natural laws of what the universe will tolerate.

This starting point distinguishes Musk’s approach from conventional analysis, which typically begins with market research or competitive positioning. Conventional analysis *assumes a destination already contended in public discourse* and examines routes to reach it. Musk’s framework begins one step earlier: *it derives through the laws of physics whether a destination exists at all in the space of physical possibility.*

Figure 4: Musk’s predictive framework for determining the end point of phase transitions focuses on achieving certainty through the laws of physics; in the figure – the laws of physics dictate further orders of magnitude improvements in computing efficiency from current levels remain achievable¹⁷

The research programs delivering computing efficiency improvements are pushing against an "open door" tailwind of long-term enhancement; datacentre efficiency per watt remains 10⁶× short of the human brain, and a further 10⁴× short of the limits of physics.



The consequence is raised forecast accuracy through the expansion of the possible end point space. Once the predictor knows something is physically possible, they are liberated from the implicit constraints embedded in how things are currently done. Musk's past identification of five domains —internet, sustainable energy, multiplanetary life, genetics, and artificial intelligence—illustrates this principle. Each domain represented a category where physics permits transformations of civilisational magnitude. The internet enables instant global information transfer because photons travel at the speed of light through fibre optic cables. Sustainable energy is possible because the sun delivers orders of magnitude more energy to Earth's surface than humanity consumes. Mars colonisation is achievable because the physics of orbital mechanics and life support, while challenging, contain no fundamental barriers.

"What are the things that are most likely to affect the future of humanity? Just thinking at a macro level... it seemed like it would be the internet and sustainable energy, making life multiplanetary, and then genetics and AI."

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

Starting with physical possibility also provides a natural filter for distinguishing real opportunities from mirages. Perpetual motion machines are not investment opportunities because they violate thermodynamics. Faster-than-light travel is not a near-term business prospect because it contradicts relativity. But reusable rockets, high-density batteries, and global satellite internet—through the laws of physics these destinations are *determinable as possible* before any analysis of probability or practicality begins.

Envision yet to be discovered laws of physics by reference to natural world adaptations

Musk's advocacy that an absence of interference can be validated by compatibility with the laws of physics nevertheless contains a drawback: the full set of the laws of physics are not yet known. The implication is that a prediction that relies on compatibility with known the laws of physics may be handicapped because it will by reject outcomes that are only compatible with the unknown laws of physics.

A partial cure to this shortcoming is the deliberate emulation of biological and natural-world processes as an alternative benchmark for feasible design. If living systems or natural phenomena accomplish complex tasks, then the task must be occurring without violating the laws of physics, and as such engineered systems designed to replicate them can be contended as meeting the laws of physics even in the absence of full knowledge of those laws.

This approach strips away artificial constraints imposed by conventional industry practices and reframes "impossible" technical challenges as solvable by aligning human technology with proven natural mechanisms.

Whilst we have not observed Musk directly advocating this cure, his ventures appear to reveal his recognition of it. For Tesla's autonomous-driving system, Musk rejected LIDAR and other sensor suites, insisting instead on camera-only vision because humans navigate roads using visual data alone. Similarly, when redesigning The Boring Company's tunnelling machines, Musk observed that a gopher in his yard digs straight down without first excavating a vertical shaft. He immediately ordered the machine re-engineered to mimic this direct-entry method, eliminating unnecessary steps and dramatically simplifying the process.

The more than four billion years of evolution that natural world solutions have had to optimise their function on earth also allows natural world adaptation to inform design enhancement. This mindset rejects the assumption that human technology must be more complex than natural analogs; instead, it treats natural processes as existence proofs that the desired outcome is physically attainable concurrent with the simplest naturally occurring form, forcing teams to confront artificial constraints and accelerate iteration toward radical efficiency.

The result is a consistent pattern: designs that are simpler, cheaper, and faster precisely because they align with how nature has already solved the same underlying problems.

Establish that the future prediction end point is economically viable

Once a destination is confirmed as physically possible, scenarios of interference from economic non-viability, essentially a form of budget blocking by capital allocators, should also be ruled out.

This step in the prediction process leans on the first principles analysis process advocated for by Elon Musk. The first principles analysis involves stripping away every assumption about cost based on historical precedent, and instead building upward from fundamental physical requirements. The question is not "What do rockets cost?" but "What could rockets cost, given what they are made of?" What are the market prices of their materials? What energy is needed to reach orbit? What is the theoretical efficiency limit of chemical propulsion? When Musk performed this analysis for orbital launch, he discovered that the raw material cost of a rocket—aluminium, titanium, copper, carbon fibre—was approximately two percent of the prevailing launch price.

As such, first principles reasoning requires decomposing a problem to its most basic, empirically verified elements and reconstructing it from the ground up. The gap between first-principles cost and actual cost represents

accumulated convention, institutional overhead, and design choices never reconsidered. That gap is not physics—it is history. And history, unlike physics, can be rewritten.

"The physics approach to thinking is very good—the first principles approach. Applying the first principles approach to thinking is a good way to figure out counterintuitive situations.

Do not reason by analogy, by looking at what other people have done—instead ask what is fundamentally true."

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

The methodological opposite of first principles reasoning is reasoning by analogy: observing what others have done and extrapolating. If those assumptions are wrong, or circumstances have changed, analogical reasoning produces systematically misleading predictions. In this regard Musk's reasoning from first principles provides counterbalance to Aschenbrenner's trendline extrapolation.

First principles reasoning also reveals which constraints are hard and which are soft. Hard constraints are imposed by physics: the energy required to reach orbital velocity, the strength-to-weight ratio achievable with known materials, the information-theoretic limits of data compression. Soft constraints are imposed by convention: the way rockets have historically been manufactured, the business models of incumbent launch providers, the regulatory frameworks governing space activities. Prediction becomes tractable when you can identify the costs which can be removed, and track that this is indeed the removal that is being achieved.

A predictive path whose attributes disincentivise interference from regulators

An additional condition in determining that a forecast trajectory can occur unperturbed by interference is that the trajectory is free from interference by regulators. Human political systems will closely scrutinize and ultimately block technologies that bring with them too high a level of hazard or catastrophe risk.

For the predictor, this principle demands that regulatory resistance risk be incorporated not as a background factor but as a primary variable. Companies and research organisations that are allocating resources to alignment in proportion to the risk—rather than in proportion to immediate commercial returns—are more likely to navigate the transition successfully. The potential misallocation is therefore simultaneously a risk indicator and an opportunity indicator: those who correct the misallocation earliest will capture disproportionate value.

Aschenbrenner distinguishes between "level risk", essentially a background level risk during any one temporal period, and "transition risk", the risk which is newly introduced by the technological transition occurring. Level risk is proportional to the danger potential of the technologies present at any one level, transition risk is proportional to the danger introduced by the process of developing new technologies within the context of an correspondingly immature regulatory framework.

“Rather than a certain level of technological development corresponding to a certain level of risk, the relevant risk is something like—say think about AI—we live in a world pre-AI and then suddenly AI determines the future of human civilisation. There’s a one-time transition we want to make this transition go well.”

Leopold Aschenbrenner, conversation with Ben Yeoh, 2022¹⁸

There are several potential nullifiers of regulatory interference during each period of transition risk. The first is transition speed – acceleration reduces the time spent at each risky technology level, as one simply passes through dangerous states more quickly. Second is the wealth generation impact of the technology transition: wealthier states of society are systematically safer, because the wealth effect motivates more stringent safety spending.

The outcome: low regulatory inference is a function of the speed and wealth generation attributes of technological progress.

“I think on the margin I’m probably a bit more convinced by the case of my paper which is the story of—you want to develop safety technologies. I think if you think about the pandemic, the way we ultimately conquered it is mRNA vaccines and better technology rather than political intervention or social distancing. But I think it’s still a very very open question and I wouldn’t be surprised if it went either way.”

Leopold Aschenbrenner, conversation with Ben Yeoh, 2022

For the investor, the transition risk framework provides a critical timing insight. The period of maximum risk — and therefore maximum opportunity—is not when a technology has matured into steady-state deployment, but during the transition itself. Capital allocated during the transition captures the premium associated with navigating uncertainty. Capital allocated after the transition has resolved captures only the steady-state return.

Recognise specific failure outcomes within controlled parameters not as interference, but as part of the statistically inevitable shape of the technological success path

Finally, determining an absence of interference does not reject that the path of technological development will retain volatility, stops and starts, as part of its statistically inevitable shape.

The path by which new technological development ultimately proceeds will still contain fundamental limits to what simulation and analysis can anticipate in advance. This is not a limitation of current tools that will be overcome with better computers; it is a structural feature of certain domains. Any attempt to build the future will encounter some element of trial and error.

As such, instances of failure within controlled parameters should be treated not as recognition of interference nor forecast error, but as a statistical inevitability to be managed. The question is not how to prevent failure but how to structure activities such that failures produce learning without producing catastrophe.

The foundation of this perspective is what Musk calls the "entropic" nature of complex systems. There are vastly more ways for things to go wrong than for them to go right. This asymmetry is not a design flaw; it is a mathematical property of high-dimensional systems. The solution space is sparse; the failure space is dense.

"There are many more ways to fail than to succeed—there's an entropic basis for this.

For a rocket, there's like a thousand ways that thing can fail and one way it can work."

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

Given this asymmetry, execution in complex domains must accommodate a necessary failure rate. Not because failure is desirable—Musk explicitly states that *"given the options, I prefer to learn from success"*—but because the alternative to accepting some failures is never attempting anything difficult enough to matter.

"You can never test the rocket completely in the environment it's actually going to experience. You can't fully recreate something moving super fast in a vacuum on the surface of Earth... any error between the simulation and reality can result in failure."

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

The practical implication: prediction in complex domains must be structured as an iterative process with aspects of a discovery process. Initial predictions may be roughly right yet also wrong with regard to their more precise assumptions. The value of making them is not in their accuracy but in their correct general orientation and in the ability of the execution path to adjust as new learnings better inform the optimisation of each successive design step. That information updates subsequent predictions. This is what Musk means by calling prediction a *"contact sport"*—arguing that the predictor must be directly engaged with the businesses building the future, experiencing successes and failures firsthand, and adjusting predictions based on what that engagement reveals.

Ascertaining the presence of a reinforcement

As this paper moves now to the third aspect driving its predictive framework – reinforcement. We define reinforcement as positive feedback, exhibiting a square-law relationship, in which the force driving change amplifies as change occurs. An intuitive analogy is a car in which the brake and accelerator are reversed: the force driving acceleration is the driver attempting to slow down by pressing on the brake, yet because the pedals have been switched, the harder the driver tries to slow down, the faster the car accelerates. That is positive feedback. It is not the same thing as a car merely having momentum.

Absent positive feedback, momentum would be subject to negative feedback – the faster the change occurs the more likely an opposite force – interference – is precipitated. Reinforcement as its name suggests therefore can be considered as the final component to allow technological change momentum to move to its end destination in a manner that overcomes all realistic remaining obstacles, including those not yet anticipated or imagined, that would otherwise lower predictive forecast accuracy.

We put forward three categories of reinforcement as necessary for the predictive analysis to address. The first is reinforcement through recursive self-improvement from the technological path itself, the second is reinforcement through government stimulus and incentive, and the third is reinforcement from the nature of the human *Superorganisations* building the technological progression.

Reinforcement through recursive acceleration when the technological path precipitates an upward adjustment function on its own development gradient

We have defined reinforcement as characterised by when the force driving change amplifies as change occurs. These conditions can be met by the dualism of a technological path itself. More powerful computers are built to satisfy the human general demand for more powerful computational tools, however, more powerful computers can also be used to accelerate the development process for building even more powerful computers. As such, the force (computing) driving change (more powerful computing) amplifies as change occurs.

This aspect of the force driving change amplifying as change occurs is seen in industries outside of computing. The US railroads were built satisfy the human general demand for heavy goods locomotion over large distances, however, the build of the railroads also, by enabling heavy goods locomotion over large distances, provided an acceleration in the pace with which even newer railroads could be built. Similarly, today, SpaceX is not only significantly lowering the cost of placing mass into orbit, but in doing so is stimulating the growth of new in-orbit industries, increasing demand for additional mass to be placed into orbit, and thereby, by economies of scale, further lowering the cost of SpaceX placing each additional kilo of mass into orbit.

However, the reinforcement effect is perhaps at its strongest in the field of computing and particularly with regard to AI. Critical differentiators include that much of computing operates at the speed of light, combined with the cost of information transfer being effectively zero resulting in the cost of information replication and distribution also trending to zero. The cognitive aspect of AI can also be directly used to accelerate the enhancement of each generation of AI hardware and software design. Furthermore, the combination of SpaceX and AI introduces a further reinforcement scenario: as space-based datacentres take hold, the further they are placed into space, the

closer they can be to the sun, increasing the cheapness of their solar-based computing power, and, through elasticity of demand incentivising additional datacentres in their proximity.

The implication is that the AI trajectory possess several dimensions driving recursive acceleration. Where most forecasters treat AI progress as an external variable—something happening to the economy—it is instead a variable that, past a critical threshold, begins to recursively accelerate its own development. This is the intelligence explosion, and not as science fiction but as logical consequence. If AI systems can automate AI research, then the rate of algorithmic improvement is no longer constrained by the number of human researchers. Millions of AI agents, running at superhuman speed, can compress what would have been a decade of human-paced algorithmic progress into less than a year.

“Hundreds of millions of AGIs could automate AI research, compressing a decade of algorithmic progress (5+ OOMs) into 1 year. We would rapidly go from human-level to vastly superhuman AI systems. The power—and the peril—of superintelligence would be dramatic.”

Leopold Aschenbrenner, Situational Awareness white paper, June 2024

“Rather than your reference point being the last 150 years of very constant growth, if you look at longer time horizons the Industrial Revolution massively accelerated growth, the Agricultural Revolution massively accelerated growth. Overall economic growth has actually been accelerating... and it's plausible that AI would do that [again].”

Leopold Aschenbrenner, interview with The Center for the Study of Partisanship and Ideology, 2024

The practical implication of recursive acceleration is that conventional forecasting methodologies—which assume the rate of change itself is roughly constant—will systematically underestimate the pace of transformation when the technological path itself also acts as an upward adjustment function on its own development gradient. In these scenarios, the trendlines do not merely continue; they steepen. With regard to AI, capital allocators who plan only for the pre-AI trajectory risk being blindsided by the post-AI acceleration.

That the AI technological path possesses a prospectively accelerating growth rate also places elevated value on any monopolistic bottleneck entities that could slow the intelligence explosion – electricity and power supply, chip production and scaling, energy efficiency in data transfer and compute, and testing and verification time.

“The most extraordinary techno-capital acceleration has been set in motion. As AI revenue grows rapidly, many trillions of dollars will go into GPU, datacenter, and power buildout before the end of the decade. The industrial mobilisation, including growing US electricity production by tens of percent, will be intense.”

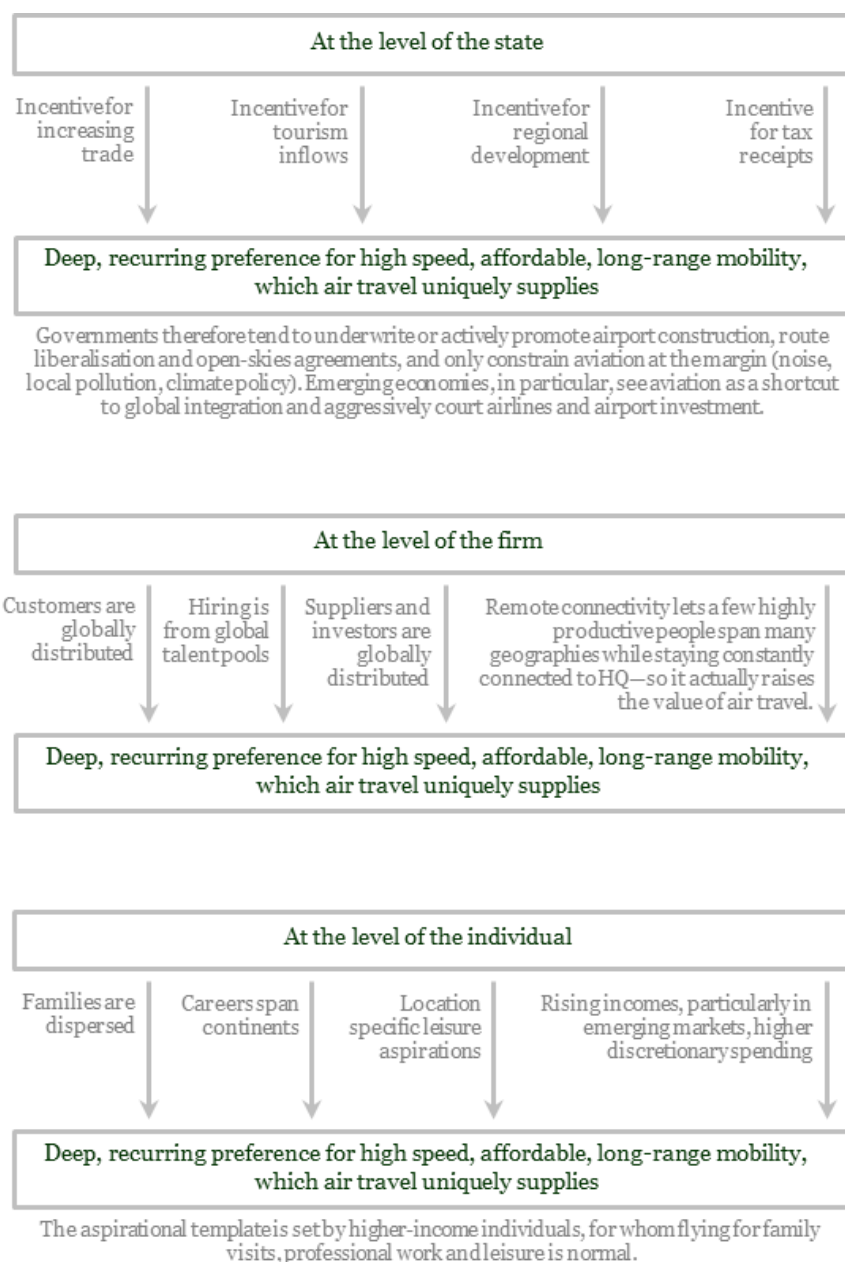
Leopold Aschenbrenner, Situational Awareness white paper, June 2024

Reinforcement through regulatory acceleration

Circumstances where the force driving change amplifies as change occurs can also be concurrent with regulatory acceleration: when government entities, in recognising the critical nature of the technological shift, implement subsidies to further accelerate its development.

Musk applies a framework he describes as the ‘usefulness filter’ to determine whether prospective beneficial outcomes are sufficient to precipitate regulatory incentives for acceleration.

Figure 5: Decades-long efficiency improvements in, and increased frequency of utilisation of, passenger air flights have been underpinned by the usefulness filter – increasing passenger flights are useful to society, and as such, governments tend to accelerate the growth of the aerospace industry by underwriting or actively promoting airport construction, route liberalisation and open skies agreements¹⁹.



The usefulness filter asks: is this problem, if solved, almost certainly going to be a good thing for humanity? Musk applied this filter to his original list of five domains and found that three passed clearly—internet, sustainable energy, and multiplanetary life—while two did not. Genetics and artificial intelligence, he concluded at the time, were more "dodgy" in terms of net benefit. Their successful development could produce enormous good or enormous harm; the direction was uncertain.

"I thought the first three, if you worked on those, were almost certainly going to be good. And then the last two—a little more dodgy in terms of the net benefit."

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

The usefulness filter has important implications for prediction. Technologies that exhibit the highest usefulness filter scores will not only eventually be developed and widely adopted, but in many cases accelerated by government stimulus. Market forces, governmental support, and social approval all reinforce in the same direction. The remaining uncertainty concerns timing and implementation, not whether development will occur.

As we detailed in the prior section, a predictive path must possess attributes which disincentivise interference from regulators. Technologies with the highest usefulness filter scores will ultimately be accelerated by regulators, even if regulatory bodies move cautiously. Legal frameworks evolve through precedent and legislation—processes that resist initially acceleration. Social acceptance shifts gradually as familiarity increases and risks become better understood.

Therefore, the predictor of the high usefulness filter technology must align their forecast path with both the regulatory permissioning phase which may slow progress initially yet ultimately be followed by the regulatory stimulus phase which thereon accelerates progress materially. Ignoring this dynamic leads to a characteristic prediction error: overestimating early adoption speed while underestimating eventual penetration.

For the predictor, this implies maintaining separate analyses for capability, permission, and stimulus, and understanding their interaction. Capability can be accelerated through engineering effort and capital investment. Permission requires different interventions: engagement with regulators, public demonstration of safety, participation in standards development, and patient cultivation of social acceptance. Stimulus requires engagement with government entities to stimulate heightened rates of capital deployment.

The reinforcement potential through government stimulus is relevant for the AI era. The forecasts put forward by AI maximalists are not solely contending that AI can increasingly address commercial use cases and that the regulatory interferences will be minimal. There is a further, inevitable, back end to the forecast, which Leopold Aschenbrenner predicts as "The Project", a specific, high-stakes government-backed acceleration initiative, with explicit parallels to the Manhattan Project (the US WWII effort to develop the atomic bomb).

The Project is not a casual term—it is a descriptive prediction of how AI development will shift from private startups to a government strategic imperative to accelerate and control. AI is not just another tech boom, it is a "weapon" with existential implications and military advantages, akin to nuclear technology. The US–China great-power contest removes any optionality. Superintelligence is not a civilian technology; it is the ultimate strategic asset that will decide economic dominance and military supremacy for the remainder of the century.

Private labs (e.g., OpenAI, Anthropic) are ill-equipped for this, so government intervention and great power state-financed acceleration – trillions in funding and international coalitions – is inevitable, reshaping the world order.

Trillions of dollars, hundreds of millions of GPUs, and tens of percent growth in U.S. electricity output are already moving from boardroom sketches to physical buildout. Export controls, security lockdowns, and industrial policy become non-negotiable because handing the model weights to the CCP is unacceptable. By 2027–28, contends Aschenbrenner, the national-security state will enter the arena with a government-scale AGI project; no startup can steward superintelligence alone.

Reinforcement from the *Superorganisation* property building the technological progression


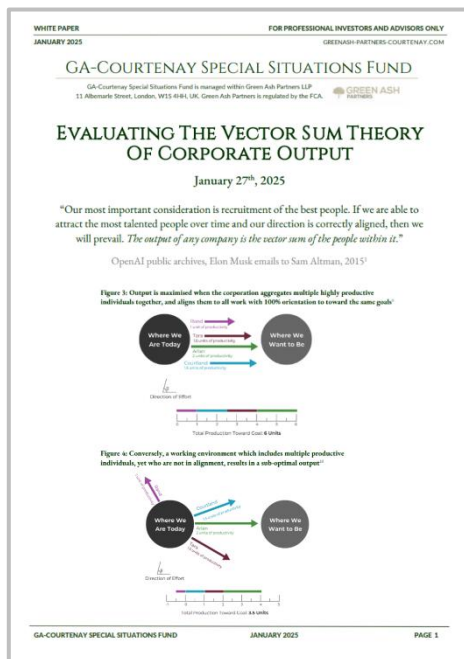
“Nothing succeeds like success.”

Alexandre Dumas, *Le Corricolo*, 1843²⁰

Finally, to identify circumstances where the force driving change amplifies as change occurs our framework recognises that the central character of technological transitions is not merely abstract, but grounded in the explicit recognition that it is highly productive groups of humans which are collaborating to drive each transition forward.

It is humans, aggregated as *Superorganisations*, whose work not only drives forward each technological shift, but through the Superorganisation entity attracts more humans to further reinforce the technological transition. As such, prediction accuracy is enhanced by the identification of *Superorganisations*, those highly productive groups of humans driving forward each technological shift with an increasing reinforcement property.

Figure 6: It is humans, aggregated as *Superorganisations*, whose work not only drives forward each technological shift, but through the Superorganisation entity attracts more humans to further reinforce the technological path. As such, prediction accuracy is enhanced by the recognition of those highly productive groups of humans driving forward each technological shift



The Super Organisation Secret (2025) was authored by Adrian Courtenay, who also serves as portfolio manager of the GA-Courtenay Special Situations Fund

In our 2025 white paper, *Evaluating the Vector Sum Theory of Corporate Output*, we identified that both Charlie Munger late in his career and Elon Musk publicly recognised that the output of a group of humans can be understood as the vector sum of individual contributions, where each worker represents a vector combining productivity (magnitude) and alignment with corporate goals (direction). When this framework is optimised, it not only leads to the form of highly productive organisation that drives forward technological shifts, but also reinforces them as it grows by continuing to be a talent magnet for new hires and a productivity outlier with regard to its growing workforce.

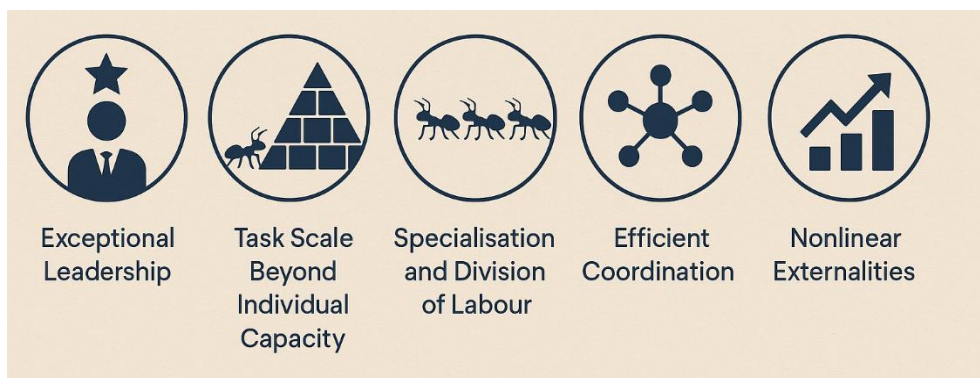
“If you get a certain kind of process going in chemistry, it speeds up on its own. So you get this marvellous boost in what you’re trying to do that runs on and on. So you get a huge boost. You accomplish A – and, all of a sudden, you’re getting $A + B + C$.”

Charlie Munger, quoted in Peter Bevelin’s book *Seeking Wisdom from Darwin to Munger*, 2007²¹

The framework, expanded in my 2025 book *The Super Organisation Secret*, argues that decisive human productivity – including in driving technological shifts – occurs only when variables beyond traditional metrics are prioritised, specifically, those variables which empower the building and maintenance of high-performing, aligned teams.

Five pillars, when operating concurrently in one organisation – exceptional leadership, scale beyond individual capacity, specialised division of labour, efficient coordination, and nonlinear advantages – create a type of lollapalooza effect which exhibits autocatalysis – as success occurs, the *Superorganisation* attributes drive an increasing reinforcement property.

Figure 7: The five pillars of *Superorganisation*²²



The significance of the presence or not of *Superorganisations* reinforcing a technological transition is clearer when we illustrate with an example. Consider, in the current era, two domains in which a technological transition is in place, space exploration and undersea exploration. Whilst both domains hold significant promise in terms of prospective addressable market sizes, as well as each domain meeting all of the predictive requirements so far laid out by this white paper, only one of the two domains – space exploration – is led by *Superorganisations*, SpaceX and its subsidiaries.

Now suppose either domain meets an unanticipated yet significant *interference* barrier, for example regulatory interference, it is only the domain led by *Superorganisations* that will have the strongest additional reinforcement characteristic to overwhelm the interference variable.

In the case of SpaceX, the company will continue obsessively toward its mission, and owner Elon Musk, overseeing the world’s largest communication platform X.com, can apply significant indirect pressure to address regulatory obstacles. The outcome: predictive accuracy of a technological transition reaching its end point is higher when the transition is led by *Superorganisations*.

Figure 8: Charlie Munger justified Berkshire Hathaway's investment in BYD by recognising many of the criteria consistent with *Superorganisation* status²³

"Wang Chuanfu went to engineering school, got a PHD, and then started off by borrowing \$300,000 from the Bank of China. And he starts going into the small batteries for cell phones which was totally dominated by high-tech Japanese firms. And he succeeds in grabbing about a third of that market from a standing start of zero. And the Japanese had all the patents but he won the intellectual property rights through litigation. And that litigation happened in Japan. He was a very remarkable man doing an almost insanely ambitious thing."

Berkshire doesn't normally do this kind of venture capital stuff, but we departed from our standard methods and did this one. So why invest in tech now? Because some models got so powerful that we thought we could make predictions. And if I said I could lift 800 pounds you'd laugh, but if you saw me do it a few times you'd change your mind. The same assessment happened with Wang Chuanfu and BYD.

And out of that, today he has 200,000 employees and a huge lithium battery plant. Last month he sold 10,000 electric cars in China which is more than Tesla sold. And of course nobody's hardly heard of Build Your Dreams "BYD". It's an interesting company.

BYD is in a position, on purpose, to benefit from this electrification trend in the world. It's been very helpful to them that people are dying on the streets of Beijing because they can't breathe the air. They have to go to electric cars in Beijing.

"I warn you the founder of BYD, Wang Chuanfu, has hired 180,000 engineering graduates, all at the top of their classes and so you get a remarkable aggregation of human talent. This is not some unproven highly speculative activity. The place is full of fanatics, and by the way they're a big supplier to Apple and Huawei. And they're a very satisfactory supplier to those companies.

The speed with which they really mastered the technology. It's just unprecedented. I've never seen anybody like it. For example, they got into the auto business in 2003, and produced their first car in 2005. That's only seven years ago. It really is a quite unusually talented group of people who were able to manage, to solve a problem like no other group of people I've ever seen. It may be a small company but its ambitions are large and I don't want to bet against 177,000 Chinese engineers led by Wang Chuanfu."

And we had prior for long felt that the auto industry – it's just a business that is too tough, you've got entrenched worldwide competitors, they are not going away."

Charlie Munger, aggregated comments, annual meeting of Daily Journal, 2013, 2015, 2018

The educational background of the founder CEO is science/technical

The prioritisation of early profitability as a business-minded protagonist

A remarkable self-confidence yet which is also calibrated to reality

A brilliant, high energy, marketeer

The founder CEO is objectively brilliant across multiple domains

A network effect exists, promoting and solidifying winner-takes-all results

Network Effects

The company grows by a positive feedback loop that displaces incumbents

Excellence in attracting top talent to the venture

The company has created a product that establishes and dominates an entirely new industry

Successful communication of a cult-like belief in their cause

The company's brand meets 'super-brand' standards and has an outstanding CEO

Excellence in graduate recruitment & accelerated mentoring by senior engineers

A dedicated, high work rate, corporate culture

A competitor set with deficiency in vector sum competency

Strength in talent retention through workplace culture and experience delivering durable competitive advantage

The emergent industry grows due to a lasting force, often tied to enhancements in computing efficiency

MOORE'S LAW

A workforce of individuals who are alignable

The company's products show 'last mover advantage', naturally guarding against new innovation from others

Motivated and dedicated, they believe their product can significantly improve the world

The central organising force of Superorganisations – exceptional leadership

The first pillar of *Superorganisation* is exceptional leadership, typically through founders who continue in active leadership roles at the corporation, and whose brilliance would not be available to its same workforce should they each choose to operate on a solo basis.

Such a team of exceptional founders and leaders will possess objective brilliance, a science/technical background, strong skills as marketeers, excellent reputations, access to capital (through their excellence, but also through their oversight of an understandable business) indefatigable energy, and strong communication skills. These in combination should be evidenced to result in decisive, confident decision making which can be effectively conveyed both internally and externally, yet which also remains calibrated to reality.

That these are in place should be also evidenced by the achievement of profitability early, in that brilliance combined with science/technical expertise should allow the core, novel product from the company to be the result of innovation that has no equivalent, and such possesses monopolistic economics even at the early stage.

Companies founded by individuals with exceptional technical expertise and marketing acumen also tend to create products and brands that dominate their categories through what can be called ‘*performance authenticity*.’ This authenticity manifests when a product becomes the undisputed leader in its performance-oriented domain—much as SpaceX represents in rocketry, or Ferrari represents in high-performance sports cars. Such products do not merely compete; they define excellence within their categories.

The objective brilliance of the leadership team of the company is also evidenced by excellence in marketing and communication, with the result that their product exhibits *super brand* characteristics. When all super brand attributes are in place, the resultant brand and product is observed, seemingly effortlessly, as travelling across borders to sell across regions worldwide.

Sufficient brilliance by the founding team of a company also acts to enhance the merit level of the company’s prospective workforce by encouraging in the highest performing graduates a strong urge to join the company, by triggering in candidates the human capacity for worship, and in the corporate sense a form of *calling*. As George Soros notes, it is the power of *attraction*, in this case through brilliance in leadership, that exceeds the power achieved through more traditional forms of authority. As such, when a founding team exhibits brilliance, they bless their venture with an additional lever through which talent attraction is achieved.

“Hard power may be needed for conquest and self-protection, but it is the power of attraction – soft power – that ensures the stability of an organisation.”

George Soros, writing in 2014²⁴

Leadership should also exhibit excellence in workforce development, including by their workforce organised to deliver specialisation through division of labour, and accelerated mentoring resulting in efficiency of raising new joiners to each level of specialisation. Furthermore, the high merit level of the workforce should be reflected in their aptitude in accelerating productivity using tools and technology.

Figure 9: Exceptional founders and leaders will possess objective brilliance, a science/technical background, strong skills as marketers, excellent reputations, access to capital, indefatigable energy, strong communication skills²⁵



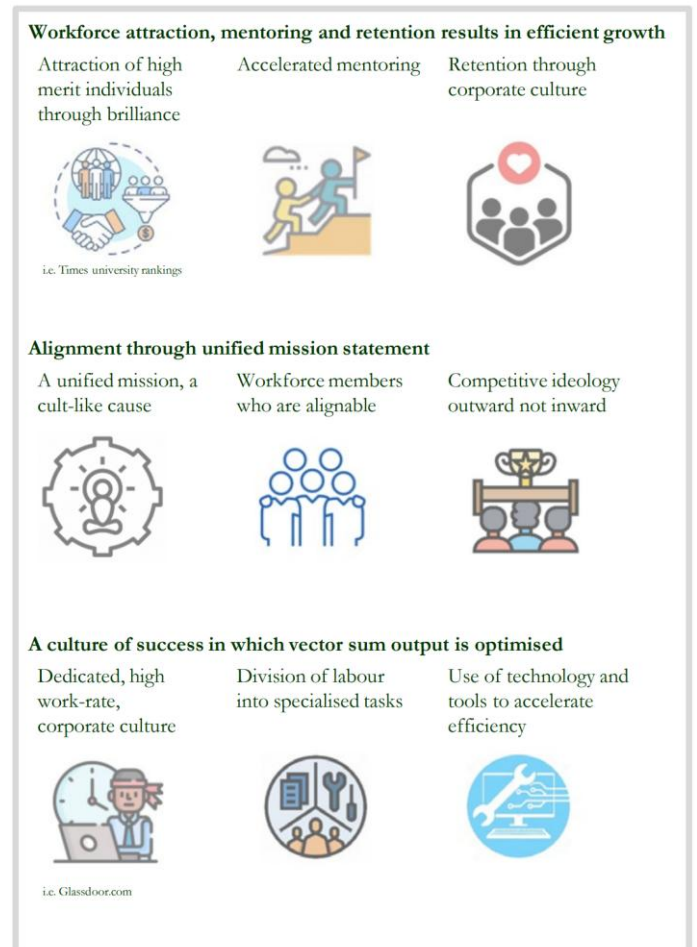
Figure 10: Performance authenticity: leadership in a performance critical domain²⁶



Figure 11: Super brand characteristics, anchored by three pillars²⁷



Figure 12: The excellence in leadership flows into excellence in workforce attraction, mentoring, retaining and aligning²⁸



The second pillar of Superorganisation – scale beyond individual capacity

The second pillar of *Superorganisation* is a form of intrinsic motivation by self-assessed effectiveness, that the corporate goals exhibit meaningful scale requirements, and as such, the goals are only achievable by its workforce aligning with the scaling corporation rather than pursuing the same goals by breaking away as individuals.

Scale requirements may be characterised by concurrent force magnitude (for example, a construction project), social aggregation (funding a spacecraft), or intellectual throughput (coding software ecosystems). Unless the scale requirement is reached, extra effort is wasted; once it is reached, the task suddenly becomes possible. Only a group, through workforce volume amplification, can aggregate sufficient labour (physical, social or cognitive) to cross such thresholds.

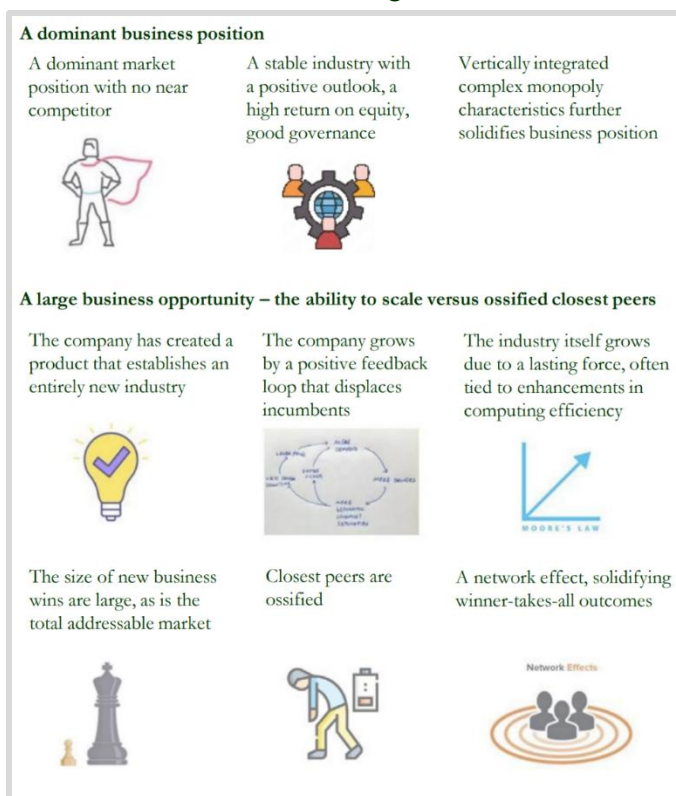
In order to capture scaling attributes, a corporation first requires dominance: an existing dominant niche position with no near competitor, and accreted in the best case by a vertically integrated complex monopoly business design which further solidifies dominance relative to otherwise new entrants. Furthermore, the industry within which the company operates should offer stability characteristics (examples include high customer switching costs i.e. with mission-critical software, essential products or services, irreplaceable physical assets, advanced intellectual property, network effects, or an installed base with recurring revenues, all leading to pricing power above inflation) and possess a positive outlook, providing an attractive substrate on which scaling can occur.

A second requirement is that the company has the ability to displace those existing, incumbent companies which possess the targeted industry opportunity. This ability is at its strongest when the protagonist company has created a product which establishes an entirely new way to address the customer requirement. Essentially a new industry model, yet addressing a customer need which has so far been provided by incumbents using a different approach.

Displacement of incumbent companies is also accelerated when the emergent company grows by a positive feedback loop, which, as it generates revenue by its superior product displacing its extant competitor set, generates revenue which is then reinvested into continuing product improvement, and as a result then precipitating additional, accelerating displacement of incumbents.

Furthermore, scaling can go to greater heights when the company’s new innovation changes the size of the industry itself. Before books could be purchased online, the friction cost of having to walk to your local book store

Figure 13: Great business opportunities by their nature target scaling. Success probability is optimised by a dominant market position, and a large prospective opportunity characterised by network effects with existing incumbents ossified²⁹



handicapped industry size by limiting impulse purchases and breadth of selection. Yet, with the launch of online book retail, friction costs were removed, and the scale of the industry benefitted from additional growth tailwinds.

A further tailwind to scaling potential occurs when the new industry is also growing due to an enduring force, often tied to enhancements in computing efficiency. For example the growth of mobile phones and tablets did not just allow a new medium on which books could be purchased, including to read digitally, but additionally the volume of mobile phones and tablets underwent rapid scaling due to the cheapening costs of computing itself, combined with the incentive for unit purchases by the growing number of other services offered on these devices at the same time.

Other powerful scale drivers include that the addressable market size is very large, particularly as measured relative to the size of the dominant yet emerging entrant. A further driver is that the incumbent competitor set is comprised of large, ossified companies with inferior technology, and unable to adapt with sufficient agility to the changes in the market opportunity that the new entrant has pioneered. And finally, that a network effect exists, resulting in winner takes all outcomes, and as such handing to the growing company the prospect of a monopolistic capture of the opportunity presented.

The third pillar of Superorganisation – the requirement for multi-domain specialisation

The third pillar of *Superorganisation* is a further form of intrinsic motivation by self-assessed effectiveness, that the business opportunity is too complex or multifaceted for smaller workforces to take on, and as such, the goal is only achievable by its workforce aligning with the multifaceted corporation rather than pursuing the same goals by breaking away in smaller groups.

Complex objectives necessitate complex co-ordination and specialised teams deployed concurrently through division of labour, enabling parallel processing capabilities unavailable to companies with a lower volume of workforce. The complex multifaceted skill-set requirement may be embedded in the product's technology itself (in the case of technology companies) or technology-like aspects of its manufacture (i.e. outside of technology companies, sourcing complex ingredients for the manufacture of a beverage, or the complex manufacturing process for high performance premium sports trainers). The result is that the company needs to encompass multiple teams each with differentiated skill sets across multiple domains. Of course, one or a group of individuals could, alternatively in principle become polymaths, but for most humans competence falls as the number of domains rises. Specialised teams side-step that trade-off by distributing expertise.

The requirement for specialised teams working in tandem is heightened by the continually changing competitive landscape that the breakthrough corporate strategy must address. The implication is that it is not enough for a company to possess one or two specialised teams which succeed in complex co-ordination to produce an exceptional product. The company must divide labour between specialists relating to its current product, and specialist innovators, who build the next generation of products suited to a dynamically changing world.

In assessing the innovation potential of a company, outcomes are optimised when division of technical labour is organised by a unifying mission statement, such that the right product segments are selected for innovation work. Furthermore, the ability of the company to innovate should not just be a conjecture from the perspective of the outsider investor, but also be evidenced by the track record of the company.

Several facets can thereon heighten the pace of successful innovation. The first is when there is a free-riding ability that comes from the counterparties of the company (i.e. community of users) also being innovators, and the company is able to replicate their innovations for the “official version” that the company produces.

An additional example of the need for a company to divide labour into groups which specialise in specific tasks emerges when we consider marketing, an entirely different activity than the technological processes that will either originate product or underpin the innovation process.

Similar to innovation, marketing is optimised firstly by the leadership of the company establishing a clear mission statement, to orientate marketing to a defined set of core values, and as such provides for internal teams to deliver efficient marketing campaign without the need for expensive outside consultants.

The efficiency of marketing is also raised when the company has dedicated community of followers, and whose contact details are possessed electronically, such that marketing can occur direct-to-consumer, and at low friction cost. In these circumstances marketing, in the AI era, can retain messaging bespoke to the customer’s attributes, yet still carried out as mass, highly efficient e-campaigns.

A critical component for high efficiency marketing is also the performance authenticity of the product itself, and demonstrating the need for division of specialised labour the delivers excellence across the company’s activities.

When a product possesses *performance authenticity*, marketing is simplified from a manipulative approach which would otherwise aim to conceal the *absence of performance authenticity*, to an efficient and reductive approach requiring marketing to merely communicate “*notice what we actually are*” messaging. This not only empowers efficient, low cost marketing, but additionally opens up those “*notice me*” marketing avenues which would otherwise be closed, such as high impact, yet also cheap to execute, marketing stunts.

Figure 14: Innovation is optimised when division of technical labour is organised by a unifying mission statement, and thereon heightened by counterparties, targeting and standardisation³⁰

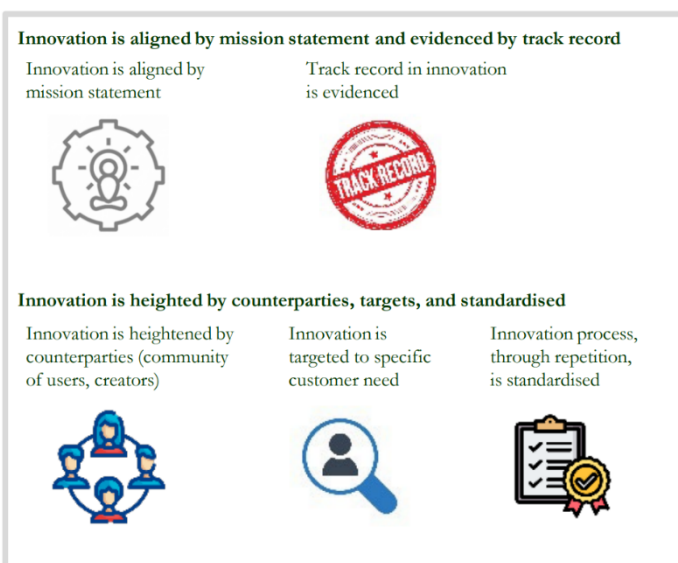


Figure 15: Marketing is optimised by clarity of mission statement and a dedicated community of followers, combined with performance authenticity and super brand attributes³¹



The division of labour product-marketing-innovation flywheel furthermore allows the company to achieve a form of *last mover advantage*. By delivering a continuing innovation flywheel, the core product has sufficient continuing innovation funding for its successive variants to remain the “last” impactful new launch in its domain, much in the way the continually upgraded and improved Apple iPhone has remained the “last” mobile phone in the modern era with no meaningfully successful challenge from new entrants.

In the best cases of exceptional, breakthrough companies, the achievement of the complex, multi-faceted division of labour business position and multi-domain specialisation results in adherence to the “*vertically integrated complex monopoly*” characteristics that Peter Thiel has advocated for³², which, when assembled in a co-ordinated and sophisticated manner, lead to tremendous competitive advantage.

Thiel’s astute recognition appears to be that a dominant *business position* alone is insufficient. Netscape dominated internet browsing, and had a meaningful software network effect advantage in its favour and a position close to monopoly; however, it was displaced by Internet Explorer, as a result of a bundling strategy by Microsoft, albeit Microsoft’s Internet Explorer was then displaced by Chrome, offered by Google for free.

Instead, enduring business dominance – extrapolating Thiel’s mode of thinking – also requires the business itself to possess a design which is non-replicable relative to the competency limitations of otherwise prospective founding teams who would seek to establish a competing new entrant. Non-replicability in design occurs when a business possesses complex assembly and multi-faceted co-ordination properties which, outside of the rare conditions that led to the isolated origin of the extant dominant player, are not realistically possible to re-create.

The fourth pillar of Superorganisation – efficient co-ordination: through feedback loops empowering unrestricted information exchange, corporate goals crisply shared, concurrent with flattening unnecessary hierarchies and minimising conflict

The fourth pillar necessary for *Superorganisation* status is that the costs of co-ordination do not exceed the benefits delivered by the same condition. One individual co-ordinates perfectly with themselves. However, as efforts move to group organisation, the larger the group, all else equal, the greater the form of co-ordination tax that must be paid: time, communication overhead, and addressing the risk of infighting. If that tax grows larger than the benefit of aggregation, the group’s advantage disappears. *Superorganisations* therefore prioritise efficient co-ordination through feedback loops empowering unrestricted information exchange, corporate goals crisply shared, concurrent with flattening unnecessary hierarchies and minimising conflict.

Best-in-class organisations begin with the premise that information is an asset whose value compounds only when it flows upwards as freely as it flows down. They therefore build deliberately decentralised feedback loops that tap the ‘distributed intelligence’ of the whole workforce—including frontline staff whose daily contact with customers gives them data executives cannot buy elsewhere. When thousands of individual observations are aggregated, contradictory errors cancel out while consistent truths reinforce one another, creating a “wisdom-of-crowds” signal that is often sharper than expert analysis – and far faster. Firms that institutionalise this flow gain a dynamic advantage over rivals whose ossified hierarchies still reward information hoarding instead of disclosure.

Best-in-class corporate leaders then ensure corporate goals are crisply shared, and model the creed by their own example working visibly harder – Musk’s 80-to-100-hour weeks or Bezos’ ultimatum that his Amazon colleagues

must work “long, hard and smart”, and not two of the three³³ – and thereby convert exhortation into social proof. Because the rule set is direct, field decisions can be pushed to the edge without the fear that initiative will veer off-mission – saving on bureaucratic cost.

Best-in-class corporate leaders also focus on recruiting people who are naturally “alignable”. Rather than cloning their own iconoclasm, these founder-CEOs seek complementary talent: professionals who are technically formidable yet happy to hitch their expertise to a shared mission and to direct any residual anti-authoritarian impulses outward – at complacent incumbents, not inward at corporate process or colleagues.

Successful emergent organisations by definition will almost always have had anti-authoritarian founders. In establishing the company, these founders will have used their wisdom and technical expertise to *reject authority*. However, for their workforce the astute selection is of individuals who are more comfortable having used their wisdom and technical expertise to *select authority* (namely, of the company’s founding team).

Alignment is further pre-wired through purpose: a mission that is bigger than compensation and strong enough to create what Charlie Munger called “*fanatics*” at BYD or the type of workforces that investors can today observe at Musk companies such as SpaceX and Tesla. Because the goal is morally resonant, elite candidates self-select, lowering screening costs and raising the probability that new hires already share the firm’s values.

Once through the door, those values are rendered in a single, memorable maxim: *maximum industriousness directed toward unified goals*. Maurice Saatchi’s dictum that “*simple ideas enter the brain quicker and stay there longer*”³⁴ defines the organisational operating system; simplicity shrinks the communicative overhead that handicaps large teams.

Alignment is also accelerated through corporate structure. Superorganisations replace expensive mid-career hiring sprees with graduate pipelines that pair novices with senior engineers in intensive mentorship loops. The novice gains tacit knowledge in weeks, the mentor gains status, and the firm gains culture carriers whose habits were formed in-house rather than imported at a premium. At the same time, in the modern era management will be able to mine AI-enriched workforce data to check – in real time – whether the workforce remains true to the mission, turning what used to be an anecdotal art into a measurable science.

This mentorship model creates mutual benefits: mentees gain accelerated alignment orientation and skillset development, mentors feel uplifted by passing on their knowledge to the next generation, and the company achieves high-quality talent and alignment at optimal cost.

“A young person needs a mentor, and I think an older person in a way needs a protégé. It’s a mutually beneficial arrangement that has gone on for millennia and I think it will continue.

So you’re lucky to have such older persons that are willing to devote a lot of energy and time to bringing along younger people. I’ve tried to do that. Whether I’m as good as the ones I had I have no idea, but I’ve also spent time, doing that..”

Tom Perkins, co-founder of Kleiner Perkins, interviewed by John Hollar, 2011³⁵

Finally, retention closes the loop. Achieving a strong retention culture requires management to recognise that it is the maintenance of a high merit workforce that will be a dominant factor sustaining the company’s competitive advantage. Breakthrough organisations profiled treat people, not patents, as their “deepest moat” and design workplaces accordingly: artisan cultures, flat hierarchies, minimal process and a felt reciprocity between individual growth and collective success. Because the same individuals stay, a shared language and set of heuristics accumulate, trimming the metabolic cost of coordination year after year.

The fifth pillar of Superorganisation – when the corporation achieves sophisticated use of tools and technology, network effects and free-riding tailwinds

Finally, *Superorganisation* status is further enhanced when the corporation achieves sophisticated use of tools and technology, network effects and free-riding tailwinds. When a firm possesses sophisticated tools and technology in-house, its teams gain a further advantage that would be absent in other career trajectories, and productivity is accreted accordingly. Positive feedback loops can also occur through the company’s business position exhibiting network effects, leading to winner-takes-all advantages. When both sophisticated tools and technology and network effects loops are present, a square law results, and the probability of the exceptional business outcome is commensurately raised.

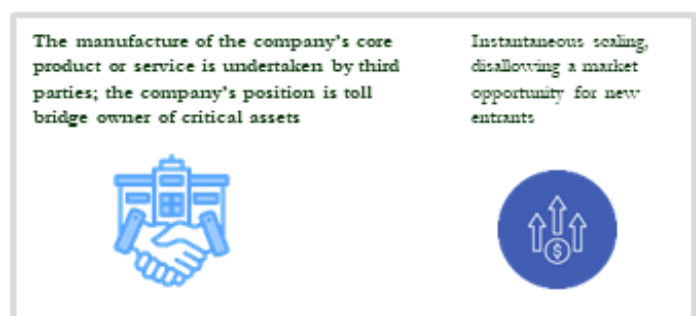
Free-riding on external resources and innovation further enhances corporate output. A common free-rider example occurs in the case of a network effect company that serves as a platform for its counterparties, be it an outsourced workforce community in the case of Uber, Netflix the Apple app store, or Microsoft Windows, or a community of users in the case of Facebook or X.com, to post content that forms the basis of the value of the platform itself. The platform exists as the gatekeeper, but the work is done by others at no ongoing cost to the platform once market leadership is established.

An additional form of free-riding occurs when a company possesses highly advantageous and unique intellectual property driving high profit margins, and yet therefore also sufficient to fund the outsourcing of manufacturing to others. The result is the use of third party capital for manufacturing, allowing the growing company to ride at a comparably low fee relative to its own high profit margins.

A rapidity advantage is realised by both forms of free-riding. When a modest size, but growing company free-rides on a considerably larger ecosystem (be it an outsourced workforce community or third party manufacturing) it also gains the ability to achieve instantaneous scaling.

The result is, as the emergent business opportunity grows, and invariably at different rates, the emergent company can instantaneously scale its size proportional to its market opportunity at each stage. This forms an additional lever in favour of *last mover advantage*. It is very hard for a new entrant to enter the market when there is never a “gap” period during which the market demand for a product exceeds its supply.

Figure 16: Outsourcing of manufacturing allows the growing company to free ride on external capital³⁶



Reinforcement from the sheer size of the technological project goal possessing such high financial magnetism that abandonment is inconceivable

Finally, reinforcement of a technological trajectory can occur from the sheer size of the technological project goal possessing such high financial magnetism that abandonment is inconceivable. As such, reinforcement is a property of technological shifts when the technological project, at its end point of completion, is of the form that would change humanity on a massive scale, reshaping how billions of people live, work, or relate to their environment.

The significance of this form of reinforcement lies in linking addressable market size to incentive for continued development work in relation to the technology in question. A technologically feasible prediction with little prospective impact on humanity will have a constrained addressable market and may be abandoned by the groups pursuing it. The outcome: the forecast is falsified.

The internet passed this magnitude filter because it transformed how information flows across humanity. SpaceX Starlink passes likewise—as Figure 17 illustrates, an addressable market size with potential telecommunications sector alone revenues exceeding \$2 trillion, connecting billions of humans to the internet who would otherwise remain unconnected, and therefore providing substantial incentive for continued engineering effort.

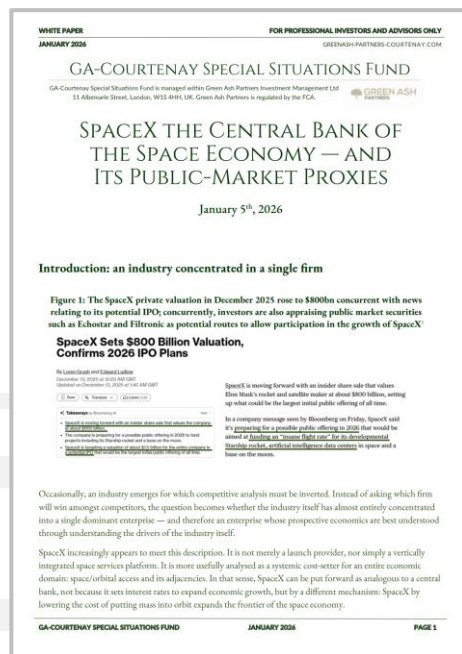
Figure 17: SpaceX Starlink passes the usefulness filter; a vast addressable market size with potential revenues upward of \$2 trillion, presented for illustrative purposes only, provides substantial incentive for continued work by SpaceX engineers in driving the technological change forward³⁷

Revenue build to 2040: Starlink targets a vast addressable market size with potential revenues upward of \$2 trillion

Important Notice — Illustrative Scenario Only. The revenue analysis below represents one internally modelled scenario based on publicly available data and GA-Courtenay assumptions. It is not a forecast, projection, or target return. Actual outcomes will differ materially. The figures should not be relied upon as the basis for any investment decision.

Bandwidth (YE2035)	Bandwidth (YE2040)
1,536 Cumulative Falcon 9 launches	1,998 Cumulative Falcon 9 launches
93 Launches pa over trailing 5yrs	93 Launches pa over trailing 5yrs
3,167 Cumulative Starship launches	9,596 Cumulative Starship launches
512 Launches pa over trailing 5yrs	1,286 Launches pa over trailing 5yrs
12 Launches per Starship per year	12 Launches per Starship per year
43 Required Starship fleet size	107 Required Starship fleet size
5 Depreciation years	5 Depreciation years
9 Required annual Starship manufacturing volume	21 Required annual Starship manufacturing volume
16,000 Falcon 9 payload capacity (kg, LEO)	16,000 Falcon 9 payload capacity (kg, LEO)
125,000 Starship payload capacity (kg, LEO)	125,000 Starship payload capacity (kg, LEO)
24,568,000 Falcon 9 cumulative upmass (kg)	31,968,000 Falcon 9 cumulative upmass (kg)
395,835,000 Starship cumulative upmass (kg)	1,199,500,000 Starship cumulative upmass (kg)
29 Falcon 9 V2 mini satellites per launch	29 Falcon 9 V2 mini satellites per launch
60 Starship V3 satellites per launch	60 Starship V3 satellites per launch
552 Falcon 9 V2 mini satellite unit mass (kg)	552 Falcon 9 V2 mini satellite unit mass (kg)
2,083 Starship V3 satellite unit mass (kg)	2,083 Starship V3 satellite unit mass (kg)
1,960 Mass (kg) per satellite	2,032 Mass (kg) per satellite
100 Falcon 9 V2 mini satellite capacity (Gbps)	100 Falcon 9 V2 mini satellite capacity (Gbps)
1,000 Starship V3 satellite capacity (Gbps)	1,000 Starship V3 satellite capacity (Gbps)
928 Average Gbps per satellite	970 Average Gbps per satellite
2,900 Falcon 9 payload capacity (Gbps)	2,900 Falcon 9 payload capacity (Gbps)
60,000 Starship payload capacity (Gbps)	60,000 Starship payload capacity (Gbps)
44,530 Cumulative Falcon 9 V2 mini satellites launched	57,942 Cumulative Falcon 9 V2 mini satellites launched
190,001 Cumulative Starship V3 satellites launched	575,760 Cumulative Starship V3 satellites launched
234,530 Cumulative satellites launched	633,702 Cumulative satellites launched
33,411 Annual number of satellites launched (trailing 5 yr average)	79,834 Annual number of satellites launched (trailing 5yr average)
31,117 Cumulative Falcon 9 V2 mini satellites de-orbited	44,530 Cumulative Falcon 9 V2 mini satellites de-orbited
36,360 Cumulative Starship V3 satellites de-orbited	190,001 Cumulative Starship V3 satellites de-orbited
67,477 Cumulative satellites de-orbited	234,530 Cumulative satellites de-orbited
13,413 Total active Falcon 9 satellites	13,413 Total active Falcon 9 satellites
153,641 Total active Starship satellites	385,759 Total active Starship satellites
167,053 Total active satellites	399,172 Total active satellites
1,341,250 Gross capacity active Falcon 9 satellites (Gbps)	1,341,250 Gross capacity active Falcon 9 satellites (Gbps)
153,640,800 Gross capacity active Starship V3 satellites (Gbps)	385,759,200 Gross capacity active Starship V3 satellites (Gbps)
154,982,050 Maximum constellation bandwidth (Gbps)	387,100,450 Maximum constellation bandwidth (Gbps)
5% Average capacity utilisation at any one time	5% Average capacity utilisation at any one time
7,749,103 Total effective deployed bandwidth (Gbps) of Starlink fleet	19,355,023 Total effective deployed bandwidth (Gbps) of Starlink fleet
5% Actual capacity utilisation of mobile usage vs 150Mbps headline speeds	5% Actual capacity utilisation of mobile usage vs 150Mbps headline speeds
1,033.2 Implied number of mobile phone bandwidth user units (m)	2,580.7 Implied number of mobile phone bandwidth user units (m)
12.9% Starlink bandwidth as a percent of total worldwide population if all used mobiles	32.3% Starlink bandwidth as a percent of total worldwide population if all used mobiles
929.9 Starlink revenue (\$bn) if all effective bandwidth at mobile pricing	2,322.6 Starlink revenue (\$bn) if all effective bandwidth at mobile pricing

GA-Courtenay white paper: SpaceX the Central Bank of the Space Economy (2026)



Revenue build estimates represent estimates by GA-Courtenay research based on public domain disclosures from SpaceX, analyses by third party research groups including Mach-33 Research, and other public corporate domain disclosures from the telecoms and satellite sectors. Full methodology available on demand from GreenAsh Partners. For total addressable market size estimates see GA-Courtenay white paper: SpaceX the Central Bank of the Space Economy (2026).

The magnitude filter also eliminates trajectories that may appear large in theoretical market terms but modest in human significance. It also elevates trajectories that might appear niche by conventional metrics but carry outsized importance. Space colonisation has no current market—there are no customers on Mars—yet it passes the magnitude filter because its successful execution would fundamentally alter humanity.

Understanding the “what” of prediction. *The what is the expected value.*

Compile prediction on an expected value basis

Finally, *A Unified Framework For Modelling the Future* requires a mode of thought calibrated to its end goal, which is capital allocation weighted to the highest probability outcomes. As such, future facing projects must be judged using expected value calculations—probability of success multiplied by magnitude of impact.

A ten percent probability of civilisational-scale impact may yield higher expected value than a ninety percent probability of incremental improvement. If making humanity multiplanetary is worth tens of trillions of dollars in avoided existential risk then a ten percent chance of contributing to that outcome justifies enormous commitment.

This framing shifts the relevant question from “will this succeed?” to “is the expected value positive?” The latter accommodates uncertainty in a way the former cannot. It acknowledges that most ambitious attempts will fail while maintaining that the portfolio of ambitious attempts, taken together, produces positive returns precisely because the wins, when they occur, are of such magnitude.

From the capital allocator’s perspective, this reframes the prediction mission. The goal is not explicitly to predict probabilities with precision, but to allocate capital and effort into projects that combine prediction with the highest expected value outcome.

For those not positioned to pursue civilisational-scale ventures, Musk articulates a corollary: optimise for usefulness.

“I don't think everyone needs to try to solve some big world-changing problem. Really it should be a usefulness optimisation—is what I'm doing as useful as it could be? Even if something's making people's lives only slightly better but it's a large number of people, the area under the curve is quite good.”

Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

The “*area under the curve*” metaphor captures expected value reasoning at a range of scales. Impact equals magnitude of improvement multiplied by number of people affected. A small improvement affecting many may aggregate to greater total utility than a large improvement affecting few. The same logic that justifies pursuing ten-percent chances of civilisational transformation also justifies pursuing high-probability chances of modest but widespread improvement.

Both SpaceX and Tesla survived their near-death experiences in 2008. The fourth rocket launch succeeded; the Tesla financing round closed on Christmas Eve, “the last hour of the last day that it was possible.” In retrospect,

these outcomes validate the original expected value calculation—but only if one accepts that validation exists at the expected value of capital allocation and effort expended, not in any single trajectory’s prediction accuracy.

Similarly, Aschenbrenner advocates for prediction through an expected value lens. He does not say “AGI will arrive in 2027”; he says it is “strikingly plausible.”

“I believe in probability distributions. I think it’s very plausible that all this stuff doesn’t go anywhere and it asymptotes and we’ll hit limits... I also just think it’s possible that on a fairly short time horizon, basically this decade, we’re gonna hit extremely extremely powerful systems.”

“I think most people just have screwed up probabilities all across the board. 10% chance is actually very high. There’s basically a zero percent chance climate change kills everybody. 10% is actually very high... and people don’t think like that.”

Leopold Aschenbrenner, interview with The Center for the Study of Partisanship and Ideology, 2024

For the capital allocator, the principle translates directly. Mastery of *A Unified Framework For Modelling the Future* does not imply a resultant capital allocation orientated to a single scenario. Instead, the outcome should be the construction of portfolios that capture value across the probability distribution. However, not all probable scenarios are equiprobable—capital allocation should be weighted toward the highest expected-value outcomes, which will often mean allocating disproportionately to tail scenarios whose magnitude, if realised, dwarfs their probability.

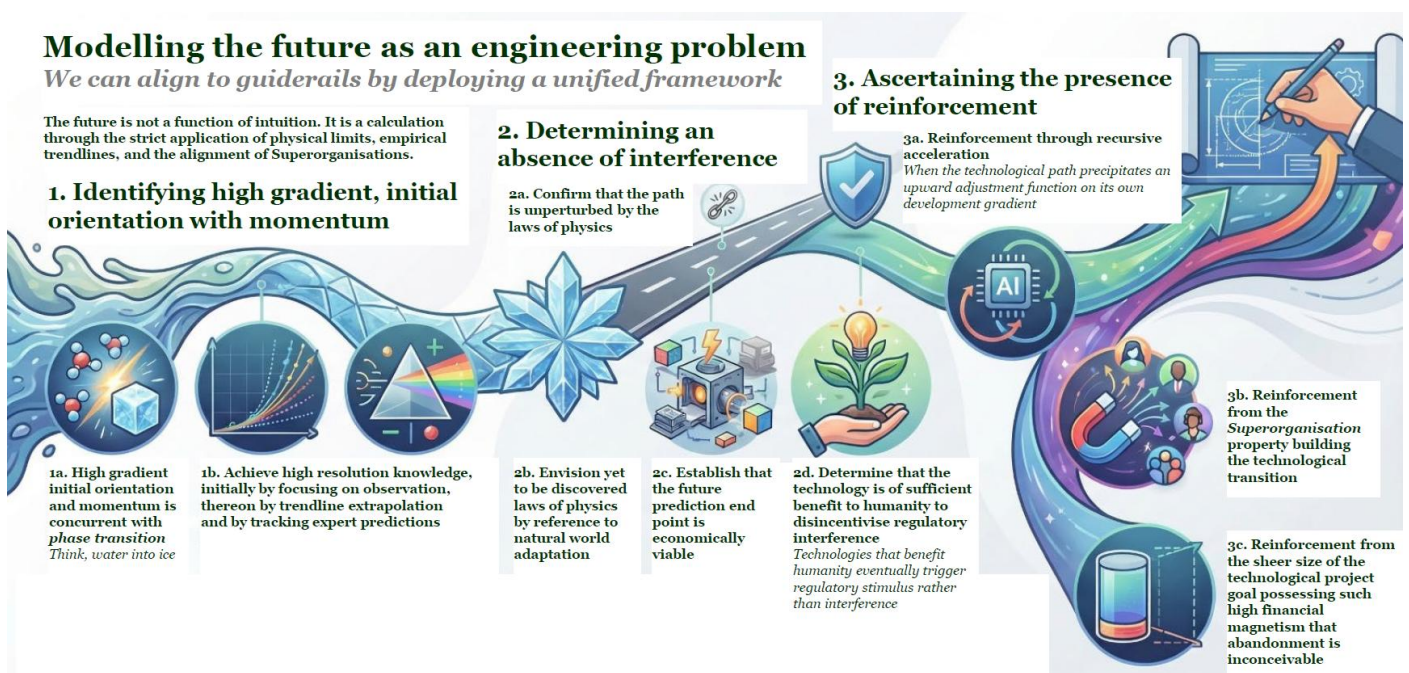
Conclusion

As this white paper has outlined, prediction can be coherently structured around the presence of an existing orientation possessing momentum, the absence of interference, and the presence of reinforcement. When these properties are present, *A Unified Framework For Modelling the Future* can be used to master the assignment of probabilities to various scenarios to a far greater degree than may be widely recognised as possible. It is the route to this mastery that this white paper has explored.

The unified framework presented in this white paper represents a synthesis of knowledge which, whilst a number of its pillars focus on the principles that have been argued for by practitioners including Charlie Munger, Elon Musk and Leopold Aschenbrenner, it also contains conjectures wholly unique, including the *Superorganisation* concept as a reinforcement mechanism, and we are aware of no other literature which aggregates all of our principles into *A Unified Framework For Modelling the Future* that this white paper presents.

The principles presented compliment, rather than oppose, each other to result in an overall stronger predictive framework. In identifying high gradient, initial orientation with momentum, phase transition recognition is prioritised concurrent with deploying multiple templates for detection concurrently: frontier proximity, trendline observation and extrapolation, and expert prediction tracking. In determining an absence of interference, principles including laws of physics adherence, natural world example equivalency, end point economic viability, disincentivisation of regulatory interference, and acceptance of failure outcomes within controlled parameters as part of the success path, are necessary. Finally, in ascertaining the presence of reinforcement, several identifiers are put forward including recursive acceleration from the technological path itself, reinforcement through regulatory acceleration, reinforcement through the identification of *Superorganisations* building the technological progression, and reinforcement through financial magnetism aspect of the sheer size ambition of the technological project.

Figure 18: An infographic for our framework – modelling the future as an engineering problem



These principles, applied with discipline, offer a systematic approach to what might be called *a biophysics of foresight*—an approach to the future grounded not in intuition or optimism but in systematic analysis of what the laws of physics and of nature permit, and what human organisational dynamics propel.

Yet embedded within this framework is a recognition that points toward both a paradox and its resolution.

Musk describes prediction as a "*contact sport*." He does not merely observe the future; he participates directly in its construction. Comparably, Aschenbrenner contends that "*situational awareness*" is necessary for prediction, an equivalency to direct participation, or at least, a close working relationship with those directly participating.

Direct participation in building the future confers epistemic advantages. The participant receives information unavailable to external observers: the texture of engineering challenges, the pace of iteration, the responsiveness of supply chains. Each failure reveals something about the failure space; each success confirms something about the path forward.

However, direct participation also imposes a severe constraint. If prediction requires participation, then prediction requires time—not merely the time to analyse, but the time to recruit teams, raise capital, manage operations, and sustain commitment through years of uncertain progress. Therefore predictive accuracy is purchased at the price of personal bandwidth.

Furthermore, the direct participant faces additional handicaps. Their equity participation is ungeared, illiquid, unhedged, and their investment allocations have high switching costs.

The principles of prediction may be transferable; yet this is of limited value if the handicaps of direct participation cannot also be overcome.

It is this aspect that sets our approach apart. Musk recognises – as a participant – that he sits in a position able to "*introduce something new to the probability stream*"; he is in the stream, not watching from the shore. The crucial inference follows: if the participant can introduce something new to the probability stream, it is because humans drive technological change. Technology does not advance automatically. It advances because specific groups of exceptional people make it advance—and it degrades when those people cease their work, as Musk himself observes regarding the regression from Apollo to the Space Shuttle to the temporary American inability to reach orbit at all.

This recognition—that humans are the engine of phase transitions—carries a methodological implication. If technological change is driven by humans, then the prediction of technological change requires the identification of those humans. Not abstractly, but concretely: which teams, led by which founders, organised in which structures, are driving each phase transition forward? The *Superorganisation* framework provides exactly this identification. It recognises that a company's trajectory is the vector sum of individual employee contributions, where each worker represents a vector combining productivity and alignment with corporate goals. Technological shifts occur when high-performing, aligned teams operating as *Superorganisations* scale. *Predicting the future is therefore inherently tied to the recognition of Superorganisations.*

This insight resolves the paradox at the heart of the Musk and Aschenbrenner frameworks.

Consider the outside minority investor who has genuinely mastered the principles laid out by this white paper—who can identify using our predictive framework initial orientation possessing momentum, an absence of interference and the presence of reinforcement – and combines their work with a deep understanding of the *Superorganisation* criteria. Such an investor possesses the analytical framework without the operational burden.

They essentially substitute their lower resolution level of Musk’s engagement in “*contact sport*”, and Aschenbrenner’s “*situational awareness*”, with their higher resolution level – through our approach – of *Superorganisation* recognition. They can recognise the founders who have assembled and aligned exceptional talent toward goals with the transformational impact of defining technological change, and allocate capital as a passive, outside minority.

This position offers structural advantages unavailable to the participant. Equipped with our framework, the investor’s time is not consumed by engineering management and operational execution. Where the participant must focus on a handful of ventures demanding his direct attention, the outside minority investor can distribute capital across a portfolio of *Superorganisations*, each driving its own technological phase transition forward. The portfolio approach converts idiosyncratic risk into systematic exposure to the future. Thereon, the investor can amplify returns through leverage, address risk through hedging, and achieve agility across deployments through liquidity—financial architecture unavailable to the participant whose equity commitment is locked to specific enterprises.

This is the ultimate implication of our framework: if the recognition of *Superorganisations* can substitute for direct participation in building them, then the outside minority investor who masters this recognition achieves an efficiency position conceptually superior to the builder or participant themselves³⁸.

Footnotes

1. GA-Courtenay white paper, Macro Protection Within A Unified Framework For Capital Allocation 2024 [link]
2. Elon Musk, The Future We’re Building – and Boring, April 2017 [link]
3. Daniel Kahneman, Thinking, Fast and Slow [link]
4. Nassim Nicholas Taleb, The Black Swan: The Impact of the Highly Improbable [link]
5. Elon Musk, speaking in interview with Kevin Rose, 2012 [link]
6. Leopold Aschenbrenner, Situational Awareness white paper, June 2024 [link]
7. Leopold Aschenbrenner, Existential Risk and Growth [link]
8. Elon Musk, The Future We’re Building – and Boring, April 2017 [link]
9. GA-Courtenay white paper, Evaluating The Vector Sum Theory of Corporate Output 2025 [link]
10. Adrian Courtenay, The Super Organisation Secret [link]
11. Figure: GA-Courtenay research and public domain disclosures
12. Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015 [link]
13. Elon Musk on his first startup Zip2, Startup Archive [link], PayPal ownership structure [link], PayPal IPO prospectus shareholder list [link]
14. Leopold Aschenbrenner, interview with The Center for the Study of Partisanship and Ideology, 2024 [link]
15. Microsoft president says no chance of super-intelligent AI soon, Reuters November 2023 [link]
16. Elon Musk, speaking with Peter Diamandis in December 2025 [link]
17. Figure: GA-Courtenay research and public domain disclosures
18. Leopold Aschenbrenner, conversation with Ben Yeoh, 2022 [link]
19. Figure: GA-Courtenay research and public domain disclosures
20. Alexandre Dumas, Le Corricolo, 1843 [link]
21. Charlie Munger, quoted in Peter Bevelin’s book Seeking Wisdom from Darwin to Munger, 2007 [link]
22. Figure: GA-Courtenay research
23. Charlie Munger, aggregated comments, annual meeting of Daily Journal, 2013, 2015, 2028 [link]
24. The Tragedy of the European Union, George Soros, 2014 [link]
- 25-31. Figures: GA-Courtenay research
32. Peter Thiel, speaking at Stanford University, 2015 [link]
33. Amazon’s 1997 letter to shareholders [link]
34. Maurice Saatchi “simple ideas enter the brain quicker and stay there longer” [link]
35. Tom Perkins, co-founder of Kleiner Perkins, interviewed by John Hollar, 2011 [link]
36. Figure: GA-Courtenay research
37. GA-Courtenay white paper, SpaceX the Central Bank of the Space Economy 2026 [link]
38. This conclusion describes a theoretical analytical framework only and does not constitute a representation regarding the past or future performance, or return expectations, of the GA-Courtenay Special Situations Fund or any other fund managed by Green Ash Partners.

Legal Disclosure

This white paper is an independent analysis conducted by *Green Ash Partners Investment Management Ltd* (“*Green Ash Partners*”) and is intended for discussion, educational and informational purposes only. This white paper is distributed to a limited number of subscribers and is not for public distribution, reproduction, or use without the express written permission of *GreenAsh Partners*.

The views expressed herein represent the opinions of *GreenAsh Partners* as of the date hereof. *GreenAsh Partners* reserves the right to change or modify any of its opinions expressed herein at any time and for any reason and expressly disclaims any obligation to correct, update or revise the information contained herein or to otherwise provide any additional materials.

The views expressed in this white paper are solely those of *GreenAsh Partners* and its authors and do not necessarily reflect the views of the companies whose logos, trademarks, or brand names are included. No representation or warranty, express or implied, is made as to the accuracy, completeness, or reliability of the information contained in this document.

The information contained herein is based on publicly available information, including filings made with the securities and exchange commission (“SEC”) and other sources, as well as *GreenAsh Partners*’ analysis of such publicly available information. *GreenAsh Partners* has relied upon and assumed, without independent verification, the accuracy and completeness of all data and information available from public sources, and no representation or warranty is made that any such data or information is accurate. *GreenAsh Partners* recognizes that the companies referenced herein may possess confidential or otherwise non-public information that could lead them to disagree with *GreenAsh Partners*’ views and/or conclusions and that could alter the opinions of *GreenAsh Partners*, were such information known. No representation, warranty or undertaking, express or implied, is given as to the reliability, accuracy, fairness or completeness of the information or opinions contained herein, and *GreenAsh Partners* and each of its directors, officers, employees, representatives and agents expressly disclaim any liability which may arise from this white paper and any errors contained herein and/or omissions here from, or from any use of the contents of this white paper.

Except for any historical information contained herein, the information and opinions included in this white paper constitute forward-looking statements, including estimates and projections prepared with respect to, among other things, anticipated company performance, the value of company securities, general economic and market conditions, and other future events. You should be aware that all forward-looking statements, estimates and projections are inherently uncertain and subject to significant economic, competitive, and other uncertainties and contingencies, and have been included solely for illustrative purposes. Actual results may differ materially from the information contained herein due to reasons that may or may not be foreseeable. There can be no assurance that any securities referenced in this white paper will trade at the prices that may be implied herein, and there can be no assurance that any opinion or assumption herein is, or will be proven, correct.

This white paper and any opinions expressed herein should in no way be viewed as advice on the merits of any decision with respect to any of the companies or securities referenced herein. This white paper is not (and may not be construed to be) legal, tax, investment, financial or other advice. Each recipient should consult their own legal counsel, and tax and financial advisers as to legal and other matters concerning the information contained herein. This white paper does not purport to be all-inclusive or to contain all of the information that may be relevant to an evaluation of the companies or securities referenced herein, or the matters described herein.

This white paper does not constitute (or may not be construed to be) a solicitation or offer by *GreenAsh Partners* or any of its directors, officers, employees, representatives or agents to buy or sell and securities referenced herein, or an offer to sell an interest in funds managed by *GreenAsh Partners*. This white paper does not constitute financial promotion, investment, advice or an inducement or encouragement to participate in any product, offering or investment, or to enter into any agreement with the recipient. No agreement, commitment, understanding or other legal relationship exists or may be deemed to exist between or among *GreenAsh Partners* and any other person, including the parties and individuals referenced herein, by virtue of furnishing this white paper. No representation or warranty is made that *GreenAsh Partners*’ investment process or investment objectives will or are likely to be achieved, or successful, or that *GreenAsh Partners*’ investments will make any profit or will not sustain losses. Past performance is not indicative of future results.

Funds managed by *GreenAsh Partners* currently beneficially own and/or have an economic interest in, and may in the future beneficially own and/or have an economic interest in the securities of any of the securities named in this white paper, as well as securities in adjacent industries. *GreenAsh Partners* intends to review its investments on a continuing basis and depending upon various factors, including without limitation, their financial position and strategic direction, the outcome of any discussions with them, overall market conditions, other investment opportunities available to *GreenAsh Partners*, and the availability of their securities at prices that would make the purchase or sale of such securities desirable, *GreenAsh Partners* may from time to time (in the open market or in private transactions) buy, sell, cover, hedge or otherwise change the form or substance of any of its investments to any degree, in any manner permitted by the law, and expressly disclaims any obligation to notify others of any such changes. *GreenAsh Partners* also reserves the right to take any actions with respect to its investments as it may deem appropriate.

GreenAsh Partners has not sought or obtained consent from any third party to use any statements or information contained herein. Any such statements or information should not be viewed as indicating the support of such third party for the views expressed herein. All trademarks and trade names used herein are the exclusive property of their respective owners.