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.

Elon Musk's Principles For Predicting The Future, And Their Implications

January 12th, 2026

Introduction

At the pinnacle of high-return capital allocation lies what might be called "god mode": the ability to make a series of accurate predictions about the future. This is not hyperbole but tautology. Equity valuation is optimised precisely when accuracy regarding the consistency, duration, and growth of future cashflows is correctly estimated. To predict the future is to master valuation itself.

The contention that an investor can achieve such mastery 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 reach this ambitious goal, beginning by identifying the theoretical conditions under which absolute determinism can occur. The future path of an object through space can be accurately predicted when an existing orientation possesses momentum or is subject to reinforcement, and interference in its path is absent. 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 has momentum, and there is no interference. 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 action, and their subsequent behaviour will be subject to interference from countless new inputs. Now consider the same person jumping from a plane mid-flight, again wearing only swimming trunks—absent a parachute. The initial condition appears similar, yet here there is both momentum and reinforcement, and the absence of a parachute eliminates plausible interference. The outcome – sadly for its protagonist – becomes deterministic.

Figure 1: Reductively "similar" human trajectories can be characterised by very different forecast accuracy. Prediction requires all of – initial orientation is subject to either momentum or reinforcement, and an absence of interference.

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 circumstances in financial markets where initial orientation, subject to momentum or reinforcement, may continue unperturbed by interference—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 human technological development must embrace a more realistic ambition: 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 2017[2]

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 discovered[4]. 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 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, 2012[5]

In October 2015, Elon Musk appeared at Stanford University's Future Fest, in conversation with venture capitalist Steve Jurvetson[6]. 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.

His framework has proven remarkably durable. The principles Musk articulated in 2015 illuminate decisions he made in 1995, explain ventures he launched in 2002, and remain visible in his strategic positioning today. They form a toolkit not merely for anticipating technological change but for participating in its creation—what Musk describes as prediction being a "*contact sport*."

Yet herein lies a paradox that this white paper will ultimately resolve. If prediction is truly a contact sport—if forecasting the future and constructing it are inseparable—then the predictor appears bound to the same constraints as the builder: the investment of time, the burden of execution, the limitation to those few ventures one can personally oversee.

However, this white paper argues that the full implications of Musk's framework, carefully examined, suggest otherwise. They point toward a position available to the investor who masters these principles that may prove superior even to Musk's own—a path to amplified participation in the future's construction, freed from the handicap of direct involvement. That resolution awaits in our conclusion.

This white paper distils the core principles from Musk's Stanford address, supplemented by our own research, and organised into a framework practitioners can apply to their own strategic thinking. These 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. They build upon one another while operating in parallel. Together, they constitute what might be called a physics of foresight—an approach to the future grounded not in intuition or optimism but in systematic analysis of what the laws of nature permit, what human needs demand, and what timing makes most profitable.

Phase transition events and their detection

<u>The gradient of technological development is highest at points of *phase transition*</u>

The first and perhaps most powerful element of Musk's framework is his treatment of timing. Musk notes that 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 irreversibly[7].

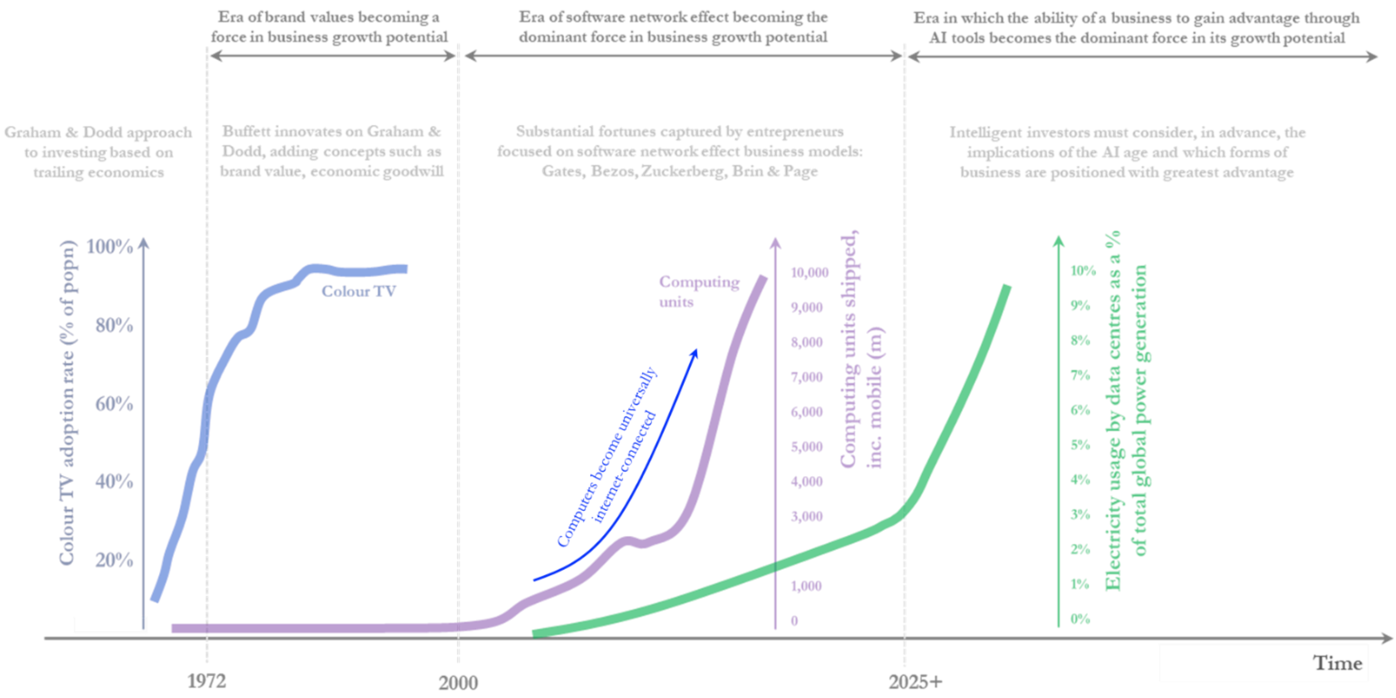

The metaphor Musk employs is water freezing into ice. The transition from liquid to solid is not gradual; it occurs abruptly when conditions cross a threshold. Before the threshold, entirely liquid. After, entirely solid. The transition point represents a qualitative shift, not merely a quantitative 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, 2015[8]

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[9]. 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[10] and PayPal[11] enabled subsequent engagement with capital-intensive hardware ventures that would otherwise have been impossible.

The detection of phase transitions

Recognising that phase transitions matter is one thing; detecting them as they approach 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.

> *"People are just not used to thinking about exponential compounding –*
>
> *because in the beginning you just don't see anything."*
>
> Bill Ackman, speaking in March 2020[12]

> *"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

Musk offers a practical methodology for detecting acceleration before it becomes obvious: *track the predictions rather than the technology itself*.

Rather than measuring technological progress against some absolute scale—which requires defining and quantifying capability in domain-specific terms—monitor whether expert predictions about arrival dates 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

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 2023[13]

> *"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 2025[14]

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.

Phase transitions are driven by super organisations

Musk's comment earlier in this white paper revealed an additional element of his thinking from which we can derive a further principle. He stated: "*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.*"

The explicit recognition is that humans drive phase transitions. By corollary, it is Musk and his workforces whose efforts may 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 2017[15]

The implication: accurate detection of phase transitions should not be merely abstract, but grounded in the explicit recognition that highly productive groups of humans are collaborating to drive each transition forward.

Figure 4: It is it is humans whose work drives forward each phase transition, and as such the verification of phase transitions is also by the recognition of those highly productive groups of humans – working as *Super Organisations* – driving forward each technological shift[16]

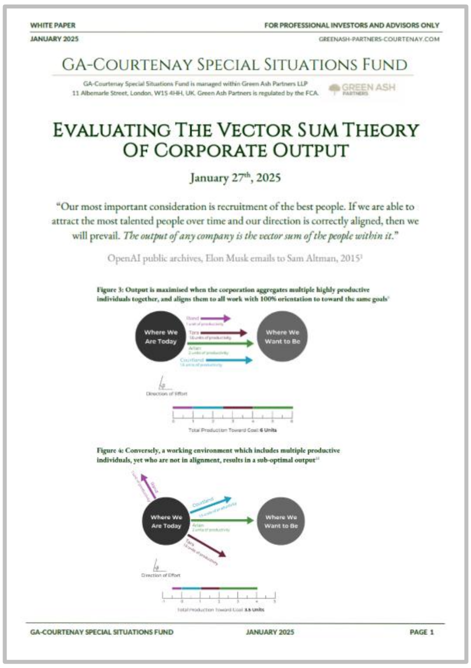
WHITE PAPER FOR PROFESSIONAL INVESTORS AND ADVISORS ONLY
JANUARY 2025 GREENASH-PARTNERS-COURTENAY.COM

GA-COURTENAY SPECIAL SITUATIONS FUND

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

EVALUATING THE VECTOR SUM THEORY OF CORPORATE OUTPUT

January 27th, 2025

"Our most important consideration is recruitment of the best people. If we are able to attract the most talented people over time and our direction is correctly aligned, then we will prevail. *The output of any company is the vector sum of the people within it.*"

OpenAI public archives, Elon Musk emails to Sam Altman, 2015[1]

Figure 3: Output is maximised when the corporation aggregates multiple highly productive individuals together, and aligns them to all work with 100% orientation to toward the same goals

Figure 4: Conversely, a working environment which includes multiple productive individuals, yet who are not in alignment, results in a sub-optimal output

GA-COURTENAY SPECIAL SITUATIONS FUND JANUARY 2025 PAGE 1

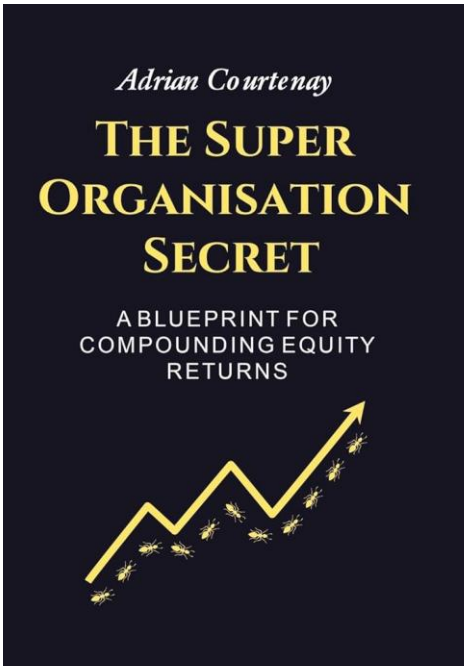

In our 2025 white paper, *Evaluating the Vector Sum Theory of Corporate Output*[17], we identified that both Charlie Munger late in his career and Elon Musk publicly recognised that the output of a group of humans can be understood as the vector sum of individual contributions, where each worker represents a vector combining productivity (magnitude) and alignment with corporate goals (direction). This framework, expanded in my 2025 book *The Super Organisation Secret*[18], argues that decisive human productivity – including in driving technological shifts – occurs only when variables beyond traditional metrics are prioritised, specifically, the building and maintenance of high-performing, aligned teams.

In identifying phase transitions, therefore, a critical second-order implication of Musk's framework is to verify the transition by identifying the presence of high-performing, aligned, and scaled groups of humans orientated as super organisations driving the phase transition forward—whether under Musk's direction or others'.

At the time of writing, the AI phase transition, the expansion of the space economy frontier transition, and the transition to autonomous military systems in our appraisal all meet this test. For other contended transitions—the shift to nuclear fusion in energy generation, or the potential shift to human life extension—the identification of high-performing, aligned, and scaled groups of humans driving the transition we find less well evidenced.

Establish that the future prediction end point is both possible and economically viable

<u>The future prediction end point must be demonstrated as permitted by the laws of physics</u>

Whilst Musk's framework for identifying phase transitions serves as an instructive initial principle, its applicability as a tool for future prediction varies with the maturity of the transition itself. Detection using the principles laid out becomes far harder at the point Musk himself is originating the phase transition—arguably the case at earlier-stage businesses such as Neuralink and The Boring Company.

Furthermore, there is a difference between recognising a phase transition and correctly predicting its end point. 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, Musk's predictive framework is to begin with certainty about physical possibility.

Instead of analysing markets, competitors, or business models, establish that the predicted destination is permitted by the laws of physics. Humans – with the right support infrastructure – can survive on another planet. 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 are facts about what the universe allows.

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 5: 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 achievable[19]

The research programs delivering computing efficiency improvements are pushing against an "open door" tailwind of long-term enhancement; datacentre efficiency per watt remains 10^6× short of the human brain, and a further 10^4× short of the limits of physics.

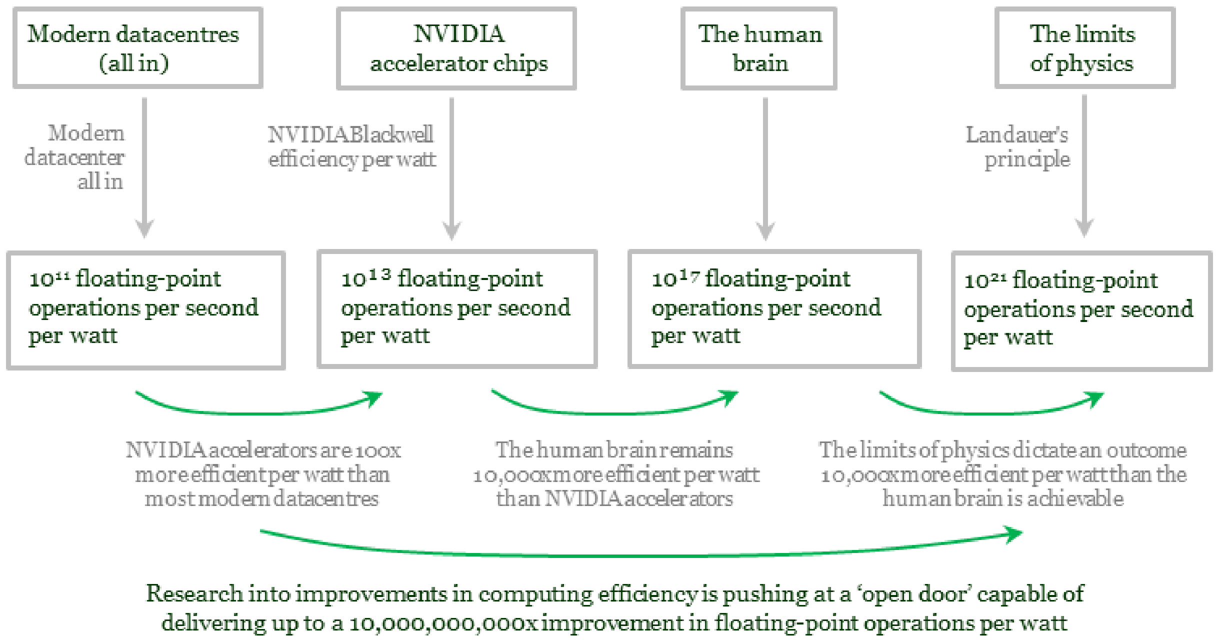

The practical consequence is an expansion of the end point space. When you know something is physically possible, you are liberated from the implicit constraints embedded in how things are currently done. The fact that rockets have historically been expensive does not mean rockets must be expensive—only that rockets have been expensive so far. The fact that electric vehicles have historically had limited range does not mean they must have limited range—only that battery energy density has historically been insufficient.

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.

<u>Establish that the future prediction end point is economically viable</u>

Once a destination is confirmed as physically possible, Musk's next step is to determine what reaching it should cost according to first principles analysis.

This means 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. When the gap between theoretical and current costs is large—as it was for orbital launch, where the ratio between actual and theoretical cost approached fifty to one[20]—the opportunity may justify enormous effort and risk.

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.

Apply the usefulness filter, and the magnitude filter

Whilst first principles analysis can identify gaps between physical possibility and current reality, not all gaps are equally worth pursuing. Musk describes a filtering process for selecting which opportunities merit commitment—two sequential tests that together determine whether to engage.

The usefulness filter and its link to the cadence of regulatory permissioning

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 pass it will eventually be developed and widely adopted, because incentives for development align with beneficial outcomes. Market forces, governmental support, and social approval all push in the same direction. The remaining uncertainty concerns timing and implementation, not whether development will occur.

Yet even when a technology functions as a general solution—performing its intended task reliably across a wide range of conditions—it may not be permitted to deploy. This is the domain of regulatory frameworks, liability structures, social acceptance, and institutional inertia.

> *"The technology works... there's a sort of second question as to when regulators would approve it."*
>
> Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015

This distinction creates a two-stage model for technological prediction. The first stage concerns capability: when will the technology work? Here, first principles reasoning, exponential tracking, and phase transition analysis apply most directly. Physics and engineering determine what is possible and when. The second stage concerns permission: when will deployment be allowed?

Regulatory bodies move cautiously. Legal frameworks evolve through precedent and legislation—processes that resist acceleration. Social acceptance shifts gradually as familiarity increases and risks become better understood. Ignoring the permission stage leads to a characteristic prediction error: overestimating early adoption speed while underestimating eventual penetration.

Figure 6: Decades-long efficiency improvements in, and increased frequency of utilisation of, passenger air flights have been underpinned by the usefulness filter; at the level of the state, the firm and individual – increasing passenger flights are useful. Accurate future predictions with regard to the widespread adoption of autonomous driving must pass the same usefulness test[20].

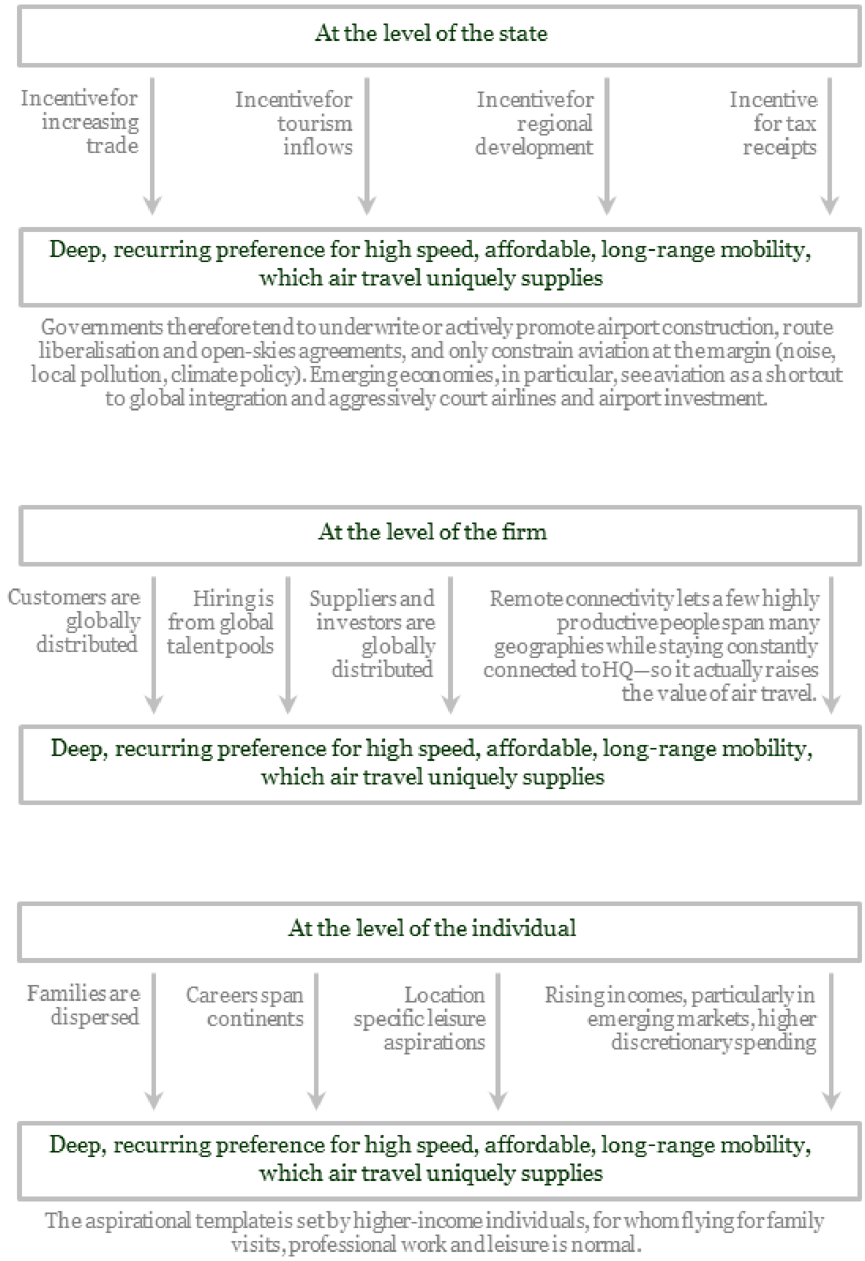

The forecaster who sees only capability predicts rapid deployment the moment the technology works, then is surprised when regulatory barriers delay rollout. The same forecaster may also underestimate long-term penetration, failing to recognise that once permission is granted, adoption can accelerate rapidly as accumulated capability is finally unleashed.

Autonomous vehicles illustrate both dynamics. Technical capability for highway driving reached commercial viability years before regulatory frameworks were prepared to permit the initial deployments observable today. The gap between "it works" and "it is allowed" stretched longer than many forecasters anticipated. Yet when full permissioning eventually arrives—through demonstrated safety records, evolved regulations, and shifting public acceptance—the penetration curve may prove steeper than linear extrapolation would suggest, precisely because so much capability has accumulated awaiting deployment.

For the practitioner, this implies maintaining separate forecasts for capability and permission, 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. The most sophisticated forecasts account for both stages.

The principle also highlights a strategic opportunity. During the gap between capability and permission, competitive positions can be established that yield advantages once permission arrives. The firm that develops and refines technology while awaiting regulatory approval will be far better positioned than one that waits for permission before beginning development.

The magnitude filter and its inverse relationship with technological abandonment risk

The magnitude filter asks: is this a problem that, if solved, would change humanity on a massive scale? This is a question about civilisational impact. Does solving this problem reshape how billions of people live, work, or relate to their environment?

The filter's significance 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 of the future is falsified.

The internet passed this filter because it transformed how information flows across humanity. SpaceX Starlink passes likewise—as Figure 7 illustrates, an addressable market size with potential 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.

The magnitude filter also eliminates problems that may appear large in theoretical market terms but modest in human significance. It also elevates problems 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's trajectory.

The combination of the usefulness filter and the magnitude filter produces a distinctive predictive advantage. The filters allow predictions to remain calibrated to the societal realities that influence the temporal gradients by which outcomes are delivered at the same time as underwriting with confidence the directional trends in relation to high-magnitude, and therefore higher certainty outcomes.

Figure 7: 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 forward[21].

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

Bandwidth (YE2035)		Bandwidth (YE2040)	
1,536	Cumulative Falcon 9 launches	1,998	Cumulative Falcon 9 launches
93	Launches pa over trailing 5 yrs	93	Launches pa over trailing 5 yrs
3,167	Cumulative Starship launches	9,596	Cumulative Starship launches
512	Launches pa over trailing 5 yrs	1,286	Launches pa over trailing 5 yrs
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	Cumluative Starship V3 satellites launched	575,760	Cumluative 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 5 yr average)
31,117	Cumulative Falcon 9 V2 mini satellites de-orbited	44,530	Cumulative Falcon 9 V2 mini satellites de-orbited
36,360	Cumluative Starship V3 satellites de-orbited	190,001	Cumluative 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)

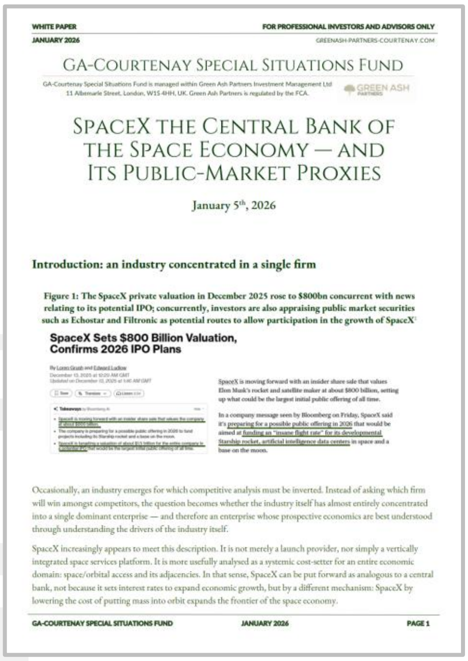
WHITE PAPER
FOR PROFESSIONAL INVESTORS AND ADVISORS ONLY
JANUARY 2026
GREENASH-PARTNERS-COURTENAY.COM

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.

GREEN ASH PARTNERS

SPACEX THE CENTRAL BANK OF THE SPACE ECONOMY — AND ITS PUBLIC-MARKET PROXIES

January 5th, 2026

Introduction: an industry concentrated in a single firm

Figure 1: The SpaceX private valuation in December 2025 rose to $800bn concurrent with news relating to its potential IPO; concurrently, investors are also appraising public market securities such as Echostar and Filtronic as potential routes to allow participation in the growth of SpaceX[1]

SpaceX Sets $800 Billion Valuation, Confirms 2026 IPO Plans

SpaceX is moving forward with an insider share sale that values Elon Musk's rocket and satellite maker at about $800 billion, setting up what could be the largest initial public offering of all time.

In a company message seen by Bloomberg on Friday, SpaceX said it's preparing for a possible public offering in 2026 that would be aimed at funding an "insane flight rate" for its developmental Starship rocket, artificial intelligence data centers in space and a base on the moon.

Occasionally, an industry emerges for which competitive analysis must be inverted. Instead of asking which firm will win amongst competitors, the question becomes whether the industry itself has almost entirely concentrated into a single dominant enterprise — and therefore an enterprise whose prospective economics are best understood through understanding the drivers of the industry itself.

SpaceX increasingly appears to meet this description. It is not merely a launch provider, nor simply a vertically integrated space services platform. It is more usefully analysed as a systemic cost-setter for an entire economic domain: space/orbital access and its adjacencies. In that sense, SpaceX can be put forward as analogous to a central bank, not because it sets interest rates to expand economic growth, but by a different mechanism: SpaceX by lowering the cost of putting mass into orbit expands the frontier of the space economy.

GA-COURTENAY SPECIAL SITUATIONS FUND JANUARY 2026 PAGE 1

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).

Tolerate failure outcomes within controlled parameters

Musk's principles for predicting the future are orientated to calibrate to reality in an additional manner: there are fundamental limits to what simulation and analysis can recognise. 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, Musk's framework treats instances of failure not as an aberration to be avoided 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

Musk translates this into a software analogy: writing a suite of modules that can never be run together, cannot be tested on the target computer, and must execute without bugs on first deployment to an entirely different machine. The gap between testing environment and deployment environment is irreducible.

This understanding transforms the response to failure. SpaceX experienced three consecutive rocket failures before achieving its first successful orbital launch. The fourth attempt—the last Musk could afford—succeeded. Those failures were not evidence of incompetence; they were the statistically expected cost of exploring a high-dimensional failure space. Each failure eliminated possibilities and narrowed the path toward success.

The practical implication: prediction in complex domains must be structured as an iterative process. Initial predictions may be wrong. The value of making them is not in their accuracy but in their falsifiability. A wrong prediction, clearly stated and tested against reality, produces information. That information updates subsequent predictions. This is what Musk means by calling prediction a "*contact sport*"—one must be directly engaged with the businesses building the future, experiencing successes and failures firsthand, adjusting predictions based on what that engagement reveals.

Assess projects on an expected value basis

Finally, Musk advocates judging future-facing projects 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.

Conclusion

The principles distilled from Elon Musk's framework for predicting the future form a coherent methodology that transcends intuition. They begin with the recognition that absolute determinism in human affairs is rare, but not absent—that under specific conditions of momentum, reinforcement, and freedom from interference, future outcomes may become knowable with near-certainty. For most human endeavours, however, prediction must content itself with the more realistic ambition of advantage in mastering the assignment of probabilities.

This mastery proceeds through a sequence of analytical steps. First, identify phase transitions—those moments when technological development crosses thresholds and fundamental states of affairs shift rapidly and irreversibly. Detect these transitions by monitoring whether expert predictions about arrival dates are converging toward the present faster than time itself is passing.

Second, establish that the predicted endpoint is permitted by the laws of physics, then apply first principles reasoning to determine what that endpoint should cost. The gap between theoretical cost and current reality represents the opportunity.

Third, filter opportunities through sequential tests of usefulness and magnitude: does solving this problem benefit humanity, and would it operate at civilisational scale? Technologies passing both filters will have a far higher probability of being be developed albeit with timing uncertainty remaining.

Fourth, acknowledge the entropic nature of complex systems—failure is statistical inevitability to be managed, not aberration to be avoided. And finally, assess projects on an expected value basis, recognising that a modest probability of civilisational transformation will yield higher expected value than near-certainty of incremental improvement.

These principles, applied with discipline, offer a systematic approach to what might be called the physics of foresight. Yet embedded within this framework is a recognition that points toward both a paradox and its resolution—one that illuminates the second half of this white paper's title.

Musk describes prediction as a "*contact sport*." He does not merely observe the future; he participates directly in its construction. His predictions carry weight precisely because he is not speculating about what others might achieve, but reporting from the front lines of what he himself is building. He recognises 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 Musk 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 that Musk demonstrates but does not fully articulate. 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 *super organisation* 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 operate as super organisations at scale. *Predicting the future is therefore inherently tied to the recognition of super organisations.*

This insight resolves a paradox at the heart of Musk's framework.

Musk's direct participation in building the future confers epistemic advantages. The entrepreneur 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.

But direct participation also imposes a severe constraint. If prediction requires building, 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. From Musk's perspective, predictive accuracy is purchased at the price of personal bandwidth. The principles are transferable; his vantage point appears not to be.

Unless, that is, the investor recognises what Musk himself implicitly demonstrates: that the locus of predictive power is not participation per se, but the identification of the humans driving change. Musk must build because building is how he identifies—and becomes—the super organisation at the frontier. But if an investor can identify super organisations without building them, the constraint dissolves.

Consider the capital allocator who has genuinely mastered these principles—who can detect phase transitions, confirm physical possibility through first principles reasoning, filter through usefulness and magnitude criteria, and critically, verify transitions by identifying the super organisations driving them forward. Such an investor possesses the analytical framework without the operational burden. They can recognise founders capable of assembling and aligning exceptional talent toward transformational goals, and allocate capital as a passive minority participant.

This position offers structural advantages unavailable to the builder. The investor's time is not consumed by engineering management and operational execution. Where Musk must focus on a handful of ventures demanding his direct attention, the investor can distribute capital across a portfolio of super organisations, each driving its own phase transition forward. The portfolio approach converts idiosyncratic risk into systematic exposure to the future. Moreover, a hedge fund investor in this position can amplify returns through leverage—financial architecture unavailable to the entrepreneur whose equity is tied to specific enterprises. The builder's returns are bounded by ownership stakes in ventures they have constructed; the investor's returns can be amplified beyond the builder's proportional to the conviction of their analysis.

If the recognition of super organisations can substitute for direct participation in building them, then the investor who masters this recognition achieves an efficiency position conceptually superior to Musk himself. They capture the analytical insights without the operational, time-allocation handicap. They gain exposure to the construction of the future without the time investment of direct participation. They achieve amplified returns on correctly identified transformations precisely because they are not the ones building them.

This is the ultimate implication of Musk's framework: that the discipline of prediction, mastered and applied systematically, offers the investor an optimised path to capturing the returns the future will generate—returns that may exceed even those available to the builders themselves.

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. Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015 [link]
7. Figure: GA-Courtenay research and public domain disclosures
8. Elon Musk, in conversation with Steve Jurvetson, Stanford University, 2015 [link]
9., 10. Elon Musk on building his first startup Zip2, Startup Archive [link]
11. Elon Musk, statement on X as to PayPal ownership structure [link], PayPal IPO prospectus lists Musk as largest individual shareholder [link]
12. Bill Ackman, speaking March 2020 [link], See Nassim Nicholas Taleb, The Statistical Consequences of Fat Tails [link]
13. Microsoft president says no chance of super-intelligent AI soon, Reuters November 2023 [link]
14. Elon Musk, speaking with Peter Diamandis in December 2025 [link]
15. Elon Musk, The Future We're Building – and Boring, April 2017 [link]
16. GA-Courtenay white paper, Evaluating The Vector Sum Theory of Corporate Output 2025 [link], The Super Organisation Secret [link]
17. GA-Courtenay white paper, Evaluating The Vector Sum Theory of Corporate Output 2025 [link]
18. The Super Organisation Secret [link]
19. Figure: GA-Courtenay research and public domain disclosures
20. Figure: GA-Courtenay research and public domain disclosures
21. GA-Courtenay white paper, SpaceX the Central Bank of the Space Economy 2026 [link]

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.