CELERITY RESEARCH

Towards the Celerity Investment Framework

From repeated performance to a causal investment system

Celerity Research | August 2026

The problem we thought we were solving

Investing has an uncomfortable habit of becoming more complicated as one learns more about it. At first the problem seems straightforward. Find a good business, understand what makes it good, avoid paying too much for it and wait. Much of the best long-horizon investment writing arrives at some version of that proposition. Long-horizon investors add time: if a genuinely exceptional enterprise can continue reinvesting successfully for long enough, surprisingly large outcomes become mathematically possible.

The difficulty is that most descriptions begin somewhere in the middle of the investment problem. A company can have excellent financial results without possessing an enduring explanation for them. It can have an excellent management team without operating in an industry that allows the value it creates to remain with the enterprise. It can possess substantial Strategic Power while slowly losing the organisational capabilities that originally created that Power. It can remain an exceptional enterprise while becoming a poor investment because the market price already assumes an even more exceptional future. An enterprise can also create enormous economic value without the holder of a particular security having an equally secure claim on that value.

Our earlier work on the Comprehensive Investment Framework (CIF) was useful because it widened the field of vision. It brought economic conditions, geopolitics, technology and infrastructure cycles, industry structure, Strategic Power, organisational quality, management, innovation, finance, valuation and portfolio construction into one research programme. The CIF was never published, so there is no reason to preserve its architecture for historical or presentational reasons. Its value now lies in the research it generated and in the questions it forced us to confront.

As that research developed, a weakness became clearer. The subjects could sit beside one another more readily than they explained one another. An economic assessment could be completed, an industry analysed, management evaluated, Strategic Power identified and valuation calculated, yet the causal relationships between those conclusions remained less explicit than they needed to be. Capability Science changed the direction of the work by asking what an organisation had demonstrated that it could repeatedly accomplish.

Repeated performance as evidence

Consider two American home-improvement retailers in the early stages of the warehouse retail revolution. Both understood the do-it-yourself customer. Both operated large stores with broad assortments. Both had knowledgeable staff, established purchasing systems, competent management and evidence of successful store replication. The older business had more stores, greater revenue, a longer operating history and considerably more evidence that its systems worked. On a conventional quality assessment it could reasonably have appeared the safer enterprise.

The younger retailer did not obviously possess a collection of capabilities that the established competitor lacked. The difference appeared in the way the system developed. New stores increased purchasing scale. Purchasing scale supported lower prices. Lower prices strengthened the customer proposition. Additional customers supported further stores. Each store generated

operating knowledge that could be applied to subsequent stores. Capital allocation funded expansion of essentially the same system. Growth was not simply increasing the size of the company; it appeared capable of making the system itself stronger.

The first ten Historical Company Assessments (HCAs) repeatedly produced versions of this problem. Strong eventual performers and credible contemporaneous alternatives often both looked good when reduced to individual capability scores. The simple scoring model did not discriminate nearly as cleanly as expected. The more informative evidence appeared in trajectory, transferability, reinforcement and the effect of scale on the organisational system. That remains a pilot hypothesis rather than a validated conclusion.

Long-Horizon Compounder Research has always taken persistent performance seriously, and rightly so. The mistake is to treat the record as its own explanation. Repeated performance can be treated as the explanation - this is a great business because it has produced great numbers for a long time - or as evidence that something has repeatedly produced unusual outcomes and therefore deserves causal investigation.

An Australian enterprise-software company illustrates the distinction. For many years it produced recurring revenue, high customer retention, sustained growth, substantial cash generation and continuing investment in research and development. The organisational history is more revealing. The company repeatedly rebuilt its software through major changes in computing architecture while preserving a common product architecture and continuing to serve specialised vertical markets. The means changed; the underlying organisational logic remained recognisable. The organisation did not merely repeat an outcome. It reproduced an outcome while the problem changed.

The distinction is particularly important because persistent financial performance can arise from very different causal histories. A company can report ten years of rising earnings because it has become progressively better at solving a difficult customer problem. Another can produce a similar record because an industry remained undersupplied. A third can maintain earnings through acquisitions that compensate for weakening organic performance. A fourth can enjoy a favourable regulatory settlement that protects returns for a period. The historical numbers may look similar while the future implications differ substantially. The record becomes useful only when it is connected to an explanation that can itself be tested.

This gives historical performance several roles rather than one. It can help identify enterprises worth investigating; it can provide evidence that a process has repeated; it can reveal the conditions through which the process has survived; and, once an investment has been made, it can establish the normal relationship between mechanism and outcome against which later evidence can be compared. The same history can therefore support selection, causal inference and monitoring. What it cannot do is relieve us of the need to understand the mechanism.

The Historical Company Assessment comparisons sharpened this point because the controls were deliberately credible. The American software comparison company had a dominant database franchise, strong customer adoption, attractive software economics and substantial commercial scale. The older American home-improvement retailer had a longer history, more stores and

mature systems. The comparison medical-products company possessed decades of orthopaedic experience and an established product portfolio. The comparison membership warehouse retailer had an excellent model, strong purchasing economics and capable management. If we had compared eventual winners with obviously poor businesses, the exercise would have told us very little about the problem an investor actually faces at the time of decision.

At the historical assessment date, the investor does not see a winner and a loser. The investor sees several plausible enterprises, each with evidence in its favour and each with unresolved uncertainty. The discipline of point-in-time assessment matters because later success changes the apparent meaning of earlier facts. A small product decision can look visionary after a company becomes dominant and merely experimental if the company later disappears. The same evidence should not acquire a different meaning simply because we know the outcome.

This is why the pilot result that capability scores were often close is more useful than an artificially clean separation would have been. It suggests that the investment problem is not solved by finding companies that possess more desirable attributes. The harder task is to determine how those attributes operate together, whether they are becoming stronger or weaker, and whether the organisation can transfer what it knows into circumstances that have not yet occurred.

Capability, limits and architecture

An activity is not a capability. Opening a store is an activity. Acquiring a business is an activity. Launching a product is an activity. Increasing research expenditure is an activity. The useful question is whether a process repeatedly produces the intended difficult outcome.

Repeated outcomes under unchanged conditions are useful evidence. Repeated outcomes across changing technologies, economic cycles, geographies, leaders or customer requirements are stronger evidence because the variation helps distinguish the persistent organisational cause from the circumstances surrounding it.

This also imposes a limit on what we can claim. An organisation that has successfully opened fifty stores within two neighbouring regions has not demonstrated that it can open five hundred stores across a continent. A management team that has successfully integrated a series of small acquisitions has not necessarily demonstrated that it can absorb a transformational merger roughly equal to its own organisational complexity.

The acquisition cases made this particularly clear. One North American logistics-software business has spent years making relatively focused acquisitions that add customers, functionality, data, geographic reach or connections to a broader network. A large dental-products combination looked equally persuasive when described strategically, yet substantial impairments followed, senior leadership changed and restructuring became necessary. Strategic adjacency did not establish integration capability.

This is where our earlier conclusion - that capability has demonstrated limits - becomes financially important. A history of integrating acquisitions representing a few per cent of existing revenue tells us considerably less about whether the same organisation can absorb a transaction

approaching half its own size. The underlying capability may be real while the proposed application lies beyond the conditions in which it has been demonstrated.

Capabilities also appear to interact. The proposition emerging from the HCA work is that some relationships may be reinforcing: successful execution generates information or economic resources that improve another capability, which in turn improves the organisation's capacity to execute again. The comparison between two American membership warehouse retailers is useful because both possessed strong capabilities. The provisional difference lay in how tightly customer trust, employee capability, purchasing, capital allocation and organisational learning appeared to connect. This remains a hypothesis, not a settled result.

The idea of demonstrated limits also provides a way to handle one of the persistent tensions in long-horizon investing. Investors necessarily make claims about a future that has not yet been observed. If every conclusion were restricted to what had already been demonstrated, the greatest opportunities would often be recognised only after much of the value had been created. Yet extrapolation becomes dangerous when it is hidden inside apparently factual descriptions of quality.

A more disciplined approach separates the demonstrated capability from the extension we are asking it to make. An organisation may have demonstrated that it can replicate a store format across several regions. International expansion is then an inference about transferability. A medical-products company may have repeatedly commercialised devices within one clinical speciality. Entry into an adjacent speciality is an inference about the transfer of customer knowledge, engineering and sales capability. A software company may have migrated successfully through several technical architectures. A new technology transition is still an inference, but the history gives us a richer basis for judging it than a company that has never had to change its architecture.

This separation has financial consequences because the market often prices the inference as though it were already demonstrated. Rapidly growing enterprises can be valued on the assumption that a local operating system will become national, that a national system will become global, or that success in one product category will transfer into several others. Sometimes those assumptions are correct. The analytical step is to identify where the valuation crosses the boundary between evidence and extrapolation.

Capability architecture may help with that judgement. A capability that depends on one individual, one product, one customer or one technological standard may be strong within a narrow setting and fragile outside it. A capability embedded across customer knowledge, operating routines, learning processes and capital allocation may have more routes through which it can adapt. This does not make the second organisation invulnerable. It does suggest that the location of knowledge and the relationships between capabilities are relevant to durability.

The American software pair illustrates the difference. Both enterprises could scale software and both had strong products. One appeared increasingly organised around a complementary portfolio of operating systems, applications and development tools, with knowledge moving across

those activities. The other remained more concentrated around a highly successful database franchise, while efforts to broaden the product portfolio were less visibly connected through a common technical and organisational system. At the time, either structure could have succeeded. The distinction lies in the evidence available about transferability.

The same issue appears in founder-led organisations. Founder involvement can strengthen strategic coherence, customer understanding and long-term orientation. It can also conceal where capability resides. If product judgement, customer relationships and capital allocation depend heavily on one person, a strong enterprise may still possess limited institutional capability. The relevant evidence is not whether the founder is impressive but whether the organisation becomes capable of reproducing difficult outcomes beyond the founder's direct intervention.

Capability architecture therefore should not become another score. It is better understood as a map of relationships: where knowledge originates, how it moves, which processes depend on which others, what growth does to those relationships and where failure in one part of the system could propagate elsewhere.

Growth, capital allocation and Strategic Power

Growth is an outcome; scaling is an organisational problem. A retailer can increase revenue by opening stores, a software company by acquiring another software company and a medical-device manufacturer by entering another clinical category. None of these actions establishes that the underlying organisation has become more capable. Growth can also consume capability as management attention fragments, products become harder to integrate or complexity increases faster than the enterprise can absorb it.

Capital allocation belongs inside this causal account. An acquisition is an activity. Due diligence and integration constitute processes. Repeated acquisition outcomes provide evidence about capability. Research and development expenditure is an activity; a product-development process that repeatedly produces commercially valuable products across changing technologies is stronger evidence.

Capital allocation should therefore be judged by realised causal outcomes. Did the investment deepen customer knowledge, improve the operating system, transfer an acquired capability into the wider organisation, create or renew Strategic Power, and increase per-share economic value after accounting for the price paid and the capital used? Focused reinvestment is not automatically good. Strategic logic is a hypothesis about what an investment should accomplish; realised outcomes determine whether that hypothesis was correct.

Strategic Power addresses a different economic problem from Capability. Capability concerns the organisation's ability to produce a difficult intended outcome. Strategic Power concerns the mechanism that allows the enterprise to retain an attractive share of the value produced. A well-run company is not necessarily a powerful company, and a powerful company is not necessarily well run.

An efficient airline can possess outstanding operating capability while competition transfers much of the productivity gain to passengers through lower fares. A highly capable supplier can find that a concentrated customer base bargains away much of the surplus. A manufacturer can improve production continuously while its product becomes commoditised. Conversely, an enterprise can inherit a valuable licence, network position or customer dependency and continue earning attractive economics despite mediocre internal capability, at least for a period.

Strategic Power therefore requires its own evidence. The mechanism may lie in customer switching, scale, network structure, scarce resources, brand preference, process replication difficulty or another structural relationship. Power is relational. A company may possess substantial Power over customers and little over suppliers, or considerable Power relative to competitors and almost none relative to government.

A global consumer-products company can retain formidable manufacturing, procurement and distribution capability while one historically important brand gradually loses consumer preference. Consolidated performance may remain respectable because pricing, geographic diversification and other brands obscure the weakening mechanism. The organisation has not forgotten how to manufacture or distribute. What has changed is one source of Power.

The distinction between growth and scaling helps explain why acquisition-led businesses require particular care. Reported revenue can continue rising even when the organic system is deteriorating, because each new acquisition adds another layer of sales. The accounting result can remain attractive while integration costs, product fragmentation and management complexity accumulate beneath it. A serial acquirer deserves confidence not because it has completed many transactions but because successive cohorts show that the organisation can select, price, integrate and improve acquired businesses without damaging the existing enterprise.

The positive logistics-software case is useful because acquisitions can plausibly reinforce a wider network. A customs-compliance product can add functionality to an existing logistics platform; new customer connections can make the network more useful; geographic expansion can increase the relevance of the same system to multinational customers. If those relationships are real, acquisition capability does more than add revenue. It strengthens the system through which the enterprise creates and captures value. Cash conversion, recurring revenue, organic performance and returns on acquired capital then become evidence about whether the reinforcing mechanism is actually operating.

The dental-products merger shows the other side. The strategic adjacency was real, but adjacency did not eliminate integration burden. Two businesses can serve the same customers and still possess incompatible systems, cultures, product-development processes and decision structures. The larger the transaction relative to the existing enterprise, the less useful a history of smaller successful integrations may become. Growth can push a genuine capability beyond the range in which it has been demonstrated.

Capital allocation also connects Capability to Strategic Power. Reinvestment can deepen a customer relationship, increase switching difficulty, improve scale economics, add network

participants or create proprietary knowledge that competitors cannot easily reproduce. In these cases, capital allocation changes the future competitive structure rather than merely increasing the asset base. Conversely, investment can increase capacity in a way that invites competition, dilute organisational focus or purchase growth at a price that transfers most of the prospective value to the seller.

This is why high return on invested capital is informative but insufficient. The investor needs to know how much capital can be redeployed at attractive incremental returns, why those returns should persist and what happens to the mechanism as reinvestment expands. A company earning extraordinary returns on a small inherited asset may be less valuable than one earning moderately lower returns while possessing a long, demonstrable runway for reinvestment that strengthens its capabilities and Power.

The per-share qualification should remain visible throughout. Management can create a larger enterprise without creating proportionate value for existing shareholders. Equity issuance, leverage, acquisition prices and repurchases alter the distribution of the economics. A strategy that looks successful at the enterprise level can therefore disappoint at the security level even before Claim Architecture is considered.

Industry, scarcity and environment

Industry structure and the external environment re-enter the analysis through causal transmission rather than as compulsory top-down categories. An external variable belongs in an investment thesis when we can describe how it enters the system.

For an American education provider whose students depended heavily on federal funding, government eligibility was not regulatory background. It was part of the revenue mechanism. If access to that funding disappeared, customer demand could remain while customers' ability to pay collapsed. The company could retain capable staff, recognised programmes and operating infrastructure and still face an existential problem because an external permission on which the economic system depended had changed.

Our electricity and artificial-intelligence research exposed another type of external relationship. Rapid growth in computation creates demand for electricity, but the existence of that demand does not tell an investor where durable economic returns will appear. Scarcity can arise in generation, transmission, transformers, interconnection, firm capacity, powered land, cooling, accelerators or other parts of the physical system. High returns attract capital, engineering solutions change substitution possibilities and governments intervene. A constraint that appears decisive in one period can become ordinary capacity in another while the bottleneck moves elsewhere.

The existence of a bottleneck is therefore insufficient. The investor has to establish who controls it, whether it can be expanded or bypassed, how quickly capital can respond and whether the owner can retain the scarcity economics before the constraint moves elsewhere. A company operating at full capacity during an industry shortage may report extraordinary margins without possessing a durable source of Power.

Capability itself is conditional. An organisation can remain extremely capable while technological or regulatory change reduces the economic value of what it knows how to do. Adaptability becomes observable when a changed condition is followed by an organisational response that changes the relevant mechanism and preserves or improves the intended outcome.

An American motor insurer illustrates the point. The tools used to assess individual risk changed substantially as data and technology developed. The enduring capability was not a particular pricing model or technology. It was the repeated ability to segment and price risk while preserving underwriting discipline.

The electricity work provides a useful example of why this causal treatment of the environment is preferable to a generic macroeconomic section. Electricity is fundamental to computation, and the rapid expansion of artificial intelligence increases the economic importance of reliable power. Yet electricity cost alone does not appear large enough, in many applications, to constrain the value of cognition in the way a simple input-cost analysis might suggest. The investment problem shifts toward the physical availability, reliability and location of power, and then toward the productivity generated by the computation that power enables.

That chain can create several temporary bottlenecks. Generation may be adequate while transmission is not. Transmission capacity may exist while transformers or interconnection queues delay usable supply. Power may be available regionally while suitable sites with grid connection, cooling and network access remain scarce. Accelerator supply can then become the limiting factor even after electricity is secured. Each constraint has a different supply response and a different set of potential beneficiaries.

An investor who simply concludes that electricity demand will rise has not yet identified an investment. The analysis must follow the scarcity through the system. Who owns the constrained asset? Can new capacity be built? What is the lead time? Can customers substitute geographically or technologically? Does regulation cap the return? Will the owner of the constraint capture the value, or will customers, equipment suppliers, labour or government take most of it? How much of the scarcity is already reflected in price?

Historical infrastructure and commodity cycles warn against treating current scarcity as permanent. High returns encourage investment, and investment changes the constraint. This can create an uncomfortable feature of cyclical investing: the strongest current financial results may occur near the point when the future supply response is becoming most dangerous. Enterprise execution can remain excellent while the environmental source of excess profitability begins to disappear.

The same logic applies to monetary conditions. The earlier Comprehensive Investment Framework research distinguished the price of credit from the quantity of credit because a high policy rate does not necessarily mean that credit creation is genuinely restrictive. That finding remains useful, but it should enter an investment only through a transmission mechanism. A leveraged property developer, a bank, a consumer lender and an asset-light software business do

not experience the same monetary environment in the same way. The analytical task is to identify the relevant channel rather than attach one macroeconomic score to all four.

Geopolitics works similarly. A semiconductor company with concentrated production in a contested region may require detailed geopolitical analysis. A domestic service company with local customers and little imported input may not. The importance of the environment is determined by causal dependence, not by the prominence of the issue in the news.

From enterprise economics to the security

Even if we understand what an organisation can repeatedly accomplish, why competitors cannot readily eliminate the resulting economic surplus and how external conditions affect that system, we have still not established what the investor will receive. Value creation, enterprise capture and shareholder capture are separate processes.

A Chinese technology platform provides a clear example. Its operating businesses could possess enormous customer reach, data, network characteristics, technological capability and cash-generating potential. Yet foreign investors could hold securities whose economic relationship with important operating entities depended partly on contractual arrangements and a sovereign regime capable of changing the rules governing them. An investor could therefore be right about the business and wrong about the security.

The principle extends beyond China. A minority investor can own shares in an enterprise controlled by another shareholder whose interests differ materially from those of outside owners. A mining company can possess an exceptional resource while the fiscal regime changes. A utility can own infrastructure of obvious economic importance while a regulator determines the return available on the capital invested. A pharmaceutical company can develop a clinically valuable product while reimbursement arrangements influence how much of the clinical value becomes corporate revenue.

Claim Architecture is a disciplined description of the path between the economic activity and the investor. It includes legal ownership, contractual rights, voting arrangements, minority protections, jurisdiction, regulatory permissions and other conditions that determine how enterprise economics become security-holder economics. In many investments the path is simple and deserves little attention. In others it is central.

Enterprise governance and investor protection are related but different. Concentrated founder control may provide strategic coherence and preserve organisational purpose while simultaneously giving outside shareholders little influence if capital allocation deteriorates. The same governance arrangement can strengthen organisational capability and weaken the protection of the investor's claim.

The separation between value creation, enterprise capture and shareholder capture clarifies why socially valuable businesses are not automatically attractive investments. An enterprise can create enormous benefits for customers while competition transfers most of the value to them. It can retain substantial economics at the corporate level while regulation, taxation or contractual

structure limits what reaches outside shareholders. The investor must follow value through each stage rather than infer shareholder value from social usefulness or corporate growth.

Essential services make this tension particularly visible. Electricity, healthcare, education, banking and communications can enjoy durable demand because society cannot easily do without them. The same essentiality can attract regulation precisely because governments care about access, affordability and continuity. Demand durability may strengthen one part of the investment proposition while weakening the enterprise's freedom to capture value or the shareholder's certainty of receiving it.

Claim Architecture gives us a place to analyse this without turning every regulatory issue into a general risk factor. In a straightforward listed company, the claim may require little discussion. In a controlled company, partnership, concession, holding structure or variable-interest arrangement, the path from operating cash flow to the security can become central. The question is whether another actor has the legal or practical ability to redirect, restrict or subordinate the economics on which the valuation depends.

The Chinese technology example separates business analysis from security analysis cleanly. An investor could correctly forecast user growth, transaction activity, margins and technological development and still receive a poor outcome if the legal or regulatory relationship between the operating businesses and the listed claim changed. That is not a failure of the operating forecast. It is a failure to model the claim.

Similar problems appear in more familiar settings. A mining company may control a world-class deposit but operate under a fiscal regime capable of changing royalties or taxes. A regulated utility may own an irreplaceable network but earn only the return permitted by the regulator. A bank may possess strong customer relationships but be required to hold more capital, reducing the return on equity available to shareholders. A pharmaceutical company may create large clinical value while reimbursement rules determine the commercial value it can retain.

These examples also show why investor governance should not be collapsed into enterprise governance. A founder-controlled structure may be excellent for long-term organisational decision-making and poor for minority protection. The same evidence can support one proposition and weaken another. Recording both is more useful than forcing the evidence into a single governance score.

Valuation, uncertainty and portfolio construction

Valuation is where our explanation of the enterprise encounters the expectations of everybody else willing to trade the security. The market price contains assumptions, whether explicit or not, about growth, returns, durability, competitive conditions, capital allocation and risk. The task is not simply to estimate a different number. It is to understand which assumptions embedded in the price differ materially from what the evidence allows us to believe.

The central uncertainty in valuing an exceptional enterprise is often not next year's earnings. It is the duration over which unusual economics can persist and the amount of capital that can be deployed into them. Small changes in the assumed duration of exceptional economics can produce large changes in estimated value. Yet the further the valuation reaches into the future, the further it often reaches beyond demonstrated organisational capability.

There is an unavoidable tension. The greatest investment outcomes are produced precisely because an organisation accomplishes things that had not yet been demonstrated when the investment was made. An investor who refuses to believe anything until it has already happened will identify exceptional enterprises too late. An investor who assumes that every successful process can continue indefinitely will eventually capitalise an extrapolation as though it were evidence. The discipline is to recognise when we have crossed from evidence into expectation.

A useful downside case should therefore describe more than lower revenue growth. It should ask how the causal system could fail. The organisation may continue producing good products but lose the mechanism that allowed it to retain attractive economics. A critical external permission may change. Growth may exceed organisational capacity. An acquisition may damage the existing operating system. A technology may reduce the economic usefulness of a historically valuable capability. The enterprise may remain successful while the shareholder's claim becomes less secure.

Position size follows from the same reasoning. Conviction is not enough. Two propositions in which we have similar confidence may deserve very different allocations if one has a relatively narrow range of adverse outcomes while the other contains a small but credible possibility of permanent impairment. Claim structure, leverage, external dependencies and reversibility all affect the distribution of outcomes.

Portfolio construction then becomes the final stage of the causal analysis rather than a separate optimisation exercise. A portfolio is not diversified simply because it contains many company names. Several apparently unrelated companies may depend on the same interest-rate environment, government reimbursement system, semiconductor supply chain, Chinese demand cycle, electricity constraint or capital-market condition. The relevant concentration is not always visible in the company name or industry classification.

There is an unavoidable tension in valuing exceptional enterprises. The greatest investment outcomes are produced precisely because an organisation accomplishes things that had not yet been demonstrated when the investment was made. An investor who refuses to believe anything until it has already happened will identify many exceptional enterprises too late. An investor who assumes that every successful process can continue indefinitely will eventually capitalise extrapolation as though it were evidence.

The historical record helps define the boundary. We can observe what an organisation has repeatedly accomplished, the circumstances under which it accomplished it and the mechanisms that appear to have produced the result. Beyond that boundary lies inference. Some inference is unavoidable. The discipline is to recognise when we have crossed from evidence into expectation.

The home-improvement retailer had demonstrated that its warehouse model could be replicated across a growing number of stores. An investor still had to infer that the model could travel across a continent. The medical-technology company had demonstrated an ability to solve related clinical problems. An investor still had to infer that this knowledge could transfer into more complex product categories and acquired businesses. The software company had demonstrated successful transitions across several technologies. An investor still had to infer that the organisation could continue adapting through technological changes it had never encountered. Valuation inevitably prices those inferences.

This is one reason apparently expensive businesses can sometimes produce extraordinary investment returns while apparently cheap businesses disappoint. The expensive security may still underestimate the duration or breadth of an organisation's ability to create and capture value. The cheap security may accurately reflect deterioration in the mechanism that produced its historical economics. A valuation multiple alone cannot distinguish the two situations.

Scenario analysis becomes more useful when scenarios represent different states of the causal system rather than mechanically different growth rates. One scenario may assume Capability continues strengthening and Strategic Power persists. Another may assume the organisation remains capable but industry scarcity disappears. A third may assume customer economics remain attractive while regulation reduces shareholder capture. A genuine failure case may involve the breakdown of a critical dependency rather than merely a lower terminal growth rate.

The structure of uncertainty also matters for position size. Two investments with similar expected returns can deserve different allocations if one has a relatively narrow range of adverse outcomes while the other contains a small but credible possibility of severe permanent impairment. Leverage, Claim Architecture, external dependencies and reversibility all influence the distribution.

Portfolio construction follows from the same reasoning. A portfolio is not diversified simply because it contains many company names or industry classifications. A generator, transformer manufacturer, data-centre developer and artificial-intelligence infrastructure supplier can occupy different industries while depending on the same capital-expenditure cycle. Several healthcare businesses can depend on one reimbursement regime. Several companies headquartered in different countries can depend economically on the same source of demand. The relevant concentration may lie in the causal dependency rather than the sector label.

This view also gives short investing a natural place in the same intellectual system. A separate theory of short selling is unnecessary. The long investor often looks for a situation in which the market underestimates the persistence, transferability or strengthening of a causal system. The short investor can look for the reverse: a market continuing to extrapolate historical outputs after the system producing them has begun to deteriorate. A consumer brand losing preference while consolidated earnings remain supported, an acquisition-led business maintaining headline growth while integration quality deteriorates, or a regulated enterprise approaching the loss of a critical permission are different versions of the same analytical problem.

Monitoring and falsification

The natural unit of monitoring is the investment proposition rather than the company. If an investment is a set of causal propositions, each important proposition should carry its own expected evidence. A retailer believed to possess a scale-driven customer-value mechanism should eventually show evidence connecting increasing scale with purchasing economics and customer value. A software business whose advantage depends on architectural reuse should show evidence that development expenditure continues producing reusable functionality without rapidly increasing complexity. An insurer whose capability lies in risk segmentation should preserve underwriting discipline even when growth becomes harder. A serial acquirer should demonstrate acceptable returns across acquisition cohorts rather than simply continuing to announce transactions.

The monitoring frequency follows the mechanism. Some evidence changes slowly enough that annual review is sufficient. Some can be observed quarterly. Regulatory permission may require event-driven monitoring. Customer behaviour can sometimes be tracked continuously. Capital allocation may need several years before the outcome becomes clear.

The purpose is not to react to every change. It is to distinguish variation in outcomes from deterioration in causes. Weak financial performance can occur while the causal system strengthens. Strong performance can continue while an important mechanism deteriorates. A capability can be impaired without being destroyed. Strategic Power can weaken while organisational capability remains. A regulatory dependency can fail while both Capability and Strategic Power remain otherwise intact.

Falsification provides the discipline that prevents monitoring from becoming narrative maintenance. Before investing, we should identify what evidence would cause us to reduce confidence in each important proposition. The threshold need not always be numerical. The important point is temporal: the conditions for changing our mind should be considered while we are still capable of imagining that we might be wrong.

Long-horizon patience is therefore conditional. Patience is justified with variance while the causal system remains credible; it becomes stubbornness when evidence says the mechanism has deteriorated and the investor responds by changing the explanation rather than changing the conclusion. Historical quality deserves respect because it contains evidence. It does not deserve immunity from new evidence.

Monitoring should begin when the investment is formulated, not after the first disappointing result. If an investment is a set of causal propositions, each important proposition should carry its own expected evidence. A retailer believed to possess a scale-driven customer-value mechanism should eventually show evidence connecting increasing scale with purchasing economics and customer value. A software business whose advantage depends on architectural reuse should show that development expenditure continues producing reusable functionality without rapidly increasing complexity. An insurer whose capability lies in risk segmentation should preserve underwriting discipline even when growth becomes harder. A serial acquirer should demonstrate acceptable outcomes across acquisition cohorts rather than simply continue announcing transactions.

The observation frequency follows the mechanism. Customer retention may be assessed annually. Underwriting can often be assessed quarterly. Acquisition integration may require several years of cohort evidence. A regulatory permission can change in one event. Brand relevance can erode slowly. A monitoring system that forces all of these into the same quarterly timetable creates an appearance of discipline while ignoring the thing being measured.

The industrial-distribution case shows why the distinction between outcome and mechanism matters. During an industrial downturn, earnings weakened while management continued investing in locations, distribution capacity and internally developed logistics. A purely financial monitor could interpret the lower earnings as deterioration. A mechanism-based monitor could observe that the customer-proximity system was still being strengthened. Later changes in the physical form of that system—from branches toward vending and embedded customer locations—could then be understood as adaptation of the means rather than abandonment of the underlying capability.

The consumer-brand case shows the reverse. Consolidated earnings could remain acceptable while a historically important brand lost share over several years. Pricing, other brands and geographic diversification could obscure the erosion. Monitoring the proposition at brand level would have produced a different signal from monitoring the enterprise only through aggregate financial statements.

Falsification prevents monitoring from becoming narrative maintenance. Before investing, the research should identify what evidence would cause confidence in each important proposition to fall. The threshold need not always be numerical. A consumer franchise might require several years of category-adjusted share loss before the evidence becomes compelling. A licence can fail in one event. A product-development capability may weaken gradually as development costs rise, release delays increase and customer adoption falls.

The timing matters because interpretation becomes harder after capital has been committed. Evidence that once would have been regarded as contradictory can be redescribed as temporary. A broken mechanism can become a long-term opportunity. Management explanations receive more weight because accepting them allows the original thesis to survive. Pre-specifying the important propositions does not eliminate this behaviour, but it makes it more visible.

Long-horizon patience should therefore be conditional on the causal system. Weak performance does not itself justify exit when the mechanism remains intact or strengthens. Strong performance does not justify continued ownership when the mechanism is deteriorating. Historical quality deserves respect because it contains evidence; it does not deserve immunity from new evidence.

Towards the Celerity Investment Framework

We began this research by trying to ensure that investment analysis did not omit important subjects. The research has led somewhere more demanding. The problem is not simply to examine more things; it is to understand how one thing produces another.

Capability Science provides a way of thinking about what the organisation has demonstrated that it can repeatedly accomplish. Strategic Power provides a way of examining why valuable outcomes can produce persistent enterprise economics. Industry and environmental analysis explain the conditions within which those mechanisms operate. Claim Architecture follows the economics from enterprise to investor. Capital allocation determines how today's economic surplus changes tomorrow's enterprise. Valuation compares the resulting future with the expectations embedded in price. Falsification and monitoring keep the explanation open to revision. Portfolio construction recognises that individually attractive propositions can still share the same underlying risk.

The sequence is causal rather than categorical. That is a substantial departure from where the work began.

There are obvious dangers in this approach. Causal stories can become more seductive than simple financial narratives because they appear to explain everything. Organisational evidence is incomplete. Historical cases invite hindsight. Relationships that appear reinforcing may simply reflect a favourable industry environment. Capabilities inferred from successful outcomes may be partly luck. Analysts can mistake elegant explanations for true ones.

The research design therefore matters as much as the theory. The HCA programme was deliberately constructed around point-in-time evidence and credible contemporaneous comparisons because retrospective success stories are too easy to write. The first ten assessments have already produced a useful warning: the original capability scores often discriminate less than the narrative architecture surrounding them. The proposition that Capability Architecture distinguishes exceptional long-duration enterprises remains something to test, not something the eventual framework should assume merely because the idea is attractive.

We can nevertheless say several things with reasonable confidence. Historical performance contains information, but it does not explain itself. Organisational capability requires evidence beyond corporate intention. The conditions under which capability has been demonstrated matter when we extrapolate it. Growth can strengthen an organisation or overwhelm it. Capital allocation has to be evaluated through realised outcomes rather than strategic rhetoric. Creating value and retaining value are different economic problems. Enterprise economics and shareholder economics can diverge. External conditions matter when they enter the causal system. Price can transform an exceptional enterprise into an unattractive investment. Monitoring is more useful when it tests the propositions on which the investment depends.

The proposition that exceptional long-term enterprises can be distinguished by a particularly reinforcing Capability Architecture remains promising but unvalidated. The pilot gives us evidence worth pursuing, not permission to declare the question settled.

There is also a practical consequence for Celerity as an investment research business. If this approach survives empirical testing, the research process does not end when a stock enters the portfolio. The original causal assessment becomes the basis of monitoring. The evidence collected before purchase establishes the propositions; subsequent company, industry and environmental evidence updates them. Selection and monitoring become one continuous process.

A causal model does not guarantee the correct answer. It gives us a disciplined way of deciding what a change means. Is it variance around an intact system, evidence of a weakening capability, loss of Strategic Power, a changed external dependency, poor capital allocation, a less secure shareholder claim, or simply a valuation adjustment unrelated to the enterprise?

There is no reason to expect the finished Celerity Investment Framework to make investing easy. Exceptional investment outcomes emerge from organisations operating within complex adaptive systems, and the evidence available to investors is incomplete. Simplification is necessary, but simplification should occur after we understand what is being simplified.

The objective is modest in wording and demanding in practice: understand enough of the causal system to know why the historical economics existed, what has to remain true for them to continue, what the security holder is actually entitled to receive, what expectations are already contained in the price and what evidence would tell us that our explanation is wrong.

That is where the research has brought us. The next task is to determine what investment architecture follows from it. That will be the Celerity Investment Framework.

The research has therefore moved beyond the question of how to assemble a more comprehensive checklist. The unpublished Comprehensive Investment Framework was valuable as a research scaffold because it forced attention onto subjects that narrower approaches can neglect. The work that followed suggests that the eventual Celerity Investment Framework should not simply preserve those subjects as adjacent modules. The more coherent organising principle is causality.

Repeated performance provides observations. Capability seeks the organisational process capable of explaining those observations. Capability Architecture examines how those processes interact, transfer and change as the enterprise develops. Strategic Power addresses why the enterprise can retain an attractive share of the value created. Industry structure, scarcity and environmental dependencies condition the economics of the system. Claim Architecture follows the value from the operating enterprise to the security holder. Capital allocation changes the future system. Valuation compares the probability-weighted future available to the security with the expectations embedded in price. Monitoring and falsification keep the explanation open to revision. Portfolio construction recognises that different securities can share the same causal dependencies.

This sequence should not yet be mistaken for the finished framework. The Historical Company Assessment pilot is small. Its most distinctive proposition—that exceptional long-duration performers may be distinguished more by reinforcing capability architecture than by the absolute strength of individual capabilities—remains unvalidated. The full historical study must be capable of weakening or rejecting it. Other findings may also change as the evidence base expands.

The discipline of leaving the architecture unfinished is useful. It prevents an attractive theory from becoming the lens through which all later evidence is interpreted. The Research Findings Map records what has survived the work so far; it does not grant those findings permanent status. The empirical programme should be able to expose missing variables, contradictory cases and conditions under which the propositions fail.

Several conclusions are nevertheless strong enough to guide the next stage. Historical performance contains information but does not explain itself. Capability requires demonstration rather than corporate intention. The conditions under which capability has been demonstrated matter when it is extrapolated. Growth can strengthen an organisation or overwhelm it. Capital allocation must be judged through realised outcomes. Creating value, retaining value at the enterprise and delivering value to the shareholder are different processes. External conditions matter when they enter the causal system. Price can turn an exceptional enterprise into an unattractive security. Monitoring is more useful when it tests the propositions on which the investment depends.

For an investment research business, this has a practical consequence. Selection and monitoring become one continuous process. The original research identifies the causal propositions, the evidence supporting them and the dependencies on which they rest. Subsequent company, industry and environmental evidence updates those same propositions. The research does not end when a security enters the portfolio; ownership creates a continuing empirical test of the original explanation.

Markets reveal changes in reported earnings quickly. The more difficult task is deciding what the change means. It may be variance around an intact system, evidence of a weakening capability, erosion of Strategic Power, failure of an external dependency, poor capital allocation, a less secure shareholder claim or a valuation adjustment unrelated to the enterprise. A causal model does not guarantee the correct diagnosis, but it gives the investor a disciplined way to distinguish among them.

The objective is modest in wording and demanding in practice: understand enough of the causal system to know why the historical economics existed, what has to remain true for them to continue, what the security holder is actually entitled to receive, what expectations are already contained in the price and what evidence would tell us that our explanation is wrong.

That is where the research has brought us. The next task is to determine what investment architecture follows from it. That will be the Celerity Investment Framework.

Research basis

This essay draws on the unpublished Comprehensive Investment Framework research; Capability Science; the Historical Company Assessment pilot; Long-Horizon Compounder Research; Strategic Power and industry-structure research; the electricity, cognition and infrastructure constraint programme; valuation and decision-risk work; and the longitudinal company cases used to test temporary weakness, capability impairment, Power erosion, external dependency, acquisition integration and shareholder claim. Company names have deliberately been omitted from the examples so that the argument rests on the mechanism rather than on the familiarity or reputation of the enterprise.