CELERITY RESEARCH

Towards the Celerity Investment Framework

Why repeated performance is evidence, not explanation

LinkedIn Article Edition | August 2026

Investing appears straightforward until one tries to explain why a successful business should remain successful. Find a good company, understand what makes it good, avoid paying too much and allow time to do the work. There is a great deal of truth in that approach. The difficulty is that descriptions such as quality, management, competitive advantage and valuation often 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 excellent management without operating in an industry that allows the value it creates to remain with the enterprise. It can retain formidable operating capability while an important source of Strategic Power weakens. It can remain an exceptional enterprise while becoming a poor investment because the price already assumes an even more exceptional future.

Our earlier work on the Comprehensive Investment Framework (CIF) tried to prevent these omissions by bringing economic conditions, industry structure, Strategic Power, organisational quality, management, innovation, finance, valuation and portfolio construction into one research programme. The CIF was never published, and the research that followed has convinced us that its architecture should not simply be refined. Its value lies in the research it generated.

Capability Science changed the direction of that work by asking a more fundamental question: what has the organisation demonstrated that it can repeatedly accomplish?

Financial statements record outcomes. They do not explain their organisational cause. Revenue growth tells us that customers bought more, but not why the organisation repeatedly persuaded them to do so. A high return on capital tells us that an attractive economic relationship existed between investment and profit, but not whether it resulted from organisational capability, temporary scarcity, favourable regulation, industry structure or luck.

This has changed the way we think about historical performance. Long-Horizon Compounder Research is right to take persistence seriously. A company that has increased revenue, earnings and cash flow while sustaining attractive returns for ten or twenty years deserves attention. The mistake is to allow the record to become its own explanation.

Repeated performance is better treated as evidence. Something has repeatedly produced these unusual outcomes. The investment problem is to understand what.

What the matched-company research changed

The first ten Historical Company Assessments (HCAs) have made this problem more difficult, which is exactly what a useful research programme should do. We assessed eventual exceptional performers against credible companies that looked attractive at the same historical point in time, using only evidence that would then have been available.

The results were not clean. Strong eventual performers and their contemporaneous comparisons often both scored well across individual capability domains. In an American home-improvement comparison, the older retailer had more stores, greater revenue, a longer operating history and mature standardised systems. An investor at the time could reasonably have regarded it as the safer enterprise. The younger retailer did not obviously possess a collection of capabilities that its established competitor lacked.

The possible distinction lay in how 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 transferred to subsequent stores. Capital allocation funded expansion of essentially the same system. Growth was not merely increasing the size of the enterprise; it appeared capable of making the system itself stronger.

A similar problem appeared when we compared two American membership warehouse retailers. Both had strong customer propositions, disciplined operations, purchasing capability, membership economics and coherent capital allocation. On a simple capability score they were almost indistinguishable. The provisional difference was the extent to which customer trust, employee capability, purchasing, capital allocation and organisational learning appeared to feed one another.

This remains a research hypothesis, not a conclusion. The pilot is too small to establish that Capability Architecture distinguishes exceptional long-duration enterprises. The full historical study must be capable of rejecting the idea. But the pilot has already shown why a single company-quality score is unlikely to be enough.

Capability has demonstrated limits

The research has also forced us to be more precise about the word capability. An activity is not a capability. Opening a store is an activity. Acquiring a company is an activity. Launching a product is an activity. Increasing research expenditure is an activity. Even possessing a formal process tells us less than corporate descriptions often imply. The stronger evidence appears when a process repeatedly produces the intended difficult outcome.

Evidence becomes more useful when conditions change. An Australian enterprise-software company, for example, repeatedly rebuilt its product through major changes in computing architecture while preserving a common product architecture and continuing to serve specialised vertical markets. Client-server computing, internet delivery, cloud infrastructure and software-as-a-service changed the means of delivery. The underlying organisational logic remained recognisable. The organisation did not simply repeat an outcome; it reproduced an outcome while the problem changed.

Historical evidence also tells us where confidence should end. An organisation that has successfully opened fifty stores within neighbouring regions has demonstrated something. It 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 demonstrated something. It has not necessarily demonstrated that it can absorb a transformational merger involving different systems, cultures and technologies.

This is where the conclusion that capability has demonstrated limits becomes financially important. The capability may be genuine while the proposed application lies beyond the conditions in which it has been demonstrated.

Our acquisition work produced a useful contrast. A North American logistics-software business has made focused acquisitions part of its normal operating model. Transactions add customers, functionality, data, geographic reach or connections to a wider network, while recurring revenue, margins and cash conversion provide evidence about whether the enlarged organisation remains healthy.

A large combination in the dental-products industry had an equally persuasive strategic rationale. The businesses served similar customers and offered complementary products. Yet substantial impairments, leadership change and repeated restructuring followed. Strategic adjacency was observable at the outset. Demonstrated integration capability was not.

Capital allocation should therefore be judged by realised causal outcomes. Did the investment deepen customer knowledge? Did it improve the operating system? Did acquired capability transfer into the wider organisation? Did it create or renew Strategic Power? Did it increase per-share economic value after accounting for the capital used?

Capability is not Strategic Power

Capability explains what an organisation can repeatedly accomplish. It does not explain who keeps the value created by that accomplishment.

A well-run airline can operate extremely efficiently 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.

Strategic Power addresses the mechanism that allows the enterprise to retain an attractive share of the value it creates. That mechanism may arise from customer switching, scale, network structure, scarce resources, brand preference, process replication difficulty or another structural relationship. Power is relational: a company can possess substantial Power over customers and little over suppliers, or considerable Power relative to competitors and almost none relative to government.

This distinction helps explain situations that otherwise look contradictory. A global consumer-products company can retain formidable manufacturing, procurement and distribution capability while an important brand gradually loses consumer preference. Consolidated results can remain respectable for some time even as one source of Strategic Power deteriorates. Conversely, an enterprise can retain a valuable licence or network position for a period despite mediocre internal capability.

The financial numbers do not resolve these questions by themselves. They become more useful when connected to a causal proposition. If management claims that scale is improving purchasing economics, the operating evidence should eventually reflect it. If an acquisition is supposed to create cross-selling, customer and revenue data should eventually show it. If a brand possesses pricing power, repeated price increases accompanied by persistent volume or share deterioration should not automatically be treated as confirmation.

The environment belongs in the causal chain

The same reasoning changes the role of macroeconomics, regulation and industry structure. They matter when they enter the causal system, not because every investment requires the same top-down checklist.

For an American education provider whose students depended heavily on federal funding, government eligibility was part of the revenue mechanism. The company could retain capable staff, recognised programmes and operating infrastructure and still face an existential problem if the external permission enabling customers to pay disappeared.

Our electricity and artificial-intelligence research produced another version of the problem. Rapid growth in computation creates demand for electricity, but it does not tell us where durable economic returns will appear. Scarcity can arise in generation, transmission, transformers, grid connection, firm capacity, powered land, cooling, accelerators or elsewhere. High returns attract capital. Engineering creates substitutes. Governments intervene. A bottleneck that appears decisive can eventually become ordinary capacity while the constraint moves somewhere else.

Scarcity is therefore not the same as Strategic Power. A company operating at full capacity during a shortage may report extraordinary margins without possessing a durable structural advantage. The investor has to understand what creates the scarcity, how quickly supply can respond, who controls the constraint and whether the owner can retain the economics before the bottleneck moves.

The enterprise is not always the security

Even after we understand Capability, Strategic Power and the external environment, we have not established what the investor will receive. Value creation, enterprise capture and shareholder capture are separate processes.

A Chinese technology platform makes the distinction unusually clear. Its operating businesses can possess extraordinary customer reach, data, network characteristics and technological capability while foreign investors hold securities whose economic relationship with important operating entities depends partly on contractual arrangements and a sovereign legal regime. An investor can therefore be right about the business and wrong about the security.

The principle is broader than one country. A utility can own economically essential infrastructure while a regulator determines the permitted return. A mining company can possess an exceptional resource while its fiscal regime changes. A pharmaceutical company can create substantial clinical value while reimbursement arrangements influence how much becomes corporate revenue. A minority shareholder can own part of an excellent enterprise while control arrangements weaken the protection of that shareholder's economic claim.

We use Claim Architecture to describe this path from operating activity to investor economics: what is legally owned, how value moves through the structure, which rights protect the investor and who has the ability to interrupt or redirect the flow.

This also means that enterprise growth and shareholder compounding are not identical. Management can create a larger company while issuing enough equity that per-share value grows much more slowly. An acquisition can increase earnings while earning an inadequate return on the capital used. Debt can improve per-share returns while increasing the probability of permanent impairment.

Price is part of the causal argument

Valuation is where our explanation of the enterprise encounters the expectations embedded in the market price. The task is not simply to calculate a different number. It is to understand which assumptions contained in the price differ materially from what the evidence allows us to believe.

The central uncertainty in valuing an exceptional enterprise is often duration. A company capable of reinvesting at high incremental returns for twenty years is worth dramatically more than an otherwise similar company whose opportunity is exhausted after five. Yet the further the valuation reaches into the future, the further it can reach beyond demonstrated organisational capability.

There is an unavoidable tension. The greatest investment outcomes occur because organisations accomplish things that had not yet been demonstrated when the investment was made. Refusing to believe anything until it has already happened identifies exceptional enterprises too late. Assuming that every successful process can continue indefinitely capitalises extrapolation as though it were evidence.

Historical evidence helps define the boundary. It tells us what the organisation has repeatedly accomplished, under which conditions, and which mechanisms appear to have produced the result. Beyond that lies inference. Investing requires inference; discipline requires recognising when we have crossed from evidence into expectation.

A wonderful enterprise can therefore be a poor investment without anything going wrong operationally. The market may simply require an even more wonderful future. The reverse can also occur: temporarily weak earnings can coexist with a strengthening causal system, leaving the market to price the output while underestimating the mechanism.

Scenario analysis is more useful when it describes alternative states of that system rather than merely changing revenue growth and terminal multiples. Capability may remain strong while industry scarcity disappears. Customer economics may remain attractive while regulation reduces shareholder capture. Growth may overwhelm organisational capacity. A critical permission may fail. These are economically different failure states.

Monitoring should test propositions

This brings selection and monitoring together. An investment is not really one thesis. It is a collection of causal propositions.

If a retailer is believed to possess a scale-driven customer-value mechanism, increasing scale should eventually connect to purchasing economics and customer value. If a software business depends on architectural reuse, development expenditure should continue producing reusable functionality without rapidly increasing complexity. If an insurer's capability lies in risk segmentation, underwriting discipline should survive periods in which growth becomes harder. If a serial acquirer is genuinely capable, acceptable economic outcomes should be visible across acquisition cohorts rather than merely in the continuing announcement of transactions.

These propositions operate at different speeds. Customer retention may be observed annually. Insurance pricing quality may be visible quarterly. Acquisition integration can require several years. A regulatory permission can change abruptly. Brand relevance may erode slowly. Monitoring should follow the natural cadence of the mechanism rather than force every proposition into the same quarterly dashboard.

Falsification then becomes part of the investment before capital is committed. The research should identify what evidence would reduce confidence in each important proposition. This need not mean artificial numerical thresholds for everything. The discipline lies in deciding while we are still capable of imagining that we may be wrong.

That changes the meaning of patience. Long-horizon investing is often described as the ability to do nothing while others react to short-term noise. A more precise discipline is to remain patient with variance while refusing to remain patient with causal deterioration. Weak reported performance can coexist with an intact or strengthening mechanism. Strong reported performance can continue while an important mechanism deteriorates.

Historical quality deserves respect because it contains evidence. It does not deserve immunity from new evidence.

Towards the Celerity Investment Framework

The work has not yet produced the new Celerity Investment Framework. It has produced something that should come first: a clearer account of what an investment framework needs to explain.

Historical performance contains information, but it does not explain itself. Capability requires demonstration rather than corporate intention. The conditions under which capability has been demonstrated matter when it is extrapolated. Capabilities may interact in ways that strengthen or constrain the enterprise, although our proposition that reinforcing Capability Architecture helps distinguish exceptional long-duration enterprises remains unvalidated. Growth can strengthen an organisation or overwhelm it. Capital allocation has to be judged through realised outcomes. Capability and Strategic Power are different. 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 transform an exceptional enterprise into an unattractive investment. Monitoring is more useful when it tests the propositions on which the investment depends.

The emerging sequence is causal rather than categorical. Repeated performance provides observations. Capability seeks the organisational processes capable of explaining those observations. Capability Architecture examines how those processes interact 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. Claim Architecture follows value from the operating enterprise to the security holder. Capital allocation changes the future system. Valuation compares that future 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 underlying causal dependencies.

This sequence should not yet be mistaken for the finished framework. The HCA pilot is small, and its most distinctive proposition must survive a much larger historical validation programme. Leaving the architecture unfinished prevents an attractive theory from becoming the lens through which all later evidence is interpreted.

For Celerity Research, there is also a practical consequence. Selection and monitoring become one continuous research process. The original assessment 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. Ownership becomes a continuing empirical test of the original explanation.

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

The next stage is to determine what investment architecture follows from those findings. That is the work now underway at Celerity Research.