

LANTERN LABS

INVESTMENT RESEARCH

www.lanternlabs.net

The Reverse Information Paradox

A Lantern Labs Strategy Paper

Why the Firms That Consume Intelligence Must Own the Intelligence They Create
July 2026

Ownership of the Learning Loop as the Competitive Asset of the AI Era

A Lantern Labs strategy paper — original first-principles analysis. Not investment advice.

Decision Intelligence, not Decision Making — a strategy paper, not investment advice.

Prepared by Lantern Labs | For public distribution | Not investment advice

Contents

Executive Brief	2
I. Arrow's Paradox, and Its Inversion	3
II. Why Investing Is the Acute Case	3
III. The Asset That Is Actually Being Contested	4
IV. What a Trust Boundary Actually Is	5
V. A Description of the Right Architecture, Not a Wish List	6
VI. The Trap — and the Commitment That Answers It	6
VII. The Universal Form — Whoever Owns the Learning Loop	7
VIII. Anticipated Objection — The Data Network Effect	8
IX. What to Do About It	8

% Document-local: sections flow to fill the page instead of hard page-breaking, % breaking only when they would otherwise start in the bottom ~14 lines.

Executive Brief

The conclusion, the reasoning, and the one thing to act on.

CORE THESIS

Kenneth Arrow described a paradox in the market for information: a buyer cannot value information until they have it, at which point they have acquired it for free — so the *seller* is the one exposed. Artificial intelligence inverts the exposure. When you buy intelligence from a model, the risk shifts to *you*. To make the model useful you must feed it your proprietary knowledge, and the better you want it to perform, the more of that knowledge you must reveal. You pay twice: once in money, and again in the private judgment you disclose to make the purchase work. In no industry does this matter more than investing — because in investing, private judgment is not a byproduct of the product. It *is* the product.

THE SHARPEST MECHANISM

An investment firm's edge is, by definition, a piece of information the market does not yet share: a variant perception, a weighting the consensus has wrong, a falsifier the crowd is ignoring. The moment that variant perception passes across a boundary you do not control — into a model provider's training set, its usage logs, its correction data — it begins its journey toward consensus. You are not merely leaking data. You are funding the erosion of the one thing that lets you outperform. And the leak is nearly invisible: it happens trace by trace, correction by correction, eval by eval, never in a single disclosable event.

THE DEEPER CLAIM

The reverse information paradox is the symptom. The underlying shift is larger: the fundamental unit of competitive advantage in the intelligence age is not data and not models, but **compounding organizational judgment** — and the AI era will be defined by *who owns the learning loop* through which that judgment compounds. Data can be bought; models can be rented; judgment can only be accumulated, and only by whoever controls the corrections, evals, and memory that produce it.

THE ONE THING TO ACT ON

Stop protecting your *data* and start protecting the *mechanism through which you learn*. Data — prices, filings, transcripts — is commodity; the whole market has it. Your learning loop — what you believed, why you differed, what you got wrong, and how the next call improved because of the last — is not. Build a trust boundary around that loop, make it model-agnostic, and ensure the intelligence you create while consuming intelligence stays yours.

I. Arrow's Paradox, and Its Inversion

In 1962 Kenneth Arrow set out a structural problem in the economics of information. Its value to a buyer cannot be known until the buyer has seen it — but once seen, it has in effect been handed over for nothing. The seller who wants to sell knowledge risks giving it away in the act of demonstrating it. Patents exist, in part, to resolve one face of this problem: they let an inventor disclose an idea and still be paid for it, converting an un-sellable secret into a licensable asset.

The AI age produces the mirror image. The exposed party is no longer the seller of information but the *buyer* of intelligence. A model arrives capable but generic. To make it perform on your problem, you supply context: your documents, your workflows, your definitions of a good answer, and — most valuably — your corrections when it is wrong. Each of those is a piece of your private knowledge. The transaction that was supposed to deliver intelligence *to* you also extracts intelligence *from* you, and the extraction compounds: the provider learns more about how you think every time you use what you bought, while you learn almost nothing about what it is learning in return.

This is the reverse information paradox. You set out to acquire an edge and end up donating one — and the more capable you want the intelligence to be, the more of your edge you must hand across the boundary to get it. Arrow needed the patent to protect the seller. The buyer of intelligence has no equivalent instrument yet, and needs one just as badly.

II. Why Investing Is the Acute Case

Every industry that adopts AI faces some version of this. Investing faces the sharpest version, for a reason specific to the business.

In most fields, the knowledge you feed a model is *operational* — how your support team handles a refund, how your code is structured, how your contracts are worded. Valuable, but not the whole franchise. In investment management, the knowledge you feed a model is the franchise itself. Alpha is *private information about value that the market has not yet priced*. Your thesis, your variant perception, the specific variable you think the consensus has mis-weighted, the precise condition under which you would change your mind — these are not inputs to your product. They are the product. The instant they cross an uncontrolled boundary, they start to decay, because information that is no longer private is no longer alpha.

Consider what actually leaks. Not the data — a competitor can buy the same filings and the same terminal. What leaks is the *decision*: the reason you differ from the market, the weighting scheme behind your conviction, and above all your *corrections* — the moments you told the machine “no, that’s wrong, here’s why.” A correction is the single richest artifact a judgment business produces, because it encodes not just what you know but *how you decide something is false*. It is the distilled residue of experience — expertise, standards, priorities, and decision process compressed into a single act. It is exactly what a competitor could never buy and exactly what escapes most imperceptibly.

There is a further irony worth naming plainly, because it is the part of the argument that survives every objection. Model providers rightly claim fair-use rights to train on the world’s public data — a genuinely valuable freedom that makes capable models possible. But the prevailing posture is then to reverse the direction of that generosity: to restrict customers from distilling or training on their *own* usage, and to reserve the right to learn from customer interaction. Learning is permitted to flow into the model from everyone, and permitted to flow out to no one. When learning runs in only one direction, economic value converges toward the owners of the learning infrastructure rather than toward the creators of the knowledge that made the intelligence useful. For an investment firm, “the creators of the knowledge” is you, and the knowledge is your edge.

III. The Asset That Is Actually Being Contested

Step back from investing for a moment, because the paradox is really a fight over which asset a firm gets to keep. Each technological age has been defined by the accumulation of a different unit of value, and the frontier has kept moving up a hierarchy:

Insight

Data → Knowledge → Learning → Judgment → Reputation

Each age has monetized a different rung, and the scarce asset has climbed steadily up the ladder:

Era	The scarce asset
Industrial	Physical capital
Information	Software
Cloud	Data
Intelligence	Learning
— <i>the competitive outcome it buys</i> —	<i>Judgment</i>

The Industrial Age accumulated physical capital; the Information Age accumulated software; the Cloud Age accumulated *data*, and taught firms to guard it. The AI age moves the contested asset up two rungs. Data is now commodity — abundant, purchasable, and increasingly used to train the very models everyone shares. *Knowledge* — structured, organized data — is only slightly scarcer. What is genuinely scarce, and genuinely compounding, is *learning*: the loop by which an organization converts its own experience into improved future performance. And the output of that loop, the thing firms have always actually competed on, is *judgment* — the ability to make a better call than a rival facing the same facts.

Above judgment sits *reputation*, and it is worth pausing on, because it completes the hierarchy rather than decorating it. Reputation is not a separate asset bolted on at the top; it is compounded judgment observed from the outside. It is the market's belief that your judgment reliably compounds — that the next call will be good because the last several were. A firm cannot manufacture reputation directly. It can only accumulate the judgment that, witnessed repeatedly over time, hardens into it.

This is why the hierarchy is also a defensibility ladder. Each rung is harder to copy than the one below, because each sits further from the raw input and closer to a track record. Data can be bought outright; knowledge can be organized in months; learning takes a full cycle of real feedback to build; judgment takes many cycles; and reputation takes all of that plus time and the market's sustained attention. The higher you compound, the less a competitor — or a foundation model starting cold — can shortcut the climb. That is the deepest reason to keep the learning loop inside your own boundary: it is the only rung from which the two most defensible assets, judgment and reputation, can be grown at all.

The strategic mistake of the moment is to fight the last war — to protect data behind the boundary while letting learning and judgment flow out of it. Data behind a firewall is the Cloud-Age reflex. In the AI age the asset that must sit behind the boundary is the *learning loop*, because that is where judgment is manufactured, and judgment is where firms actually compete. A company that owns its learning loop can afford to lose its data advantage; a company that owns only its data will watch its judgment quietly relocate into someone else's model.

IV. What a Trust Boundary Actually Is

The answer is not stronger data protection. Data protection guards a store of files. The reverse information paradox is not about a store of files; it is about a *flow* — the continuous exhaust of prompts, tool calls, evals, and corrections through which an organization’s judgment is externalized and, under the current regime, captured.

What a firm needs is a **trust boundary**, and it deserves a precise definition:

Insight

A trust boundary is not simply where information is stored. It is where organizational judgment compounds without involuntary leakage.

Inside that perimeter, a firm’s data, its traces, its evaluations, its adapted models, and its decision memory accumulate and improve *together* — and across it nothing crosses, not even the intelligence exhaust, without consent. In the cloud era, firms learned to accumulate data behind such a boundary. In the AI era, the thing that must sit behind the boundary is not the data but the *learning*: the mechanism by which the firm gets smarter, adapts, and compounds judgment over time.

Five properties make a boundary real rather than rhetorical. A firm can test itself against each.

Control. You own your evals, because evals define what “good” looks like *inside your firm* — what makes a research note sharp, what makes a valuation trustworthy, what makes a thesis falsifiable. You retain ownership of your memory, your traces, your feedback, your decisions, and your institutional context, and you keep the right to use the outputs of your own tasks and queries. If a provider’s terms forbid you from learning from your own work, you do not have control; you have a rental.

Capability. You can operate proprietary learning environments *inside* the boundary — systems where models are evaluated, supplied with institutional memory, and improved through structured feedback from real workflows without the firm’s decision knowledge leaving the perimeter. This does not require owning or fine-tuning the underlying model. It requires owning the evals, corrections, traces, context, and orchestration through which the model becomes useful to the organization. The model may remain general; the accumulated capability remains specific to the firm.

Choice. Your orchestration layer is decoupled from any single model. The test is concrete: *if any one model you use today were taken away, could you still operate and still optimize for your evals using another?* Does your firm’s accumulated “veteran” capability stay with you even when a given “generalist” model departs? If the answer is no, your edge lives in someone else’s data center.

Cost. A decoupled orchestration layer also lets you route each task to the most efficient combination of context and model without sacrificing quality — you are no longer paying a premium to a single provider for the privilege of teaching it your business.

Compound. Bring the four together and you have a **hill-climbing machine**: a loop in which every interaction improves the next one, producing a compounding optimization process rather than isolated automation. Each engagement makes the next one sharper, and the improvement is yours. This is the property that turns AI spending from an operating cost into an accumulating asset. It is also the property that is impossible to fake: if you cannot show that this quarter’s judgment is better *because of* last quarter’s, you have an archive, not a moat.

V. A Description of the Right Architecture, Not a Wish List

None of the above is speculative. It is a description of how a judgment business should be built when its raw material is proprietary decision data.

At Lantern the governing metaphor is deliberately narrow: *SemiAnalysis tears down chips; Lantern tears down decisions*. SemiAnalysis's moat is proprietary primary data about the physical world — a teardown lab a model cannot prompt its way into. Lantern collects a different primary data: the anatomy of a decision. What did we believe, what did the market believe, why did we differ, what was the key variable, what would have proven us wrong, how good was our evidence, what did we decide, what actually happened, and — the field most firms never record — what *type* of error did we make when we were wrong. Those structured, timestamped fields are retrieved into every future call, so that later calls demonstrably improve because of earlier ones. The model writes the prose; the decision data compounds.

Insight

The prose is commodity and depreciates with every model release. The decision data is proprietary and appreciates with every engagement.

That is the trust boundary made concrete for an investment firm. The five properties above are not aspirations bolted onto a research product; they are the product's spine. Model-agnostic orchestration is Choice. Private evals are Control. Private evaluation and decision-memory environments over real coverage are Capability. The compounding decision-memory ledger is Compound. A firm built this way does not rent intelligence from a foundation model — it accumulates judgment the foundation model never saw, and therefore cannot reproduce.

One honest qualification keeps the claim defensible, and stating it is a strength rather than a hedge. The *logic* of this architecture — the decision-memory layer, the evals, the model-independent operating procedure — is built and running today. Two things deepen over time rather than arriving fully formed. The first is the *longitudinal proof* that later calls measurably improve because of earlier ones, which can only accumulate as real outcomes resolve; the mechanism to compound exists now, the statistical track record is still being earned. The second is the *hardness of the boundary* itself, which today is enforced contractually and procedurally rather than by cryptographic, tenant-isolated infrastructure. Neither qualification weakens the thesis; both locate the firm honestly on the maturity curve. The learning-loop logic and the decision-memory architecture exist now; the harder infrastructure and the longitudinal proof are what deepen the boundary from here. That is a claim a firm can defend today — and it is a more credible one than pretending every layer of the future architecture is already built.

VI. The Trap — and the Commitment That Answers It

Here is the uncomfortable turn, and it is the part that separates a credible position from a clever one.

If a firm offers *its clients* a trust boundary — a place where the client's own investment judgment can compound safely — then that firm now stands exactly where the model provider stood in Arrow's inversion. It holds the client's decision data. It is, structurally, in a position to become the very extractor the whole argument warns against. A sophisticated client will see this immediately and ask the obvious question: *aren't you now doing to me what you told me to fear?*

The answer cannot be a reassurance; it has to be an architecture. The trust boundary must apply *recursively*. The five properties must be offered *to the client*, not merely practiced internally. The client owns their frame, their corrections, and their decision history. They can port it out. Any learning the firm derives from the client's data is consented to and visible. The boundary protects the client's edge from the firm with the same hardness it protects the firm's edge from the model provider. This is precisely the promise a foundation-model provider is structurally unable to make to a portfolio manager — and it is therefore not a concession but the sharpest possible differentiator. A firm that answers the extraction question honestly, in its architecture, is selling the one thing the current regime cannot.

VII. The Universal Form — Whoever Owns the Learning Loop

The paper has argued the investment case, because investing is where the paradox bites hardest and soonest. But the mechanism is not specific to investing, and the honest conclusion is broader than the frame it arrived in.

Medicine, law, engineering, consulting, scientific research, defense — every profession whose product is judgment runs on the same loop. A diagnosis corrected, a brief revised, a design rejected in review, an experiment that falsified a hypothesis: each is a correction, and each reveals the expert's standards and decision process far more richly than any raw record of what they did. In every one of these fields the same four questions now decide where value accrues:

Insight

Who owns the feedback? Who owns the corrections? Who owns the evaluations? Who owns the memory?

Whoever owns those owns the learning loop, and whoever owns the learning loop compounds the judgment. This is the deepest claim in the paper, and it outlives any particular model-provider policy: **the AI era will be defined by the ownership of learning loops.** Firms will not, in the end, be valued primarily by the data they hold or the models they can access — those commoditize. They will be valued by their ability to improve judgment faster than their competitors, continuously, inside a boundary they control. That is a profound shift in what a durable advantage even *is*, and the reverse information paradox is simply the first place the shift becomes impossible to ignore.

There is an even deeper way to state the principle, and it is worth ending the general argument on. What the trust boundary ultimately protects is not learning narrowly but *adaptation* — the organization's capacity to become better over time. Learning is one form of adaptation; so are the ways processes, incentives, and people evolve around it. Framed this way the claim turns almost biological: an organism that cannot keep its own adaptive machinery inside its own membrane does not remain an organism for long. The trust boundary is, in the end, a membrane around the right to keep getting better.

Pushed to its limit, this exposes the first principle sitting underneath the entire argument. Data, software, capital, and now AI are not ends in themselves; each is an instrument a firm acquires to make better judgments than its rivals.

Insight

Organizations exist to compound judgment.

Everything else — the data, the software, the capital, the models — is a means to that end. Once that is granted, the rest of this paper is a consequence rather than a claim: if compounding judgment is what a firm is *for*, then the loop through which judgment compounds is the one thing a firm cannot afford to rent, lease, or leak. It is not a component of the business. It is the business, expressed as a process.

VIII. Anticipated Objection — The Data Network Effect

The sharpest objection to everything above runs like this: foundation models improve through a data network effect — hundreds of millions of users contributing knowledge every day — and a model trained only on your own data, without everyone else's, would be useless. Pooling the world's knowledge into shared models is better for each user and effectively unstoppable; a firm that hoards its data to protect an edge is siloing itself out of the snowballing intelligence and will be left behind. It is a strong argument. It is also aimed at a position this paper does not hold.

The objection collapses two layers that have to be kept apart. There is a **general-capability layer** — a model trained on pooled human knowledge — and here the network effect is real, positive-sum, and worth riding: contribute to it, use the best model available, do not build your own. This paper argues *for* that, not against it. And there is a **firm-specific judgment layer** — your evaluations, your decision history, and above all the corrections that encode how *you* decide something is wrong. That layer is not a network good. Pooling it does not hand you a better shared model; it moves your distinctive view toward the consensus. For any business whose product is non-consensus insight, that is not a contribution — it is a transfer of the one thing that made it different.

Insight

Ride the network effect on general capability. Do not feed it the judgment that makes you different. Own the loop, not the model.

This is why the familiar network-effect analogies do not carry. Pooling your social graph, your documents, or your telemetry makes the shared product better *for you* — positive-sum, and the skeptics who told consumers to opt out deserved to fail. But a firm competing on insight is not a consumer of the pool; it is a competitor whose edge exists only while it is not shared. When its thesis enters the pool, the market prices the thesis in and the edge is gone. Network effects win wherever the pooled good benefits the contributor. Alpha is, by definition, the case where it does not.

Two concessions keep this honest. The general layer *will* keep snowballing, and a firm that fights it by training its own model is finished — but that is this paper's premise, not a rebuttal of it. And anyone who argues this has an incentive to: a firm that sells a trust boundary benefits from your believing you need one, exactly as the platforms that harvest the exhaust benefit from your believing you don't — weigh both accordingly. The strongest form of the objection is not about today's models but tomorrow's: that a sufficiently capable general model will reconstruct your judgment from public data and light context, dissolving the edge regardless of any boundary. That is the real bet, and the answer is structural rather than confident. Your *realized outcomes* and your corrections are private data that never enter the public pool, so the model cannot learn what it never sees; and a model trained on the consensus cannot, by construction, hold your non-consensus view for you. If those two facts stop being true, the thesis fails with them. Until they do, the boundary is where the edge lives.

IX. What to Do About It

The practical instruction reduces to a single reorientation. In the cloud era, enterprises accumulated data and learned to protect it. In the AI era, enterprises accumulate *learning*, and the boundary must evolve to protect the *mechanism* through which they learn, adapt, and compound — not merely the information they hold.

For an investment firm the checklist is short and testable. Own your evals, because they encode what good judgment means on your desk. Keep your traces, corrections, and decision history, because they are the distilled record of your experience and the richest thing you produce. Decouple your orchestration from any single model, because your veteran capability must outlive any one generalist. Route for cost without surrendering quality. And close the loop, so that every decision makes the next one sharper and the improvement stays inside your boundary.

A firm should be able to use a model without giving up the knowledge that makes it unique. In investing, that knowledge is the only durable private information there is. Protecting it is not an IT policy. It is the whole game — and increasingly, in every judgment business, it is the game itself.

In consuming intelligence, you create intelligence. In this business, the intelligence you create is your alpha. It should belong to you.

Insight

Models reason. Organizations remember. Competitive advantage belongs to whoever can combine the two without surrendering either.

Disclaimer. This paper is published by Lantern Labs for informational and educational purposes only. It sets out a conceptual and strategic framework and the author's opinions; it does not constitute investment advice, a recommendation, or an offer or solicitation to buy or sell any security or to adopt any investment strategy. Any companies, products, or examples named are illustrative and are not recommendations. Lantern Labs makes no representation or warranty as to the accuracy or completeness of the information herein and accepts no liability for any reliance placed on it. Readers should conduct their own analysis and consult appropriate professional advisers before making any decision. © 2026 Lantern Labs. www.lanternlabs.net