


Original Intelligence



What Stays Valuable When
AI Can Deliver Anything



APPLIEDXL RESEARCH

SPECIAL REPORT · JULY 2026

Original Intelligence

What Stays Valuable When AI Can Deliver Anything

For two centuries, information companies were paid for two things at once: originating what is true, and delivering it. AI has split the bundle and priced the halves at opposite extremes. Delivery is becoming a commodity. Origination is becoming the whole business.

BY FRANCESCO MARCONI

For the full interactive version of this report and the downloadable data, visit www.appliedxl.com/research/original-intelligence

CONTENTS

I	The lineage	5
II	The split	10
III	The craft	12
IV	The playbook	14
V	Delivery	16
VI	The prices	17
VII	The inflection	19
VIII	The domains	22
IX	The moves	24
X	The field	25
XI	The choice	26

In the spring of 1850, a man named Paul Julius Reuter set up shop in the German border town of Aachen, at the point where the telegraph line ran out. The wire reached Aachen from the east and reached Brussels from the west, but between the two lay about a hundred miles of Europe with no line at all, and that gap was Reuter's whole business. Each afternoon he waited for a bird to come in over the rooftops. Tied to its leg was the closing price of the Paris stock exchange, carried by wire as far as Brussels and then flown the rest of the way, and the pigeon crossed the gap six hours faster than the mail train. For those six hours Reuter knew what the Paris market had done and no one else in Germany did. A fast pigeon was worth more than a telegraph office, because it reached the one stretch the telegraph could not.

The pigeon was never the point, and Reuter never mistook it for the point. His edge did not last two years. The wire was extended across the gap he had built his business in, and the day the line was complete the birds were finished. The men who lost were the ones who had fallen in love with the bird. Reuter had not. When the telegraph swallowed his advantage he took up the telegraph, moved to London, and opened an office beside the exchange, and the company he built there is the one we still call Reuters. Each turn of the machinery left another set of rivals behind: the pigeon men undone by the telegraph, the telegraph men by the wire, the wire by the terminal. Every generation mistook the machine for the prize. Every generation was wrong the same way.

We are living through the largest such turn yet, and the industry has mistaken it for a funeral. Machines can now write without end and answer nearly any question, so the newsrooms and data companies are bracing to lose. They are hoarding what they have, suing over what has been taken, preparing for a smaller and meaner version of the work they used to do. Watch where the money agrees with them. The world is on track to spend roughly \$600 billion a year to justify what it pours into AI compute, and almost nothing to expand the supply of verified, original facts that compute exists to process. Every dollar is going to the machinery that reads. Almost none is going to the thing worth reading.

That is the misread, and it is expensive, because the companies bracing to lose are optimizing the exact half of their business the machines just made free. When any sentence can be generated for nothing, the sentence stops being the valuable thing. What holds its value is knowing what actually happened, and being able to prove it: the first contact with a true thing, before anyone else has it. That was always the real business. It is the one thing the machines cannot do. The telegraph could carry that Paris price anywhere on earth in seconds, but it could not stand in the exchange and watch the price settle. A person had to do that, and write the number down, before the wire had anything to send. We have built a far more powerful version of the same machine, and its limit is the same one.

Machines can retrieve the truth now. They cannot produce it.

There are two kinds of work in the information business.

- **Origination:** being first to a fact about the world. Finding what has newly become true, proving it at the source, and turning it into a signal someone can act on.
- **Delivery:** searching the recorded world, summarizing it, recombining it, and moving it back out in useful form.

Origination puts a new true thing into the record; delivery moves what is already there. For nearly two centuries the two were bundled and sold as one product, at one price, and no one had to ask which half they were paying for. AI has pulled them apart and priced each at an opposite extreme, and if you cannot tell which half your revenue rests on, you cannot tell whether it is the half that is vanishing or the half about to be worth more than ever. The originating half has a name: Original Intelligence (OI). It is the scarce half, and the rest of this essay is about what that changes.

AI is becoming extraordinary at delivery, moving through the recorded world at a speed no newsroom can match and drawing connections a person would never have time to find. What it cannot do is observe the world first, sit in the courtroom, call the source, notice the anomaly the instant it lands. However brilliant, delivery always runs on material that origination has already produced.

The companies that win the next decade will be the ones that treat their origination as the scarce asset and use AI to produce more of it than humans ever could by hand. That is a claim about winners, and it can be tested two ways: against the last two centuries, which show how value has moved and who captured it, and against the machine itself, which shows why this time the thing left standing is origination. Start with the pattern. ■

I. THE LINEAGE

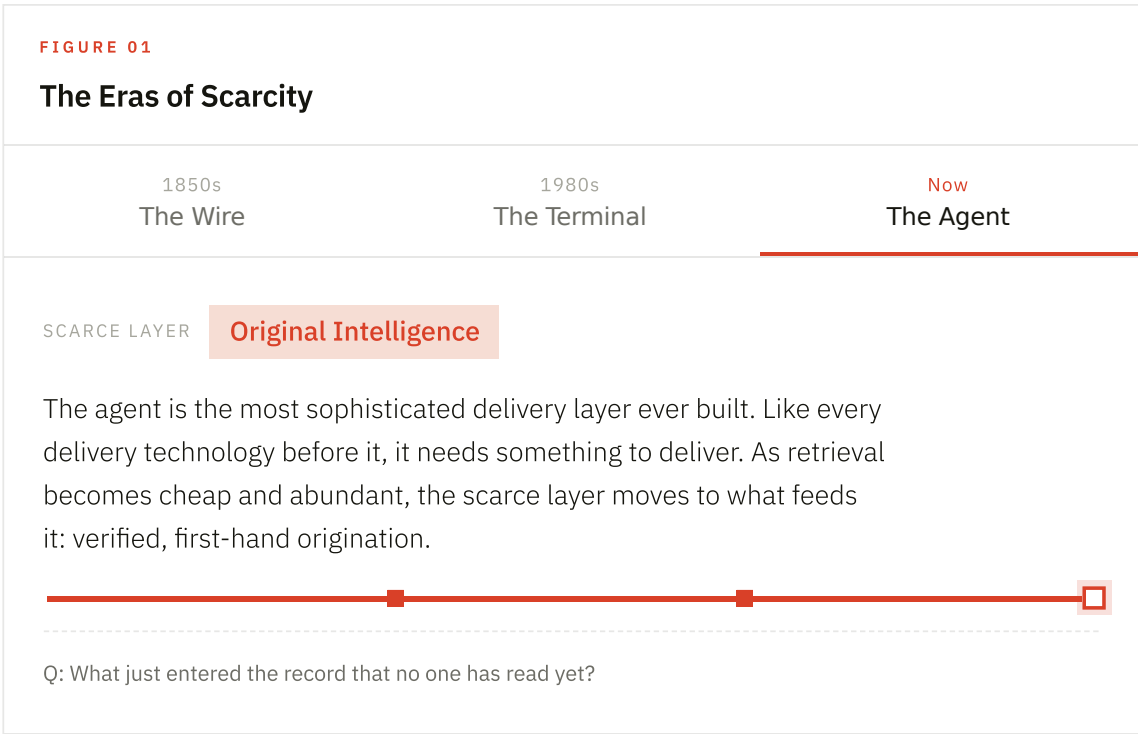
What has each era of moving a signal really rewarded?

The Wire. Before the telegraph, distance was the enemy; news moved at the speed of a horse. The telegraph collapsed distance to near zero, and that changed what a message was worth. When everyone receives the same news at the same moment, the advantage shifts from transport to selection and speed. This is the era Reuter was born into, and the lesson of his pigeons generalizes: the winners owned the fastest way to move a valuable signal, and when a newer machine made that speed ordinary, the advantage moved on. The business was never the pigeon or the cable. It was the edge itself, wherever it had gone next.

The Terminal. By the late twentieth century the wire's output had become overwhelming. There was too much to read and no way to act on it in context. LexisNexis built on networked document delivery, FactSet on databases, Bloomberg on networked computing. The value moved again, from the feed to structured data, tools, and workflow. A terminal organized the world into fields you could interrogate, and put that intelligence where decisions were made.

The Agent. The value is moving once more, and this time what changes is the receiver. In every earlier era the endpoint was a person: someone read the wire, someone worked the terminal. Now the endpoint is a mix: agents acting for people, agents answering other agents, sometimes no person involved at all. The person has not left; the person has moved up a level, from reading the signal to directing what reads and acts on it. And this is the moment the whole industry is bracing for as if it were the end — the reader replaced by a machine. Read correctly, it is not the end at all but the clearest signal yet of where the value is about to go.

The agent is the most sophisticated delivery layer ever built. Like every delivery technology before it, it needs something to deliver. As retrieval becomes cheap and abundant, the scarce layer moves to what feeds it: verified, first-hand origination.




Now put the eras side by side, and one rule appears. Each dominant information company started narrow, with one slice of data. Dun & Bradstreet, the business-credit bureau, began in 1841 with credit reports on merchants – and its innovation was not the data but the network of accountable named correspondents who filed it, a roster prestigious enough to include four future presidents, Lincoln, Grant, Cleveland, and McKinley, filing reports on the merchants of their districts. The name behind the claim was the product from the very first day. Poor’s, later half of S&P, began in 1860 with a book of railroad accounts.

Dow Jones began in the autumn of 1882, in a basement beside the New York Stock Exchange, where three men wrote the news by hand. Charles Dow, Edward Jones, and Charles Bergstresser worked over stacked tissue and carbon, so one pass of the stylus threw a dozen faint copies at once. Boys ran the copies out to the banks and brokerages several times a day, each slip a few minutes fresher than anything else on the Street. They called them flimsies. The afternoon roundup took a name, the Customers’ Afternoon Letter, and seven years later that sheet became The Wall Street Journal. The company was never the tissue or the boys. It was the few minutes.

Platts, today’s oil price benchmark, began in 1909 with a trade sheet for the young oil business. Each rode an emerging technology to scale, and each shift moved the value one step closer to the decision: first transport (moving the signal), then selection (picking out what matters), then structure and workflow (turning it into data and feeding it into the tools where work happens). A fifth step has opened above them, called simply action, where no person reads the signal at all; a system makes the decision itself, the trade placed, the approval issued. Picture these steps as a ladder, each step closer to the moment a decision gets made. That is the first of two patterns in the history, and it is this simple: over two centuries, the value has climbed the ladder, one step at a time.

FIGURE 02

The Decision Ladder



The rung is set by what the product does with the record, not by the technology that delivers it.

Every step, all the way up to a machine acting on its own at the top, does the same thing: it moves, organizes, or acts on a fact that already exists. Not one of them produces the fact. The higher the value climbs, the more it depends on the one thing none of the steps contain, the first observation that put a true thing into the record. Automate the whole ladder, let machines trade and approve and decide with no human involved, and you have only made the bottom step more valuable, because every step above it is still consuming what the bottom step produces. The ladder does not lead away from origination. It leads back to it.

It is easy to draw the wrong lesson here. Reuters rode the wire and Bloomberg the terminal, so the moral looks like: own the transmission layer. But the layer was never the point. These companies won by owning whatever was scarce at the time, and scarcity does not stay put. Each transmission technology, once it spreads, stops being the bottleneck and becomes cheap infrastructure everyone has, and the advantage moves to the next hardest link.


But does it stay scarce, when the same law just made transport, selection, and structure cheap in turn, turning each, once everyone had it, into a commodity no one will pay a premium for? This link is different. Every layer that got cheap was a way of moving or organizing information that already existed; each was a technology, and technologies spread until everyone has them. Origination works differently. It is the act of first contact with the world: reading the primary record, judging what in it matters, calling the source, putting a name behind the claim that it is real. That act does not spread that way, because there is no version of it that ends in software everyone can copy. It has to be performed again for every new fact, at the moment the fact appears.

Delivery is a machine you build once and run until the next machine. Origination is a cost you pay every single time. It is the one link in the whole chain that never turns into a cheap commodity, because it is not a way of moving information – it is the act that puts the information there to move.

There is a second pattern in the history, and it matters even more. Look at who settled each new layer. The Pony Express did not build the wire. Western Union did not build the terminals. Bloomberg was born on the trading desk: founded in 1981, its first terminals sat on Merrill Lynch desks by the end of 1982, and it never had a lower layer to climb from. We traced 235 information companies founded between 1792 and 2023 and logged the layer each began on. Of the 54 founded before 1950, none began at the two layers nearest the decision, workflow and action. Since 2005, more than one in four have, and the 2010s produced 53 new information companies, the largest cohort in the record. A layer opens only when its technology arrives; workflow had to wait for the computer. The pattern is what happens when it opens: the new layer is settled by companies born on it, not by the incumbents, the established players who already dominate the older layers, climbing up to it. Each generation is born closer to the decision than the last.

FIGURE 03

Companies by Starting Layer and Year Founded



Source and methodology: AppliedXL analysis of 235 information companies, compiled from the lineages of the major news, financial and market data, credit and risk, and scientific publishing businesses. Each is dated to the founding of its original product and coded by the layer that product occupied: transport, selection, structure, workflow, or action. Dots are jittered within each layer for legibility. The red curve is a kernel-weighted running average of the starting layer, drawn from 1845, where the record becomes dense enough to average. Across the set, founding year and starting layer correlate at $p = 0.41$ (Spearman, $n = 235$, $p < 0.001$).

Incumbents do move, but look closely at which half of the work moves. Bloomberg has rebuilt its own delivery again and again, from desk terminal to data feed to machine-readable pipe, always in place. What it has never grown in place is a seat at a new domain's record. Even Bloomberg News, built internally in 1990, read the same markets for the same desks. The new records were bought: clean energy intelligence with New Energy Finance, legal and regulatory intelligence with BNA. Thomson, a newspaper chain, became a legal and financial information company by buying West and then Reuters. Two centuries, one pattern: companies rebuild their own delivery, and they buy, found, or partner to reach a new record. They build their own pipes; they go elsewhere for new water.

Through all of it, two kinds of work were bundled and priced as one. Every one of these businesses did original work, finding and verifying information, then wrapped it in a technology that carried it to the people who needed it. You never had to separate the value of the origination from the value of the delivery, because no technology had pulled them apart. One finally has. ■

II. THE SPLIT

What can retrieval do, and what can't it?

A model does more than look things up. It reasons, synthesizes, follows a chain of instructions, and recombines what it has read in ways no one bothered to before. All of that is real, and getting better every month. But at bottom it works the recorded past: it operates on what has already been written down.

Retrieval has a hard limit, and it is the structural kind, not a gap the next model closes. A model works only with what has already been recorded, because it cannot observe the world directly. Give it two permits that are already filed and machine-readable and it will connect them well, often faster than a person. What it cannot do is surface the fact that is not yet in the record it can see, or the reading of scattered public facts that no one has yet assembled into a signal.

On November 28, 2023, the Supreme Court of Panama struck down a single mining contract, and a mine that produced roughly one and a half percent of the world's copper began to go dark. The company that owned it lost roughly half its value in the weeks that followed. Everyone called it a shock. It was not a shock. It was the last page of a story that had been unfolding in public for months: a disputed law, protests in the streets, a vote in the legislature, a court challenge anyone could have followed. All of it was visible. None of it had been assembled into anything you could act on until the price had already moved.

The clearest sign was hiding in the most boring place imaginable. Every shipment of copper that leaves Panama gets logged, and the size of those shipments is public record – and copper was nearly three-quarters of everything the country exported. Month after month it held steady; then it collapsed toward nothing, visible in the shipping data in near-real time. The company that ran the mine did not even withdraw its own production forecast for the year until the first of December, days after the ruling. The supply was stopping in plain sight, in a public ledger, before the official numbers had caught up to it. Anyone who thought to watch could have watched it happen. Almost no one did, because noticing means knowing to read a customs ledger against a court ruling against a street protest, and knowing that the three are the same story. The facts were public. The reading was the rare part.

Someone has to observe that and write it down first. Retrieval always comes after that first act of origination.

Doesn't AI already do discovery, the drug discovery, materials discovery, scientific discovery we keep hearing about? It does, in one sense: it generates novel candidates by searching and recombining what is already known. A model can propose a molecule no chemist wrote down. But a proposed molecule is a hypothesis, not a fact. What makes it true is a clinical trial: someone dosing real patients and observing what

happens in the world. The model cannot run the trial, and cannot know the result until the result exists and has been recorded. That is the divide in miniature: AI is extraordinary at generating hypotheses from the record, and still cannot originate the fact that settles them. Origination is that second act, first contact with a truth the record does not yet contain.

The honest complication is that machines are starting to observe too. Sensors, satellites, and real-time filing feeds now capture raw events cheaply and around the clock, and that part spreads and gets cheap like any technology. But raw capture is not origination. A satellite counts cars in a lot; it does not know the retailer is about to miss its quarter. A feed logs a filing the instant it posts; it does not know the buried clause changes the outcome. The alt-data business already lived this in miniature. For a while a satellite photo of a retailer's parking lot was worth real money, because almost no one had one. Then everyone had one, the picture became ordinary and ignored, and the money moved to the few analysts who could look at the same image and know which chain and which quarter it was quietly predicting.

The camera got cheap. Knowing what the camera was looking at did not.

That knowing is a bundle of judgments no sensor makes: deciding what is worth observing, structuring the raw signal into a fact, verifying it well enough to put a name behind it. That is the layer that stays scarce, and it is why AI is leverage rather than threat: the machinery industrializes the capture and frees the scarce judgment to be spent across far more facts than a newsroom could reach by hand. The cost per fact does not vanish. It moves higher up, to the one link that was never a technology.

So AI, for all its power, is the most sophisticated delivery technology ever built, the terminal's natural successor, and like every one before it, it needs something to deliver. It also moves the whole competition. The old advantage was comprehensiveness of the archive; the new one is freshness, how fast a true thing enters the record and reaches a decision. That is a frontier retrieval cannot hold on its own. ■

III. THE CRAFT

What do information specialists actually do?

Origination is a craft, and it belongs to a particular kind of worker. Call them information specialists: journalists and analysts trained to find what is new, prove what is true, and understand what it means before it is obvious.

An information specialist finds what is different. Out of a flood of filings, disclosures, permits, trial records, and transcripts that all look routine, the job is to spot the one that is not: the anomaly, the number that should not be there, the change in wording, the quiet update that alters the picture. Machines are good at flagging the outlier; that part is pattern-matching, and it keeps improving. The harder judgment is which outlier matters, which deviation is a story and which is noise. That call turns on knowing the domain and the stakes, and it is a different kind of work than detection.

An information specialist gets the dot that isn't online yet. A single record is rarely the story. The story appears when a filing is set beside a lawsuit, a permit, a hiring pattern, a registry update. A model can join records that sit in front of it, and often finds links a person would miss, so the edge that stays human is not the joining. It is getting the dot that is not yet online, not yet filed, not yet written down, so that there is something new to connect at all. The best connections are the freshest, and the freshest have to be fetched from the world, not the archive.

An information specialist thinks about implications. Knowing what happened is only the beginning. The value is in reasoning forward: if this is true, what does it change, who does it affect, and what does it signal before the market, the public, or the institution has absorbed it?

Those three together reveal something. Original Intelligence is not only a talent; it is a *procedure*. A good analyst does not simply know things, they know the *steps*: which records to pull, in what order to cross-reference them, how to weigh each source against the others until the picture holds. Call that ordered sequence the human algorithm, the part a model can run once it has been written down but cannot supply on its own. And here is the concession the honest version of this argument has to make: yes, the method itself is software, and software spreads, so a competitor can copy the steps. But the steps are not the scarce thing. What the steps run on is: the live access to the record that a newcomer has not negotiated, the domain knowledge it took years to encode, and the accountable name that makes a buyer trust the output.

You can copy the recipe. You cannot copy the kitchen, the suppliers, or the chef's reputation — and those are what the recipe needs to produce anything.

Suppose the question is whether an already approved drug could treat a disease it was never designed for — the kind of overlooked opportunity that never becomes news because no company has a reason to chase it. An expert does not answer by looking it up. She runs a sequence.

Below the headline result of the trial the drug was approved on sits a secondary endpoint, a health measure the trial tracked but was not designed to prove, that moved when it had no business moving, a few points in a direction that means nothing unless you happen to know the biology of the second disease. She opens the trial's registry record and finds an amendment slipped in eight months after patients were enrolled: the eligibility criteria widened, then narrowed again, the fingerprints of a sponsor who saw something and chose to say nothing. She pulls the filings that log the drug's side effects and finds one that, read against the second disease, is not a side effect at all. It is the mechanism. Then she calls someone who ran a trial site, and learns the study was stopped for a reason that never reached the published paper.

None of that is retrieval. It is origination, synthesis, and judgment applied to what is happening now. A model can assist every part of it, but left to run on its own, it can only report what has already been recorded. Being first to see what changed is a different job, and it is the one that pays.

FIGURE 04

Case Study | The Human Algorithm: How an Analyst Finds a Repurposed Drug

THE QUESTION

Could a drug the FDA already approved treat a different disease — one it was never designed for — that no company is bothering to pursue?

01

02

03

04

05

06

STEP 01

FDA approvals · drug labels

Start with a drug that already works

Pick a medicine the FDA has already approved for something. It's already proven safe in people — that's a huge head start.

IV. THE PLAYBOOK

What does an Original Intelligence company build?

Most information companies, staring at AI, are asking how to become an AI company. That is the mistake. The winning move is to become an *OI company*, and pick up AI as the tool that makes it possible. An AI company competes on the delivery layer, against the best-funded firms on earth. An OI company competes on the one thing those firms cannot manufacture: verified origination, a true thing found first.

The template is not proprietary; it is simply what we built AppliedXL to do, and the shape of it is available to anyone. Use AI heavily, the way a newsroom once used the printing press, as the machinery that lets a small team work at a scale that would otherwise be impossible, but aim it at producing new information rather than repackaging what exists. AI is the tool. OI is the product. Any company that already originates could run the same play.

The method underneath is the same whoever runs it, and it starts with a rule most information businesses violate: go straight to the primary institutional record, the filings, clinical results, permits, and registries where market-moving signals surface before they become news, with no scraper and no third-party aggregation layer sitting between you and the source. Own that first contact, then turn it into product. What ships is one of four things: a signal that something material just changed, a benchmark, a probability forecast, or a resolution record – the settled outcome of a defined question with the proof attached.

FIGURE 05

Source → Deliver

The OI Architecture

- 01 Source**
 Pull the primary institutional record — filings, clinical results, permits, registries — the moment it lands.
- 02 Structure**
 Turn dense, technical documents into machine-readable events, each field tied back to the source.
- 03 Detect**
 Find what is different in the record: the anomaly, the number that shouldn't be there, the exception worth a second look.
- 04 Aggregate**
 Join those events across time and entities — the connection a single filing can't make alone.
- 05 Verify**
 Validate every output against the source record, check provenance, and route it through human review. Speed without verification is only noise.
- 06 Contextualize**
 Reason about what it means: who it affects, what it signals about a market, and why it matters before the signal is obvious.
- 07 Deliver**
 Ship it where a decision is made — as a signal, indicator, forecast, or resolution record.

The rigor is journalistic, not merely algorithmic. Outputs are validated against source records, traced back to where each fact came from, and reviewed by human analysts before they ship, because speed without verification is only noise. OI is editorial work with verification built in from the start, not a data pipeline with a check bolted on at the end.

If that sounds too modest to matter, consider what a journalistic method already prices. The benchmark that much of the world's physical crude oil settles against is not set by an exchange. It is set by reporters. Every trading day, in a defined window that closes at 16:30:00 London time, price reporters at Platts collect the bids, offers, and confirmed deals that traders and brokers relay to them, many literally by phone, the exact time of the call recorded, and, applying published methodology and editorial judgment, publish a single number: the day's assessment for Dated Brent. That number then flows through official selling-price formulas and derivative settlements into hundreds of billions of dollars of physical trade. It is not a summary of the market. For much of the market, it is the price. A reporter's sourcing, structured into a defined procedure and stamped with a name, became the settlement layer for one of the largest physical markets on earth. That is journalism installed as infrastructure, not journalism decorating a market, and it is what origination becomes when it is done to a standard the market can build on. ■

V. DELIVERY

How is delivery rebuilt for the agent era?

O I is only half the product. The other half is getting it to a decision in the form the decision needs, and that is the half everyone else is also racing to build, which is why it cannot be where you win. Delivery is the minimum price of entry. You have to do it excellently and it will never, by itself, set you apart.

But one thing about delivery has genuinely changed, and it is stranger than it sounds. For nearly two centuries, delivery meant formatting a signal for a person to read. In the agent era the reader is increasingly not a person at all – it is another model, querying and cross-referencing and acting with no human looking at the page. That inverts the whole design. The signal now has to be structured, source-linked, and directly interrogable, so a machine can consume it and trace it back to the record it came from.

And follow what that does to the economics. When the reader is a machine, delivery gets commoditized faster than ever, because any capable model can retrieve, summarize, and format – those were the hard human skills delivery used to charge for, and now any competitor's model has them too. When every delivery surface is roughly as good as every other, the only thing left to compete on is the quality of what flows through them: the originated fact underneath. So the machine audience does not threaten origination. It strips the value out of everything above origination and leaves it nowhere to pool but the one input a machine cannot produce.

FIGURE 06

Two Ways at Once

Delivery, Rebuilt for the Agent Era

PUSH

PULL

Push — structured feeds, tuned to an audience of one.

You pick the sources, events, format, and cadence. Intelligence arrives the moment something matters — the customization of the best terminals, narrowed from a shared screen to a feed built for one reader.

SOURCE



YOU

Push delivers the origination the moment it happens. Pull lets you interrogate it on demand.

VI. THE PRICES

Why is content going free while intelligence stays priced?

Somewhere in every one of these companies is a number that used to be safe and is now falling. A subscription that renews a little less often. An ad rate that softens each quarter. A licensing fee a customer is quietly starting to question, because the thing it paid for is now something a model hands out for free. That is not a distant threat; it is the current quarter. This section is about why that number is falling, and why a different number on the same income statement is about to rise.

Commodity content is heading toward free. A model generates unlimited amounts of it, and the cost of one more article, summary, or explainer is falling toward zero. Not all of it collapses at the same rate (distinctive reporting and trusted brands still hold subscription value), but the vast undifferentiated middle does, and it is that middle most attention businesses were built on. You cannot charge for, or advertise against, what the reader gets for nothing. And the erosion runs one layer deeper: as models ingest and reproduce open and licensed content, raw information itself stops being defensible, and durable value migrates to structured, decision-grade intelligence that drives a customer's revenue.

Intelligence has not moved the same way. It still commands high prices, because it is a different product: a structured, verified signal a professional pays for because it drives a decision with money on the other side. A portfolio manager does not want more content. They want the one thing that changes what they do next, in a form they can act on and trust.

This split is older than the technology forcing it. John Moody published a manual of statistics, tables of railroad figures anyone could tabulate, and the 1907 panic wiped it out; he lost the business. He came back in 1909 selling something the tables could not: a letter grade, a compressed judgment on whether a railroad's bonds would pay. The compilation had died in a panic; the opinion survived every panic since, and in 1975 the SEC wrote the ratings agencies into regulation, converting private judgment into required infrastructure. Data commoditized first, and judgment kept its price, a century before a model made the same thing happen to everything else.

FIGURE 07

Price gap: 70pts

The Divergence



The fair objection is that AI will not stop at content; it will get cheap for the part that is really retrieval and formatting. But the scarce core is different. What a professional pays a premium for is not the packaging of a signal but the assurance that it is real: that someone with domain expertise found it in the primary record, verified it, and put their name on it. As content collapses in price, the assurance does not. It becomes the product, and it attaches to a specific class of material: signals structured from primary institutional records, the regulatory filings, court documents, permits, and registries, before they surface as news. You cannot reproduce that by ingesting the open web, which is why it survives the repricing while commodity information does not.

Return, now, to the number this essay opened on, because the divergence explains it.

FIGURE 08

Reading vs. Creating

2026 · Where the Money Goes

AI infrastructure
THE DELIVERY LAYER

~\$700B

Producing new
intelligence
THE SCARCE LAYER

The world is building the machinery to read at unprecedented scale — and there is no comparable line item for creating the signal worth reading.

Sequoia's David Cahn estimates ~\$600B in annual revenue is needed just to justify the compute spend, and calls GPU computing "increasingly a commodity, metered by the hour." Value migrates to what feeds a commoditizing layer.

VII. THE INFLECTION

What survives when the bundle splits?

The essay's central claim has a sharp consequence for the companies that dominate professional information, and it is not the one usually drawn. The reflex is to say they are exposed. What is really happening is that their old bundle is being split into its two halves, and the halves are moving in opposite directions at once. One is losing its scarcity. The other is becoming the scarcest asset in the industry. Everything depends on which half a company decides it is in: whether it believes the value was in holding the record, or in being the one who read it first.

The half that is losing scarcity is the stored past. The durable advantage used to be the archive, a structured record accumulated over decades, and that record is now something an AI-native competitor can license and cross in a single acquisition. But note what commoditizes, because it is narrower than it first appears. What gets cheap is the warehouse of already-recorded facts. What does not get cheap is the live act this essay keeps returning to: first contact with a fact the record does not yet contain. The archive was delivery, frozen. Origination is the live version, and it does not sit in the warehouse that just lost its value. The archive stops being the product and becomes the fuel: decades of structured, domain-specific records are exactly what an AI system needs to recognize what is anomalous in the record arriving today. Sold as a static pile, the archive is commoditizing. Turned on the present as context for producing new signal, it is an advantage a newcomer cannot buy.

That is the split, and it lands on the incumbent as a live capability, not a stored one. The frontier moves from how complete your history is to how fresh your signal is, and freshness is not something you own, it is something you perform, every time, at the moment the fact appears. Think of what that word actually means. It is a person, or a system a person built, reading a filing the hour it posts and knowing what in it matters – the same act as a reporter standing in a courtroom, the same act as Reuter watching a price settle. It cannot be done once and stored away. It has to be performed again tomorrow, and the day after, for every new fact the world produces. The organizations built to optimize the completeness of the record now find the competition has moved to the one link they always treated as a cost: the standing apparatus for being first.

One half is losing its scarcity. The other is becoming the scarcest asset in the industry.

That is why the surviving half belongs to them more than to anyone else. Two things survive the split, and the established players hold both. The first is the customer: the installed base, the accounts already won, the reach a new entrant would spend a decade building. The second is the standing apparatus itself, the negotiated feeds and held credentials that keep the primary record flowing in the moment it is written, the one input origination cannot proceed without and cannot be gotten by scraping the open web. So the incumbent's position resolves into a single asymmetry: the hard half to build is the intelligence layer, which takes editorial discipline, domain mapping, and standing presence at the source that no tool hands a competitor; the easy half is reach, which the incumbent already owns. The two halves need each other, and unequally. That is why, again and again, the companies that hold the origination layer get partnered with rather than competed away.

There is a third property of the seat that the history makes unmissable, and it is one capital cannot buy: neutrality. In 1986 Citicorp bought Quotron, then the icon of market data with a hundred thousand terminals and sixty percent of the market, on the theory that banking was an information business. Its largest customer, Merrill Lynch, would not keep feeding its quotes and orders through a box owned by a competing bank, declined to renew, and put its money instead into a startup called Bloomberg. Quotron lost money every year after; Citicorp eventually paid Reuters over a hundred million dollars to take it away. The technology aged badly too, which is the honest other half of the story, but the founding wound was ownership. A seat at the record has to be trusted by the whole market it serves, and the moment a market participant owns it, the rest of the market leaves. That is no historical curiosity but a structural fact: the origination layer is most defensible in independent hands.

None of this forces a particular path, but two centuries of the industry price each option: incumbents have always rebuilt their own delivery in place, and new origination has always arrived by transaction, an acquisition, a purpose-built entity, or a partnership. So assembling origination in-house is not a retooling project. It is a founding, with a founding's timeline and odds, run inside a company optimized for something else. Buying or partnering delivers the same layer at the speed of a contract, which is the route the record favors. And waiting for origination to commoditize is waiting for the one thing the pattern has never delivered. On the credit question alone, forty-one companies have re-answered it since 1841. Gaps between new entrants once ran nineteen years; since 2000 the median is two. The forty-one were not fighting over one chair; the question kept opening new records and new layers, and each was claimed as it opened. Seats do occasionally flip – in 2009 Saudi Aramco moved its US crude pricing off a Platts marker it had used since 1994 to Argus's Sour Crude Index – but look at what that took: a sovereign oil producer, a structural dislocation at Cushing, and even then only one regional formula, with Brent left untouched. The price of moving a seat is measured in sovereign decisions and decade-long dislocations, not in a competitor shipping a better tool. The seat is per-domain, and it is filled once.

That last phrase was once a literal contract. In 1859, Reuter and Wolff (the latter had earlier worked inside Havas's Paris office) met the Havas heirs at the Hôtel Bullion and carved Europe into exclusive territories; in 1870 the three formalized it into the Ring Combination, dividing most of the world outside North America among them and holding the arrangement for roughly seventy-five years. AP spent decades trapped inside it, and its general manager Kent Cooper fought his way to full independence only in 1934, writing the campaign up in *Barriers Down*, arguing that exclusive news territories were barriers to the truth itself. The seat, in other words, has always been a thing that gets claimed and held per domain. The only difference now is that it is enforced by an installed workflow rather than a treaty, and contested by product rather than by cartel.

Print monetized distribution. Digital monetized data. Social monetized attention. The next era will not monetize information at all. It will pay for producing it. ■

VIII. THE DOMAINS


Where does the next record open?

If seats are claimed per domain, the useful question is which domains are opening. The same dataset that tracks the layers answers it, coded a second way: not by how close each company sat to the decision, but by the subject it covered and whether it served a professional or a general audience. Read that way, the birth record tells a story most people in the industry would get wrong.

Start with the founding era, because the common picture of it is inaccurate. The reflex is that professional information began as a finance story, Wall Street data for Wall Street firms. It did not. Before 1900, health and life sciences accounted for 32 percent of professional-information births, exactly matching finance. Lippincott began publishing medical texts in 1792; The Lancet launched in 1823. They stand in the record beside Poor's and Dun & Bradstreet, from the very beginning. The origin was never Wall Street alone. It was Wall Street and the clinic, side by side.

FIGURE 09

Where Information Companies Are Born, by Domain



Source and methodology: AppliedXL analysis of 235 information companies, 1792–2023. Shown: the 170 professional-information companies; the 65 audience-media businesses are excluded pending subject coding. Each company is dated to the founding of its original product and stacked by decade and domain. Finance: market data, ratings, indices, research. Health & Life Sciences: scientific publishing and clinical intelligence. Security & Systemic Risk: cyber, ESG, sanctions, risk ratings. Energy: commodity prices and energy intelligence. Legal & Policy: legal and regulatory research. Other professional: market research, supply-chain and enterprise data.

What happened to health next is the mechanic worth understanding. It fell from a third of professional births to zero for the fifty years from 1900 to 1949, and stayed near zero for eighty. Then, between 2000 and 2007, trial registries and mandatory results disclosure arrived, and a domain that had produced almost no companies for three generations began producing them again within five years. Health is now 23 percent of professional births since 2005, and five of the ten companies born so far in the 2020s are health companies.

A domain, in other words, is not a permanent market. It is a function of whether its primary record exists and is readable, and you can watch it open and close. And that turns the record into a forward signal. Wherever a mandatory record is being created right now, a birth cohort is roughly five years behind it. Carbon registries, emissions disclosure, AI incident reporting: each is a record coming into existence, and each marks a domain that has not yet been fully claimed.

Finance needs a correction too. Its dominance was real, but it was one phase, not the nature of the business. Finance peaked at 55 percent of professional births in the terminal era of 1985 to 2004, and has fallen about twenty points since. The instinct that “information business” means “financial data” describes a twenty-year window, now closing, not a permanent fact.

Two patterns in this second cut do more than decorate the argument; they validate it. First, the rule that new layers are settled by new companies holds one level up: new *domains* are settled by new companies too. Security and systemic-risk intelligence barely existed before 1985 and has run at 12 to 15 percent of professional births since, built almost entirely by companies born into it, not by incumbents extending sideways. The thesis is true at two levels of aggregation at once. Second, the entire modern acceleration is professional. Audience media has held flat at roughly a fifth of births for seventy years; every bit of the climb to the 2010s peak came from the professional side. Building B2B intelligence does not go against the direction of the industry. It is what the birth record has been doing for two generations.

Put the two dimensions together and the framework becomes a map. Cross a domain's momentum with how recently its highest layer was settled, and the open seats become visible: health at the action layer is being claimed now, while energy and materials show records being created without a corresponding wave of companies yet, which is either a gap in the data or a gap in the market. Either way, it is where to look. ■

IX. THE MOVES

What do you actually do on Monday?

B uild B2B products. Sell structured intelligence to the professionals who will pay for it, instead of selling advertising against content that is increasingly free. The reader economy is being commoditized by the same models flooding it. The professional economy is not, because what a professional needs is the one signal that changes a decision, and that was never a page view.

Monetize data with AI. Most established information organizations already sit on something rare: a body of records they have the authority to read, gather, or produce. That material used to sit idle, too vast to turn into product economically. AI changes the math – turning a cost center into a franchise. The winners are not the ones with the most data; ownership alone no longer protects anyone. They are the ones that execute intelligence on the data they already have the right to.

Know who is buying. Scarce is not the same as valuable; something is only worth producing if someone pays for it. Origination has three buyers forming at once, and all three are buying the same thing, a verified fact defensible line by line, for different reasons.

- **Professionals pay for decisions.** A portfolio manager, a regulator, a clinician needs one accountable answer with money or lives on the other side, and will not stake that on a model's unsourced guess.
- **AI platforms pay for ground truth.** The models flooding the world with cheap content need verified, source-linked facts to check themselves against; a system that generates fluent text has every incentive to buy the one thing it cannot generate, proof that a claim is real.
- **Prediction and resolution markets pay for settled outcomes.** A market on whether a drug was approved, a contract awarded, a target met needs an authoritative source to resolve against, and that source is precisely an origination layer that can say what happened and show the evidence.

Three buyers, one product. None of them is paying for delivery.

Together these mark the real shift: the advantage of information companies is moving from data ownership to intelligence execution. Anyone can run a record through a model. Turning that into a signal a professional will stake a decision on is the hard part, and the part that compounds.

Execution is the new ownership. ■

X. THE FIELD

If everyone is building the same pipes, what gets rare?

Look across the market for professional information and the convergence is striking. Very different companies are all racing to build the same thing: the best AI-driven way to find, query, and reason over information. AI companies are building workflow tools for specific professions. Established information companies are putting AI in front of their own data. Newer entrants compete on search.

This is not a story of winners and losers among them. Each is competing to be the best way to retrieve and reason over information, the delivery layer, and each is only as valuable as the Original Intelligence flowing into it. As the market fills with excellent, competing delivery surfaces, the layer that grows scarcer is the one beneath them all.

Everyone is racing to build the pipes. The water is what gets rare. ■

XI. THE CHOICE

What is in front of every Original Intelligence company?

News organizations and information companies are the original OI companies. They have spent decades, some more than a century, building the one thing AI cannot manufacture: the ability to find what is true and new and turn it into something a person can act on. That is not a sentimental claim about the value of journalism. It is the coldest economic fact in this essay. The scarcest asset in the information economy is the thing these companies already know how to make, and too many of them are spending this moment on defense, mourning a business that is not the one that is dying.

Because here is what should change the whole calculation: the winning move asks almost nothing they do not already have. It does not mean building an AI lab or out-engineering the frontier labs. The hard, compounding part was never the technology – it was the editorial judgment, the verification, the standing at the source, and the trust that comes from a name attached to being right. Those took a century to build and cannot be bought. The machinery that turns them into intelligence at scale is the part that is now cheap, and it can be assembled in a quarter. The industry has the impossible half already. It is being asked to pick up the easy half and refusing, because it has mistaken the easy half for the whole game.

And the window does not stay open. Every domain has a primary record and a first mover who claims it. In each one, the intelligence layer gets built once, on whoever holds the sources and moves first, and everyone who arrives after buys access to a seat someone else is already sitting in. It stays one seat because the buyer wants one accountable answer per question, and a verified feed, wired into a workflow, is an installed fact the next vendor has to argue against. That is not a forecast. It is already happening, vertical by vertical – one domain's seat claimed this quarter, another's the next.

So the choice is smaller and starker than the funeral makes it look. The tool everyone fears is the tool that makes the scarce thing scalable for the first time in the industry's history. Point it at the cheap half and you help commoditize yourself. Point it at the expensive half, the finding, the proving, the being first, and the thing that felt like the end becomes the largest lever anyone in this business has ever been handed.

The facts are public. They always were. The reading was always the rare part, and it is about to become the only part anyone will pay for. Machines can retrieve the truth now. They cannot produce it. That work belongs to the originators, and it always has. The only question that remains is whether they will pick up what is theirs before someone else sits in the seat. There is a clock on that answer, and it is running now. ■

Original Intelligence

What Stays Valuable When
AI Can Deliver Anything

AppliedXL

