

Who Interprets the Company Now?

AI, Machine-Mediated Inference, and the New Contest for Investor Conviction

Recently Jennifer Parmelee Witt, Managing Director, interviewed Mark Hayes, Founding Partner of Breakwater Capital Markets, about a Capital Markets consequence of AI that goes well beyond generative content. As machines increasingly participate in screening, monitoring, comparison, thesis testing, and the reconstruction of company economics, public markets are beginning to shift from a competition for information toward a competition over interpretation. The challenge for management is not to communicate to algorithms instead of investors. It is to ensure that the company's economic reality survives an information system in which machines increasingly influence what investors notice, how evidence is categorized, which assumptions are challenged, and when an existing view of the company is revised.

Jennifer: Mark, we have heard almost every possible claim about AI. What is the argument here that you think the market still has not fully absorbed?

Mark: The underappreciated change is that AI is transforming both sides of the public market information system simultaneously. It is changing the economics of the companies that investors are valuing, but it is also changing the machinery through which investors construct an understanding of those companies. Those are fundamentally different developments. One belongs inside the operating model: revenue, productivity, capital intensity, ROIC, competitive advantage, and risk. The other belongs inside the epistemology of the market: how evidence is collected, compared, prioritized, interpreted, and ultimately converted into conviction. For a long time, the Capital Markets advantage was partly an information-access problem. Who had the right datasets, who read the filing fastest, who had the deepest historical model, who could assemble competitor commentary most efficiently? AI compresses a meaningful portion of that advantage because retrieval, comparison, summarization, and basic synthesis become cheaper. The research captures the migration clearly: management judgment under uncertainty scores 96 in information value, changed assumptions and beliefs 94, decision and reallocation signals 92, disconfirming evidence 90, and scenario intelligence 89, while generic peer benchmarking is only 48. As machines become better at telling investors what is publicly knowable, the scarce information increasingly becomes what machines cannot confidently infer: why management changed its mind, which evidence it trusts, what decision followed, and what would cause management to reverse that decision. AI is therefore not eliminating the value of management judgment. It is stripping away many of the lower-value informational activities surrounding it and making judgment more visible as the scarce resource.

Jennifer: So, are you saying information itself is becoming commoditized?

Mark: I think a large category of information is. That does not mean facts cease to matter; it means possession of broadly available facts becomes less differentiating when sophisticated systems can retrieve and organize them almost instantly. Historical financials, previous guidance, transcript language, peer margins, competitor references, and many forms of public-domain benchmarking increasingly become infrastructure rather than edge. The investor does not stop needing them. The investor stops earning much of an intellectual premium merely for assembling them. That changes the nature of analytical advantage because if everyone can retrieve substantially the same evidence, the important differences increasingly lie in causal interpretation. Which evidence should cause the investor to change an assumption? Which apparent signal is actually noise? Which external datapoint deserves more weight than management commentary? Does a competitor's price reduction indicate temporary tactical behavior or a structural deterioration in industry economics? Does rising AI spending imply a widening moat or a new cost of remaining competitive? Those are not retrieval problems. They are problems of judgment under uncertainty, and AI can actually make them harder by increasing the speed and volume with which contradictory evidence enters the system.

Jennifer: You are using the word “epistemology,” which is not common IR language. What exactly has changed about how the market knows what it thinks it knows?

Mark: Historically, investors had to make difficult judgments partly because information was incomplete or difficult to obtain. Increasingly, they will make difficult judgments despite having extraordinary amounts of information. That is an important inversion. More information does not necessarily produce more knowledge because the marginal problem shifts from acquisition to selection, weighting, causal interpretation, and belief revision. An investor may simultaneously have management guidance, transaction data, competitor pricing, channel commentary, supply-chain evidence, expert analysis, policy developments, customer reviews, alternative data, and machine-generated synthesis of all of them. The problem is no longer whether there is enough evidence. The problem is whether the evidence belongs in the same causal story, whether some signals deserve much greater weight than others, and whether the model being used to interpret them is still appropriate. An AI system can ingest a thousand signals and still misunderstand the business if it is solving for the wrong relationship. That is why I think AI pushes Capital Markets away from a primarily informational contest and toward an interpretive one. The advantage increasingly belongs to whoever can construct the most accurate causal model, recognize when that model is failing, and update it before the evidence becomes obvious to everyone else.

Jennifer: If machines are now participating in interpretation, are they becoming an investor audience in their own right?

Mark: Not in the sense that they bear fiduciary responsibility or own the capital, but machines are becoming an increasingly important intermediary in the formation of human judgment. An AI system does not need autonomous authority over a portfolio to influence capital allocation. If it decides which changes to surface, which disclosures to compare, which risks to flag, which historical passages to retrieve, or which assumptions appear inconsistent, it has already shaped the environment in which the portfolio manager thinks. The emerging system is therefore neither purely human nor purely machine. Institutional research workflows increasingly use AI-assisted screening, monitoring, thesis testing, and risk analysis; scenarios can be tested across more variables and at greater speed; and the formal reporting cycle no longer controls the interpretation cycle because investors continue refreshing the thesis between disclosures. The practical consequence is that company information increasingly passes through an inference layer before reaching the ultimate decisionmaker. Management still communicates with investors, but what the investor encounters may first have been selected, compressed, classified, and compared by a machine.

Jennifer: That seems to create a new risk: machine-mediated misclassification. What does that actually look like?

Mark: It is the risk that a system reconstructs the wrong economic meaning from information that may be individually accurate. Suppose margins decline because management is funding a temporary implementation program expected to reduce service costs materially over the next six quarters. If the causal relationship between current spending and future cost economics is poorly specified, a machine may classify the margin decline as structural deterioration. Suppose CapEx rises because the company is building an advantaged infrastructure position, but the disclosure does not distinguish growth investment from maintenance or defensive spending. The system may infer worsening capital intensity when management believes it is extending competitive duration. These errors matter because once the wrong classification enters a research workflow, subsequent evidence can begin to be interpreted through it. Higher spending confirms the “deteriorating capital intensity” thesis. Temporary margin weakness confirms the “structural pressure” thesis. A machine does not have to invent a falsehood for the analysis to go wrong; it can combine accurate facts inside an incorrect causal model. That is why economic misclassification may ultimately matter more than ordinary summarization error. It can create a coherent but wrong representation of the business.

Jennifer: But is management responsible when an external AI system constructs the wrong model from accurate disclosure?

Mark: Not automatically. Management cannot assume responsibility for every external inference, just as it has never been responsible for every analyst model. Investors own their judgments. But there is a meaningful difference between disagreement caused by genuine differences in economic interpretation and disagreement caused by avoidable ambiguity in the company's own information architecture. If one investor believes the competitive advantage lasts ten years and another believes it lasts five, that is healthy market disagreement. If one investor thinks a metric includes a particular segment while another does not because the definition drifted without a clear reconciliation, that is an information integrity problem. If one system treats an expenditure as maintenance and another as growth because management itself has not clearly distinguished the two, unnecessary uncertainty has entered the underwriting process. The standard should therefore not be whether management can control the conclusion. It should be whether the information allows sophisticated interpreters to begin from a reasonably common reconstruction of the underlying economics before they disagree about value.

Jennifer: That sounds almost like external model governance. Are you saying IR becomes responsible for governing investors' models?

Mark: No, and that boundary is important. The company does not control the investor's model, probabilities, valuation methodology, or conclusion. What it increasingly needs to govern is the integrity of the economic objects from which those external models are built. Stable KPI definitions, clear bridges, explicit causal relationships, data lineage, uncertainty ranges, and well-defined strategic drivers all make it less likely that external models begin from avoidable ambiguity. I think that creates a useful conceptual role for IR that I would call representation governance. Management has an internal representation of the enterprise: how it believes revenue is generated, what drives margin, where capital earns its return, what risks matter, and how decisions change when assumptions move. The capital market constructs an external representation of the same enterprise. IR increasingly must understand where those representations diverge, whether the divergence is economically meaningful, and whether better evidence can narrow it. That is much more strategic than ensuring the quarterly message is consistent because it is ultimately about whether the company is being modeled through the right causal architecture.

Jennifer: Could widespread use of similar AI models make markets less intelligent by creating correlated interpretation errors?

Mark: Absolutely, and I think this deserves much more attention. Portfolio diversification does not necessarily imply diversification of interpretive machinery. Ten investment firms may own different portfolios while relying on overlapping foundation models, common data providers, similar retrieval architectures, comparable prompts, and broadly similar classification systems. Their holdings may be diversified while part of their cognitive infrastructure is not. If those systems identify a real structural change, the market may incorporate the information more efficiently. But if they share a flawed classification or an incomplete causal model, the error can become correlated across institutions that appear analytically independent.

That creates a potentially important systemic distinction between diversity of capital and diversity of interpretation. Markets benefit when participants arrive at different conclusions through genuinely independent models of the world. If a large share of the market begins relying on common machine-mediated inference layers, apparent analytical diversity could narrow even while the number of investors remains large. The research already identifies machine-speed capital, trust compression, and disclosure interoperability as emerging market variables. The next question is whether common interpretive infrastructure creates a new form of model concentration. If it does, clean inputs and explicit uncertainty become even more important because an error can propagate across many portfolios before human challenge catches up.

Jennifer: That suggests AI could create path dependence in how a company is perceived. Once a system classifies a business a certain way, could subsequent evidence get interpreted through that classification?

Mark: Yes, and that may be one of the most consequential risks. Every model has priors, whether they are explicit or not. Once a company is classified as structurally capital intensive, cyclically impaired, ex-growth, low quality, or strategically advantaged, new evidence tends to be interpreted relative to that frame. A higher CapEx number looks different if the model already believes the company is entering an unattractive capital cycle. A margin decline looks different if the business has already been classified as losing pricing power. This is not uniquely an AI problem, humans do the same thing, but AI can make the classification more persistent because the model can repeatedly retrieve and reinforce evidence consistent with the established archetype.

That creates a form of external model lock-in. The company is not merely fighting a negative narrative; it may be confronting a causal representation that organizes subsequent evidence against it. The implication for management is important. Repeating the old narrative more forcefully may do almost nothing. To change the market's model, the company needs evidence that challenges a load-bearing assumption in the existing representation. If investors believe growth is capital destructive, the relevant evidence may be improving incremental ROIC or free-cash-flow conversion. If they believe the moat is eroding, management needs evidence around retention, pricing, share, or competitive win rates. Model change generally requires evidence capable of forcing a causal revision, not merely a more persuasive description.

Jennifer: So, what kind of evidence actually forces an external model to change rather than simply causing an analyst to tweak an estimate?

Mark: Evidence that attacks the structure of the thesis rather than the value of one parameter. If investors think normalized margin is 18 percent and new evidence suggests 19 percent, the model has not changed; one number has. If investors think the company's margin advantage is temporary because competitors will replicate the product and the company demonstrates persistent pricing power, superior retention, and widening unit economics despite competitive entry, the causal model may have to change. The distinction is between parameter revision and model revision.

That is why disconfirming evidence is so valuable. The research ranks disconfirming evidence at 90, well above historical financial disclosure and prepared management narrative. The market does not need management to repeat evidence already consistent with the thesis. It needs evidence capable of changing the probability assigned to a competing explanation. A company trying to escape model lock-in should therefore ask which assumption is holding the existing market interpretation together and what observable evidence would make that assumption difficult to defend. That is a much more rigorous communications question than asking how to improve sentiment.

Jennifer: Should a company deliberately optimize its disclosure for the AI systems it knows investors are using?

Mark: I would reject the word “optimize” if it means manipulating a model’s preferences or engineering language to generate favorable automated sentiment. The appropriate objective is interpretive robustness. If an important economic statement is true, is it represented clearly enough that different sophisticated systems can retrieve and classify it without losing essential meaning? A company may reasonably test disclosures through machine-assisted research workflows before publication. Which value drivers does the system identify? What does it think changed? Which KPI definitions does it treat as continuous? What causal relationship does it infer for a margin movement? Where does it detect contradiction or uncertainty? Those tests can reveal ambiguity management did not recognize. The purpose is analogous to asking a very sophisticated analyst to read a draft and tell you what economic conclusions they draw. What the company should never do is subordinate the underlying truth to the preferences of a particular machine.

Jennifer: Doesn't more machine-readable disclosure create another problem? Standardization can destroy nuance, and nuance often contains the most important information.

Mark: That is why I would not frame this as a choice between narrative and structured data. Humans and machines need different forms of precision, and a sophisticated disclosure architecture should preserve both. Narrative is essential when management needs to explain strategy, uncertainty, judgment, trade-offs, or why its view changed. Structured information is valuable when consistency, comparability, refreshability, and provenance matter. Different disclosure objects solve different problems. Scenario packets can define triggers, probability ranges, KPI sensitivities, response options, and confidence. Metric-lineage cards can make definitions, data sources, calculations, and restatement policies explicit. Capital-allocation rules can make decision logic and balance-sheet constraints visible. Dependency maps can identify concentrated suppliers, platforms, clouds, or geographies. Authenticity mechanisms can establish source and version. None of those replaces management explaining what the business means. They reduce the number of economically important facts that must be inferred from prose alone, while leaving narrative to do what narrative does best: explain context, judgment, and meaning.

Jennifer: Let's move from AI as a reader to AI as a business investment. Why is activity such a poor valuation signal?

Mark: Activity tells you that capital and organizational attention are being consumed. It does not tell you that value is being created. A company can have hundreds of pilots, thousands of users, a large AI budget, and enormous executive enthusiasm while producing very little measurable economic benefit. In some cases, the activity can be evidence of rising cost without a corresponding return. The stronger framework moves through investment, adoption, operating change, financial translation, and competitive consequence. What is being spent on data, compute, models, infrastructure, and talent? Is the technology being adopted in workflows or products? What changes operationally: conversion, cycle time, labor productivity, error rates, service cost, product velocity? Where does that show up financially? Finally, does the change create a durable advantage, merely preserve parity, or accelerate commoditization? The research explicitly shifts AI proof toward revenue growth, reinvestment, margin, incremental ROIC, and moat duration rather than pilot counts and generic spending.

Jennifer: What if the economics cannot yet be measured? Some of the most important AI investments may be strategically necessary years before the financial return becomes visible.

Mark: Then management should be precise about what can be known at the current stage rather than manufacturing premature ROI. An experiment does not need a mature ROIC to be rational. It needs a learning objective. What uncertainty is management trying to reduce? What would successful adoption look like? What evidence would justify the next stage of capital? What result would cause management to stop? The burden of proof should rise as the commitment becomes larger and less reversible. The danger is allowing “too early to measure” to become a permanent shelter from accountability. A small experiment can be judged by what it teaches. A scaled workflow should eventually be judged by adoption and operating change. A major infrastructure commitment should be judged by capital requirements, expected returns, and the durability of the advantage it creates. The maturity of the proof should match the maturity and irreversibility of the capital commitment.

Jennifer: Productivity is probably the easiest AI claim to make and one of the hardest to validate. If employees are saving time, why isn't that automatically economic value?

Mark: Because saved time is capacity, not necessarily realized value. If an employee saves eight hours a week but nothing changes in headcount, output, customer experience, innovation velocity, or revenue capacity, the productivity improvement may be operationally real without being financially material. The company has created capacity but has not yet determined how that capacity becomes economic value. The analysis also must capture cost migration. Labor expense may fall while compute, data, infrastructure, cybersecurity, governance, and model costs rise somewhere else. A workflow can become more efficient in isolation while enterprise-level economics remain unattractive. The more rigorous AI productivity discussion therefore asks what capacity was created, what management did with it, what new costs accompanied it, how quickly the organization redesigned the workflow around the technology, and what financial effect remained after those costs. Investors are not trying to deny operating improvement; they are trying to determine whether the improvement deserves to enter the valuation.

Jennifer: How should investors distinguish AI investment that builds a moat from AI spending that is simply the new admission price for staying competitive?

Mark: By asking whether the investment produces differential economics or merely maintains parity. If every competitor must spend roughly the same amount on similar infrastructure to remain credible, then much of the industry's AI investment may become a new cost base rather than a source of excess return. The strategic necessity of the spending does not automatically make it attractive capital. AI becomes moat enhancing when it compounds an advantage that is difficult to replicate: proprietary data, distribution, embedded workflows, switching costs, customer ownership, unique infrastructure, or a faster learning loop. The relevant question is whether each incremental cycle of AI investment causes the company's advantage to widen or merely prevents it from narrowing. The research appropriately places data advantage, switching cost, and speed of learning inside the discussion of AI and moat duration. The investor therefore must separate "AI is strategically necessary" from "AI is economically differentiated," because those can be completely different propositions.

Jennifer: Could AI destroy the very advantages companies are spending money to protect?

Mark: Certainly. A technology can strengthen one layer of a value chain while commoditizing another. If a capability historically required scarce expertise and AI radically lowers the cost of reproducing it, the economic surplus may migrate away from the incumbent. If switching becomes easier because customers can recreate workflows across platforms, customer captivity can weaken. If everyone gains access to similar intelligence, the differentiator may shift toward proprietary data, distribution, trust, brand, or another scarce asset. That is why management should not evaluate AI only as an efficiency tool. It should ask how the technology changes industry structure and where the economic surplus migrates. A company can execute AI exceptionally well and still face a declining moat if the technology makes its historical advantage easier to replicate. Conversely, a company that already owns a scarce complementary asset may see AI increase the value of that asset dramatically. The relevant valuation question is therefore not simply how much AI improves the company, but how AI changes the company relative to the competitive system and whether management has identified the layer of the value chain where scarcity will remain.

Jennifer: You also connect AI governance, authenticity, and provenance to economic value. Isn't there a danger of turning all of that into a generic governance premium?

Mark: There would be if governance were treated as an abstract virtue. It becomes valuation relevant only when there is a credible economic pathway. Weak AI controls can create customer harm, privacy exposure, model failures, regulatory cost, cyber vulnerability, reputational damage, or business interruption. At the same time, governance can become excessively restrictive and impose an innovation penalty. The economic question is whether the control architecture allows the company to scale AI rapidly enough to create value while keeping the probability and severity of failure within acceptable bounds. The broader trust environment is also changing because synthetic voices, documents, identities, and messages make verification a business process rather than simply a cybersecurity control. Misinformation can affect brand, executive credibility, and investor confidence, while boards face expanding responsibility around AI, cyber, privacy, and information provenance. As false information becomes cheaper to create, verified information becomes more valuable. That does not justify a generic governance premium; it means the company's trust architecture can increasingly influence risk, operating continuity, and the credibility of information entering the market.

Jennifer: At some point, does AI arbitrage away the value of traditional IR almost completely?

Mark: AI may arbitrage away the parts of traditional IR that depended on scarcity of access, manual synthesis, routine information delivery, and generic benchmarking. If the value proposition is primarily knowing where to find the numbers, summarizing what management already said, or assembling ordinary peer comparisons, intelligent systems will increasingly do that cheaply. The research already suggests that generic peer benchmarking carries dramatically less informational value than management judgment, changed assumptions, decision signals, disconfirming evidence, and scenario intelligence. But I think that raises rather than eliminates the potential value of a world-class IR function because the enterprise still knows things the market cannot infer confidently: how management interprets operational evidence, what assumptions are changing internally, which strategic trade-offs are being considered, where capital is being reallocated, what has disproved a previous belief, and what leading indicators matter before those changes appear in the financial statements. The low-value layer of IR is vulnerable to automation. The high-value layer becomes more valuable because it sits closer to causal knowledge and decision judgment.

Jennifer: If you had to define the company that ultimately earns the highest level of investor trust in a machine-mediated market, what would distinguish it?

Mark: It would recognize that the scarce capital markets asset is no longer information volume but the quality of the system through which information becomes belief. Economically important claims would have visible evidence, causal relationships would be explicit enough to survive compression, important metrics would have stable lineage, uncertainty would be represented rather than cosmetically removed, and management would continuously test the relationship between its internal model of the enterprise and the model the market is building outside it. The objective would not be to engineer agreement. Investors should disagree about intrinsic value, probabilities, duration, and risk. The objective would be to make sure the disagreement concerns those genuinely uncertain economic questions rather than confusion over what the company's own information means.

The strongest company would also understand model lock-in. It would know that once the market has classified the business through a particular causal lens, rhetoric alone rarely changes the classification. Management would identify which assumptions are carrying the external model and produce evidence capable of challenging them. It would understand that in a world of overlapping machine infrastructure, correlated interpretation can arise even among nominally independent investors, making clarity and disconfirming evidence more consequential. It would not try to game machines, but it would design information for interpretive robustness so that the economics remain intelligible after being retrieved, decomposed, summarized, compared, and recombined.

CONTINUED

Internally, the company would refuse to let AI become a language strategy detached from capital discipline. Management would know which investments are experiments, which have achieved adoption, which have changed operating economics, and which deserve scaled capital. It would distinguish productivity capacity from realized financial value, necessary defensive spending from differentiated advantage, and technological enthusiasm from measurable return. What AI ultimately changes is not the need for human judgment but the location of its scarcity. Retrieval becomes cheaper, comparison becomes cheaper, and summarization becomes cheaper, while the ability to determine what matters, recognize when the existing model is wrong, change a belief when evidence demands it, and make a consequential decision under uncertainty becomes more valuable. The company that earns trust in that market will be the one whose economic reality is sufficiently coherent that it remains intelligible even after machines have taken the information apart and reconstructed it through thousands of independent, or seemingly independent, investment processes.