



AI and the Fiduciary Test

**A Guide for Institutional
Investors in Evaluating
AI Proxy Voting Solutions**

Arik Brutian, Senior Vice President, Artificial Intelligence & Data
Kaveh Beigi, Vice President, Digital & Content Strategy

April 2026

Contents

Executive Summary	3
Five Key Takeaways.....	5
Section 1: AI Is Reshaping Proxy Voting	6
Section 2: Why Proxy Voting Makes Unique Demands on Any AI System	8
Section 3: Two Approaches to AI in Proxy Voting	13
Section 4: Five Questions Every Asset Manager Should Ask	18
Section 5: The Glass Lewis Approach	24
What's Next	28
Appendix: Key Terms	29
Notes and references.....	33

Executive Summary

A Guide for Institutional Investors in Evaluating AI Proxy Voting Solutions

AI tools are delivering real benefits to proxy voting workflows.

The ability to process disclosures at scale, apply voting policies consistently across large portfolios, and surface governance risks earlier in the research cycle represents genuine progress for institutional investment stewardship. These capabilities are now widely available, and most major providers in this space have deployed some form of AI-assisted workflow.

The question for institutional asset managers is not whether to use AI in governance: it is what standard governs the AI applied to fiduciary responsibilities, and whether the systems being offered are actually built to meet that standard.

Setting a high bar takes expertise and investment.

Glass Lewis has spent more than 23 years developing the institutional expertise, analyst depth, and governance methodology that defines what investment-grade proxy research requires. That expertise, in which data points are material across more than 100 regulatory regimes, how governance concepts translate across markets, where the non-routine situations are and why they matter, is the foundation on which the firm's AI strategy is built. Sophisticated AI pipelines are increasingly available across the governance industry.

What determines their fitness for fiduciary purposes is the depth of institutional knowledge governing what those pipelines produce, and whether that knowledge is sufficient to support defensible fiduciary decisions.

 **The question for institutional asset managers is not whether to use AI in governance: it is what standard governs the AI applied to fiduciary responsibilities, and whether the systems being considered are built to meet that standard.**

This paper makes three arguments.

First, it argues that proxy voting has structural properties, including fiduciary non-delegation, cross-jurisdictional data complexity, and the frequency of genuinely non-routine decisions, that make the quality of the underlying architecture a fiduciary matter, not merely a technical one.

Second, there are two fundamentally different approaches to AI in proxy voting, and the difference between them is material: one embeds human expertise, judgment, and oversight in the production process, the other applies human review to outputs after the fact.

Third, there are five specific questions that any asset manager can use to distinguish between the two approaches and determine whether any AI governance solution is built for the accountability context that fiduciary obligations create.


Five Key Takeaways

1. The critical variable in any AI proxy voting solution is not the sophistication of the model. It is the quality of the institutional expertise, data governance architecture, and production accountability underneath it. An AI system can only produce outputs as defensible as the foundation it is built on.
2. There is a meaningful operational difference between AI systems that embed human expertise in the production process and systems that apply human review to AI-generated outputs after the fact. That difference determines whether the human is exercising governance or performing quality control, and it is the difference that matters in a fiduciary context.
3. Fiduciary accountability in proxy voting cannot be delegated to an AI system. The regulatory framework is increasingly explicit on this point. The EU AI Act's human oversight requirements take effect in August 2026. Stewardship codes in the UK, EU, and Japan all require that voting decisions be explainable and defensible at the process level, not just the output level. The regulatory direction across every major institutional market is toward a higher standard of accountability.
4. Investment-grade governance data is the product of governed architecture: data models defined by domain experts, normalization rules applied consistently across regulatory regimes, and traceability maintained from source document to final output. Not all AI extraction approaches are designed to this standard, and the difference is not always visible at the recommendation level. It becomes apparent when a specific voting decision is scrutinized.
5. Institutional investors should assess five important aspects of any AI proxy voting solution: the data governance architecture, the nature of human production roles, investment-grade data standards, the AI system's designed scope, and exception-handling accountability.

Section 1: AI Is Reshaping Proxy Voting

What That Actually Means

AI is delivering real, measurable benefits to governance research and proxy voting workflows. Machine learning and large language models can surface relevant data from regulatory filings in minutes, rather than hours. Automation tools can apply voting policies consistently across thousands of ballots, reducing the operational risk of human inconsistency at scale. AI-assisted analysis can easily identify patterns in governance data, board composition trends, compensation structures, and environmental disclosure trajectories, that would require significant effort on the part of analyst teams to detect through traditional research approaches.

 **AI is delivering real, measurable benefits to governance research and proxy voting workflows.**

These are genuine AI capabilities, and they are changing what is possible in investment stewardship. Glass Lewis's new Climate Intelligence platform, for example, uses AI to evaluate how companies are managing climate transition risks and opportunities by assessing both their strategy and their execution. The result is a forward-looking view that stewardship teams can use to prioritize engagement, identify advanced and emerging companies within a sector, and ground their voting and engagement decisions in something more analytically rigorous than disclosure language alone. The question facing institutional investors is not whether to use AI in governance workflows. It is how to evaluate the quality and fitness of the specific systems being considered.

In 2026, that evaluation question has become more urgent. A few of the world's largest asset managers have announced that they are building or deploying internal AI voting tools, some explicitly to reduce reliance on external advisory services. New platforms have

entered the market promising AI-powered governance research at lower cost and higher speed. Every major provider in the proxy voting and governance research space now describes its approach using some variation of human oversight language.

The question before asset managers is no longer whether AI belongs in governance workflows, but what standard should govern the AI used to carry out fiduciary responsibilities, and whether the solutions being offered are actually built to meet that standard. Not all human oversight is equivalent. Not all AI data architectures are investment-grade. The evaluation criteria in what follows are designed to make those distinctions more clear.

Glass Lewis's own direction reflects this reality. Beginning in 2027, Glass Lewis will no longer issue a standard benchmark recommendation, moving instead to AI-powered, client-specific perspectives that reflect an organization's specific investment philosophies rather than a single advisor-determined house view. That transition is underway. It is accompanied by deliberate investment in the AI infrastructure, governed data architecture, and analytical capability required to ensure that every client-specific recommendation is analytically rigorous and traceable, so that when clients defend their voting decisions to their own stakeholders and to regulators, the research behind those decisions can withstand scrutiny. This paper describes the standard that investment demands.

Section 2: Why Proxy Voting Makes Unique Demands on Any AI System

Proxy voting sits at the intersection of fiduciary duty, cross-jurisdictional regulation, and corporate accountability. It is among the most demanding AI applications in institutional investment, and understanding what makes it distinct is essential to evaluating any AI system designed to support it.

Fiduciary Accountability Cannot Be Delegated

What was once seen as merely a compliance exercise, proxy voting is now an important expression of investment stewardship and fiduciary duty. When an asset manager casts a vote at a shareholder meeting, that vote is an act of fiduciary duty that can be scrutinized by regulators, clients, and courts.

The regulatory framework is clear across jurisdictions. Institutional investors are required to maintain documented rationales for voting decisions and to anchor those decisions to objective factors that can withstand regulatory and client scrutiny. The UK Stewardship Code 2020 requires signatories to explain their voting decisions for significant votes and to demonstrate that voting policy has been implemented with genuine judgment.¹ The revised Japan Stewardship Code, updated in June 2025, explicitly states that voting policy must not be a mechanical checklist and expects case-by-case assessment of individual proposals.² SRD II, in force across EU member states, requires proxy advisors to disclose the essential features of their methodologies and models, their primary data sources, and the procedures ensuring research quality, including the qualifications of the staff responsible for it.³

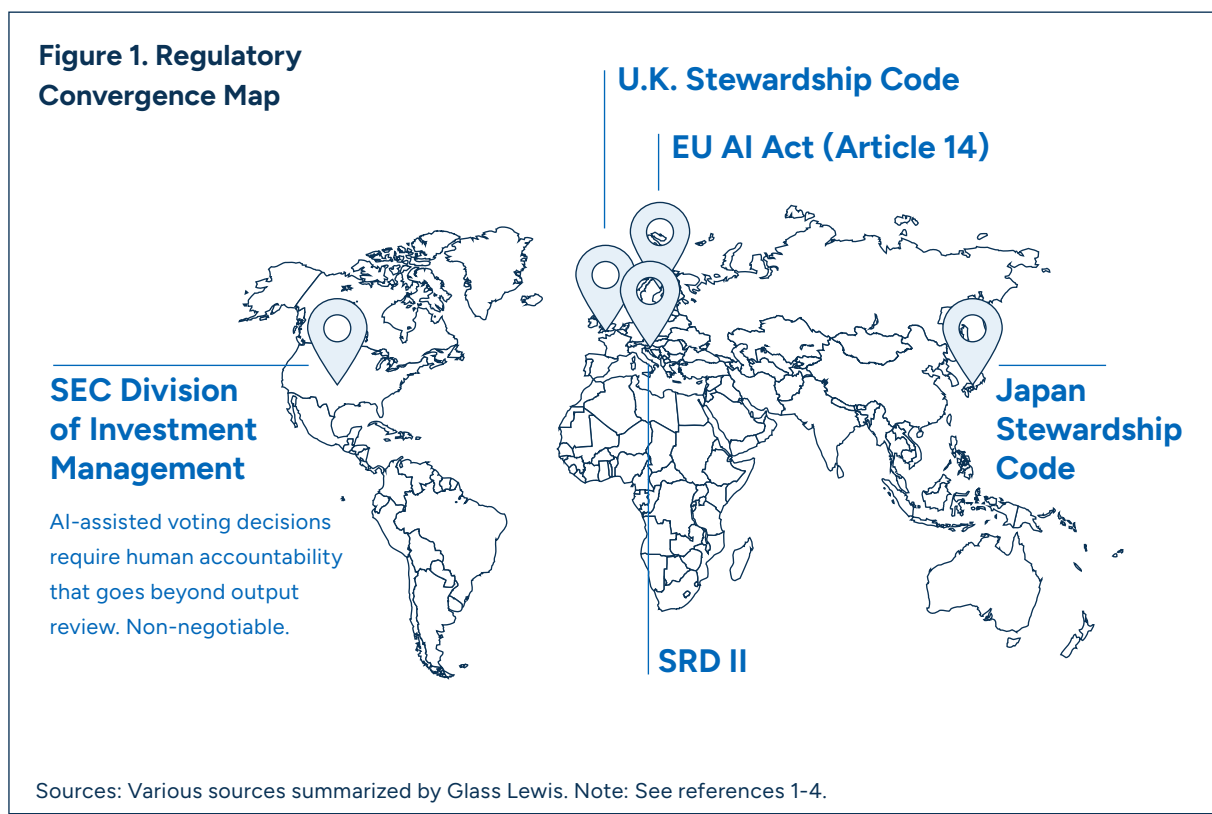


The common thread across these frameworks is accountability at the process level.

The common thread across these frameworks is accountability at the process level, not just the output level (see Figure 1). An AI system that generates recommendations and has a human approve them is not the same as a system where human expertise governs the process that produces the recommendation. The first system creates a record of approval. The second system creates a record of judgment.

In the United States, the Securities and Exchange Commission’s Division of Investment Management has acknowledged that AI agents reviewing proxy materials against an adviser’s expressed recommendations represent a near-term reality, while emphasizing that outputs must be reviewed for fiduciary consistency and that human accountability for the voting decision remains non-negotiable.⁴ Across jurisdictions, the regulatory direction is consistent.

The EU AI Act, whose high-risk provisions take effect in August 2026, codifies this standard for European markets.⁵ Article 14 requires that AI systems used in high-risk contexts be designed so that human oversight is built into the architecture, not documented as a policy overlay. The persons overseeing the system must be able to understand its capabilities and limitations, recognize when outputs are likely unreliable, and override them. The regulation’s implication for the use of AI in proxy voting is this: a provider cannot say the AI determined the board was insufficiently independent, the specific logic must be documented, and a human must be able to override it before the vote is cast.



Summary of EU AI Act, Article 14 on Human Oversight:

High-risk AI systems must be designed so that the persons assigned to oversight can understand the system's capacities and limitations, remain aware of automation bias, correctly interpret outputs, and decide, in any **particular situation**, not to use the system or to override its output. Oversight must be by design, not by documentation. Effective date for Annex III high-risk systems: August 2, 2026.⁶

Investment-Grade Data Requires Governed Architecture

Even a well-designed AI system can only produce outputs as reliable as the data foundation it operates on. The data environment that supports enterprise-level global proxy voting is extraordinarily complex. Governance information must be collected across more than 40 languages, interpreted within the context of more than 100 distinct regulatory regimes, normalized across different disclosure frameworks, and linked to stable corporate identifiers that account for mergers, spin-offs, and structural changes over time.

 **Governance information must be collected across more than 40 languages, interpreted within the context of more than 100 distinct regulatory regimes, normalized across different disclosure frameworks, and linked to stable corporate identifiers that account for mergers, spin-offs, and structural changes over time.**

Getting this right is fundamentally a data governance issue.

The specific challenge is normalization. For example, a board independence assessment means different things in different markets. A non-executive director in the UK is not the same as an outside director in Japan or an independent director under U.S. securities rules. Another issue is that executive compensation benchmarking requires different peer group methodologies in different regulatory environments. Or, environmental data that conflates company-reported figures with model-estimated figures produces inconsistent comparisons that are invisible at the recommendation level but surface when a specific vote is challenged and the underlying data is examined.

A governed normalization layer is what allows AI to produce outputs that remain internally consistent across markets and what makes the underlying data defensible when a specific voting decision is scrutinized. Without that layer embedded in the architecture by design, inconsistencies accumulate in the data itself, where they are difficult to detect at the recommendation level, but become material when the provenance of the data is examined.

Investment-grade governance data requires six properties that are distinct from what many AI extraction pipelines are designed to produce: 1) accuracy validated against source evidence, not inferred from model confidence; 2) consistency maintained through deterministic normalization rules, not probabilistic interpretation; 3) completeness verified against known coverage gaps, not assumed from extraction completeness; 4) traceability preserved through an unbroken chain from source document to final output; 5) validity, meaning the degree to which governance data conforms to specified formats, value ranges, and logical reporting rules; and 6) timeliness, meaning the degree to which data is available for analysis within the specific timeframe required for the research to remain actionable.

Non-Routine Situations Require Institutional Judgment

The most consequential proxy voting decisions are rarely the straightforward type. Consider what genuinely demanding governance situations look like in practice:

1. A contested board election at a company where the incumbent directors and the activist slate both have credible governance arguments, and the decisive factor is the board's track record on a specific strategic question that requires historical context to evaluate.
2. A shareholder proposal on an emerging systemic risk, such as an AI governance framework, where there is no established market precedent and evaluating the proposal requires forward-looking judgment about which oversight structures are substantively meaningful, versus those that are cosmetic.
3. A say-on-pay vote at a company that is simultaneously undergoing a strategic restructuring and facing a governance controversy, where the compensation outcome must be assessed in the context of both financial circumstances and board conduct during the controversy.
4. A contested acquisition vote where the board has completed a go-shop process and assessing the situation requires asking not whether the formal steps were followed, but whether the structural mechanics of that process gave potential bidders a genuine opportunity to compete, or effectively preserved the incumbent acquirer's position.

Go-Shop Process Integrity: A Recent Example⁷

A medical device company's board sought shareholder approval for an all-cash sale to a strategic acquirer. Investors holding more than a third of the company's shares had publicly opposed the deal, and the original vote was adjourned at the eleventh hour when defeat appeared imminent. The board subsequently initiated a go-shop process and returned to shareholders with a revised offer at a roughly 10% premium, characterizing the outcome – 21 outbound contacts, two NDAs, and no competing offers – as validation of the original process and evidence that the deal represented the best available value.

Glass Lewis reached a different conclusion. The go-shop's mechanics required the company to provide the acquirer with unredacted copies of any competing proposal within one business day of the period's close, and gave the acquirer a four-day window to respond before the company could accept a superior offer. Any party considering participation would have known that its terms would immediately be visible to the incumbent bidder. The formal attributes of a competitive process were present. The functional reality of one was not.

Reaching that conclusion required something an AI extraction pipeline cannot provide: accumulated knowledge of how acquirers and potential bidders actually behave in go-shop contexts, and how structural mechanics that appear neutral on paper can systematically favor a preferred bidder. Glass Lewis maintained its recommendation against the transaction through multiple amendments and a final price increase.

All of the examples illustrated above share a common feature: they require nuance and specific consideration. They require the kind of judgment that comes from years of covering governance events when they are contested, economically complex, or without established precedent in a specific market – observing how companies respond to shareholder engagement, and building a working knowledge of what good governance looks like in practice.

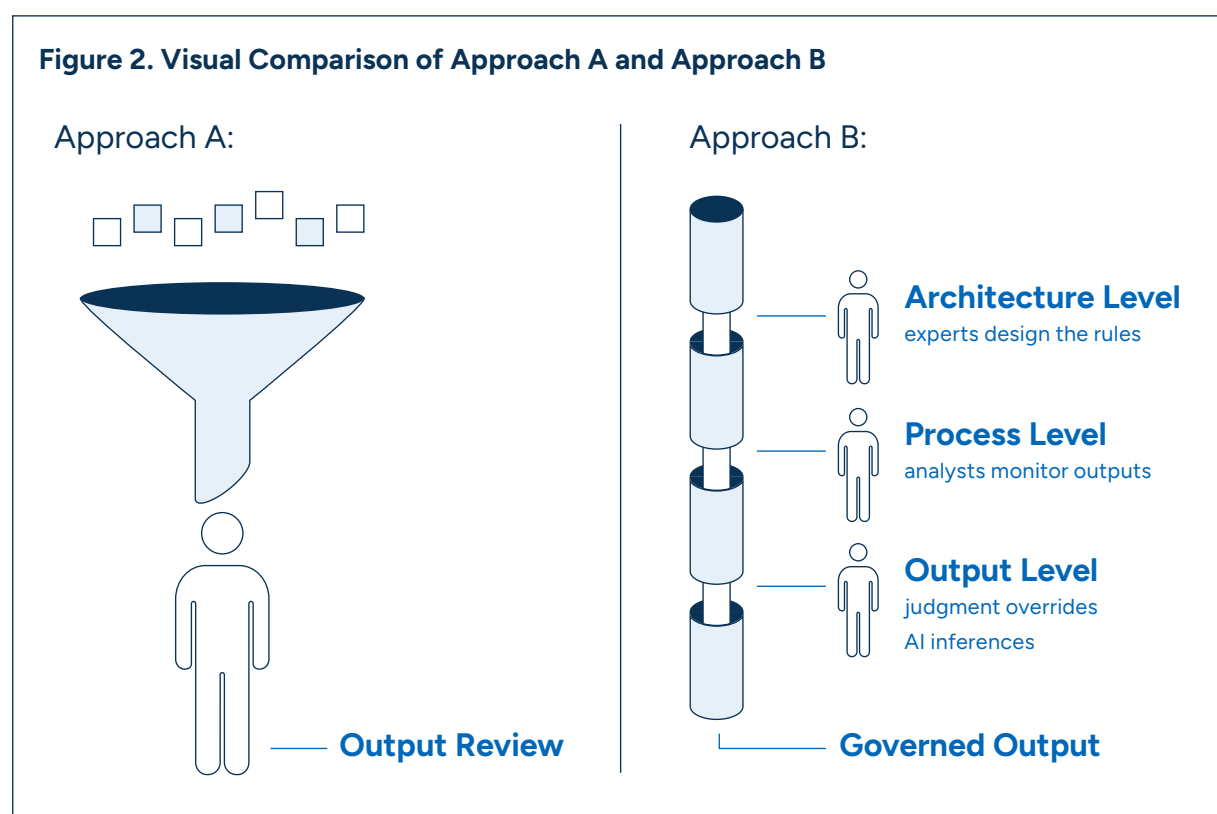
AI systems built on and informed by historical governance data and institutional knowledge have real analytical value in these situations. They can surface relevant precedents, synthesize disclosure patterns, and identify structural similarities to past events. What they cannot replicate is the institutional perspective that evaluates a specific situation on its own terms, in the context of everything else that is known about the company, the board, and the market environment.

The value of human expertise in this context is not the review it provides at the output stage. It is the accumulated knowledge that shapes the methodology governing what the AI is permitted to conclude, and the judgment that guides the AI's reasoning when a situation requires more than pattern recognition. That is the institutional expertise dimension of the proxy voting AI challenge, and it is the dimension that takes the most intensive effort to build.

Section 3: Two Approaches to AI in Proxy Voting


The phrase human-in-the-loop (HITL) is now used across the governance and stewardship industry to describe AI systems that look quite different from one another in practice. The term as commonly used encompasses a wide range of implementations, from systems where human expertise governs the production process to systems where human review is applied to AI outputs after the fact. Understanding the operational difference between these implementations is the most important analytical task for any asset manager evaluating AI proxy voting solutions. Human-Centric AI (HCAI)⁸, as Glass Lewis defines and practices it, is something more specific: a framework in which the institutional knowledge of governance professionals shapes, governs, and checks the methodology the AI follows, not just the outputs it produces.

Figure 2. Visual Comparison of Approach A and Approach B



Approach A: The AI-First Pipeline

In the most common architecture entering the market, AI handles the production of governance research and voting recommendations from source data through to output (Figure 2). Data is extracted from regulatory filings using machine learning. Policy rules are applied algorithmically to generate voting instructions. Human analysts review outputs to check for obvious errors, flag edge cases, and confirm that recommendations align with the client's stated policy.

 **A human reviewing AI-generated recommendations at volume is exercising a different kind of judgment than an analyst whose expertise shaped the process that produced those recommendations.**

This architecture has genuine capability. It can process thousands of shareholder meetings faster than manual research allows. It can apply policy consistently across large portfolios. It can identify potentially contentious votes before they require deeper analysis. For asset managers focused primarily on operational efficiency in routine voting situations, these are meaningful benefits.

The limitation is structural, not technical. In an Approach A architecture, humans are positioned at the output stage. Their role is to evaluate what the AI has produced. The methodology governing what the AI is permitted to conclude was designed as a largely technical exercise: schema choices, extraction rules, model training decisions. The data feeding the model reflects what was extractable from source documents, not necessarily what a governance expert would specify as the relevant inputs for a given decision type.


This matters in practice. The EU AI Act's explicit concern about automation bias, the documented tendency of human reviewers to defer to AI outputs even when those outputs are incorrect, is a known risk that applies wherever AI-assisted decisions are made under time pressure. In proxy voting, that pressure is acute, with thousands of ballots requiring attention in compressed windows during a proxy season. Regulators across the EU, UK, U.S., and Japan are converging on the same expectation: that human oversight must be meaningful and substantive, not limited to the output stage. A human reviewing AI-generated recommendations at volume is exercising a different kind of judgment than an analyst whose expertise shaped the process that produced those recommendations. The first is checking. The second is governing.

Approach B: The Expertise-Governed Architecture

In this more rigorous architecture, AI operates inside a methodology framework built by governance domain experts before the AI was applied. The data models defining which data points are relevant, how they are normalized across markets, and how they map to specific governance scenarios were designed by analysts who have spent years covering these markets. The rules governing what the AI is permitted to conclude in a given situation were developed from accumulated institutional knowledge by experts and methodologists and embedded into the AI architecture by design. Human review is embedded in the AI-powered production process at defined decision checkpoints, not appended to the output.

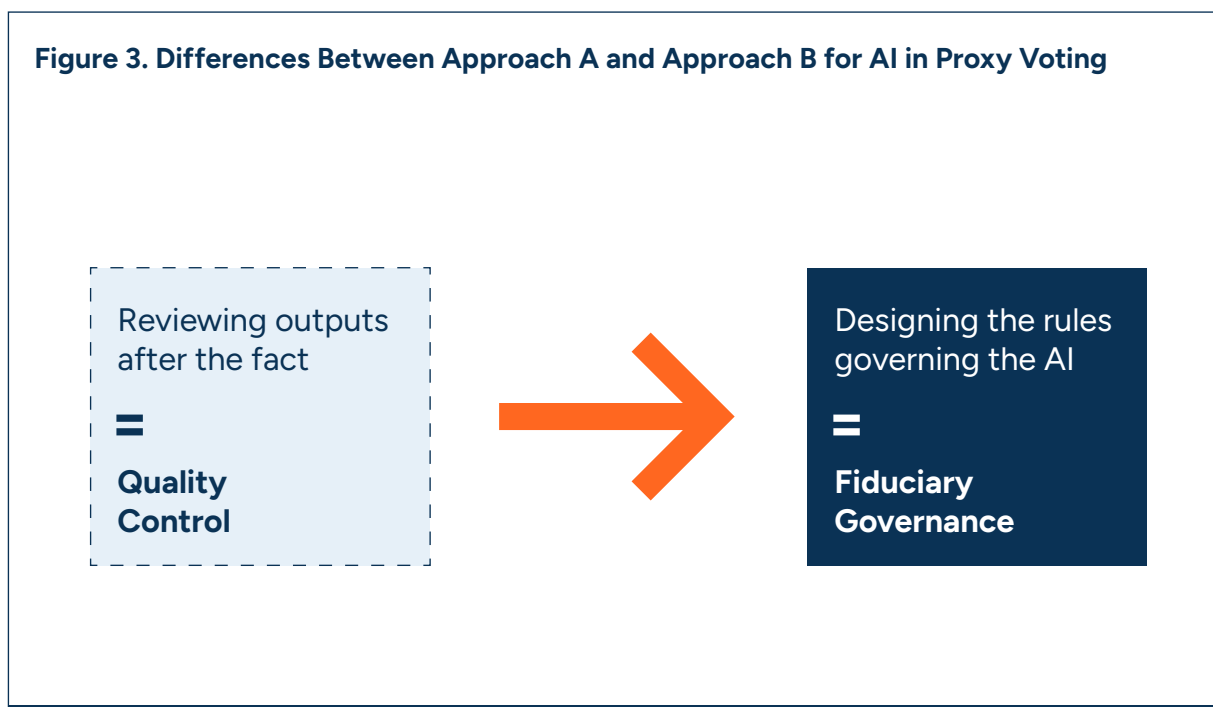
The difference is significant. In an Approach B architecture, the analyst's expertise governs the system's behavior before a single output is generated. In Approach A, it evaluates outputs after they have been produced. The first creates a system whose outputs reflect institutional knowledge by design. The second creates a system whose outputs are filtered through it (Figure 3).

Consider a board oversight assessment at a company where the relevant governance information is scattered across the annual report, committee charters, and investor day presentations, none of it structured in a way that permits a direct algorithmic conclusion. These situations are the norm, not the exception. Corporate disclosure rarely arrives in a form that allows reliable inference without contextual interpretation.

 **Glass Lewis has built an AI architecture governed by expert-defined rules that operationalize the firm's methodology into a structured reasoning framework. The system reads the relevant disclosure, extracts the evidence, applies inferential logic, and produces an assessment with explicit justification for each conclusion.**

For situations like this, Glass Lewis has built an AI architecture governed by expert-defined rules that operationalize the firm's methodology into a structured reasoning framework. The system reads the relevant disclosure, extracts the evidence, applies inferential logic, and produces an assessment with explicit justification for each conclusion. At every stage, the analyst has full visibility into which documents the AI selected, which passages it used to support its inferences, and the reasoning path it followed to reach its conclusion. The analyst can modify the governing ruleset, calibrate the system's logic, and override or edit any finding. Every inference is traceable to the specific passage in the source disclosure that generated it.

Figure 3. Differences Between Approach A and Approach B for AI in Proxy Voting



This approach reflects a broader principle in how Glass Lewis structures human involvement across its AI production environment. The architecture embeds oversight at three levels: at the individual output level, where high-stakes findings require analyst review before reaching clients; at the process level, where analysts monitor streams of outputs for anomalies and recalibrate the system when patterns fall outside expected parameters; and at the architecture level, where methodologists assess whether a given AI approach is fit for its purpose and retain the authority to modify or halt it if it is not. Each level serves a distinct governance function. Together they ensure that human expertise governs not just what the AI concludes, but how and whether it operates.

There is a practical test that reveals the difference in approaches. Ask any proxy voting AI provider to describe, for a specific governance scenario, what their analyst is actually doing in the production process. **For a contested board election at a non-U.S. company, at what point does a human engage? What methodology framework are they checking against? What source evidence are they expected to verify? What happens if their judgment differs from the AI output? What documentation is produced?** A provider with a genuine investment grade Approach B architecture will answer that question with specificity. A provider applying human oversight as a label over an Approach A architecture will describe a review process.

This distinction has direct regulatory alignment across major markets. In Europe, the EU AI Act's Article 14 requirement that oversight be by design describes an Approach B architecture.⁹ ISO/IEC 42001, the international AI Management Systems Standard, requires that responsible AI deployment include upfront decisions on which stages of AI production human oversight is to be incorporated.¹⁰ In the UK, the Financial Conduct Authority (FCA) places accountability on senior managers to demonstrate appropriate governance of AI-enabled processes.¹¹ In Japan, the Stewardship Code explicitly rejects mechanical checklists

as a substitute for genuine judgment.¹² In the United States, the SEC has emphasized that AI-assisted voting decisions require human accountability that goes beyond output review.¹³ The U.S. National Institute of Standards and Technology (NIST) AI Risk Management Framework requires human oversight over AI systems and places responsibility for AI governance in organizations with fiduciary duties on senior leadership.¹⁴ Independently and across jurisdictions, regulators are converging on the same standard.

Table 1. Comparison of Approach A and Approach B by Dimension

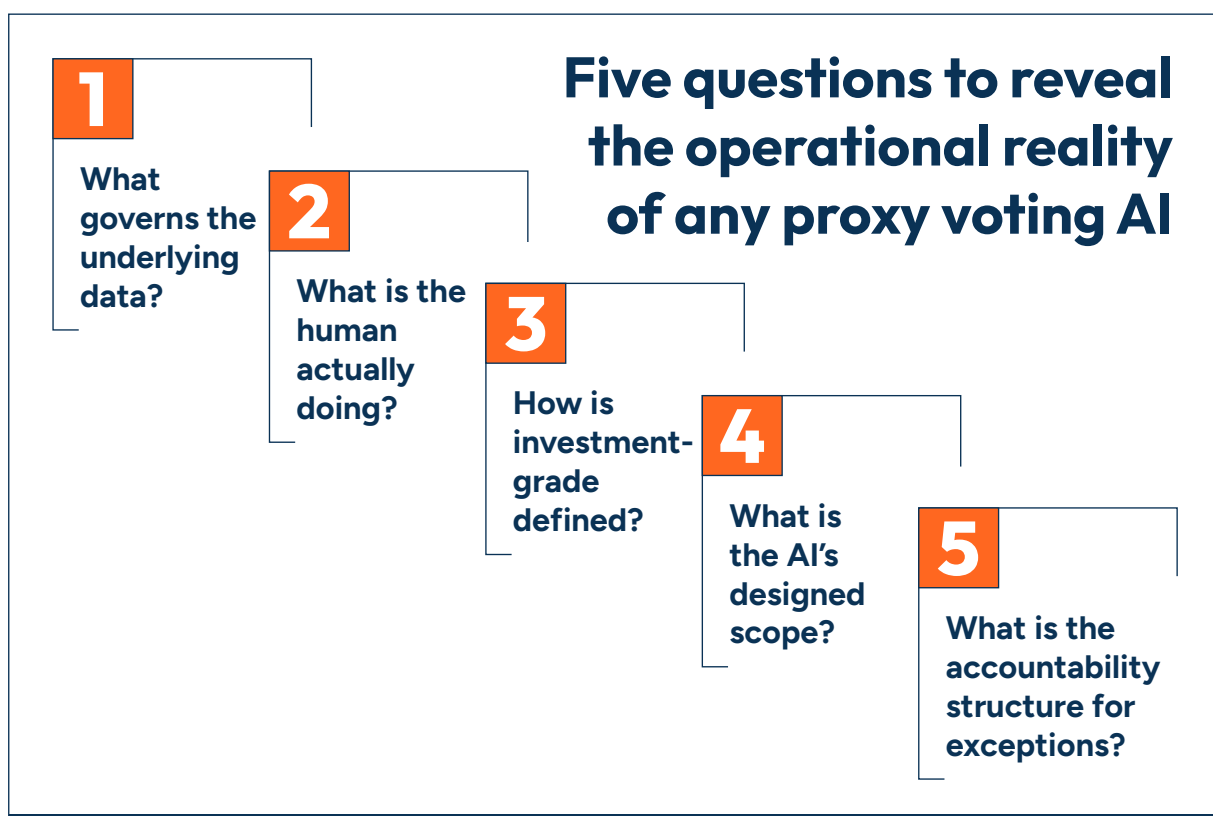
DIMENSION	APPROACH A: AI-FIRST PIPELINE	APPROACH B: EXPERTISE-GOVERNED ARCHITECTURE
How AI operates	From data extraction through to recommendation output	A methodology framework designed by governance domain experts governs the AI production process – from document acquisition to recommendation output
Where humans engage	At the output review stage	At multiple defined checkpoints throughout the production process
What humans are doing	Reviewing AI-generated recommendations for errors and policy alignment	Exercising decisive judgment against a methodology framework, source evidence, and AI reasoning
How the methodology was built	Designed primarily as a technical exercise	Designed by governance analysts whose institutional knowledge shaped the rules and then operationalized as part of the AI system
Documentation produced	Output approval record	Governance checkpoint record with source traceability and rationale for each inference
Regulatory alignment	Human oversight applied to outputs	Human oversight by design, built into the architecture at all required decision points
Primary vulnerability	Automation bias at output review volume	Requires upfront investment in analyst expertise, methodology development, and AI systems design

In practice, AI systems often blend elements of both approaches, and the boundary is not always visible from the outside. The decisive question is not what category a provider claims. It is where judgment actually governs the process. That is what the five questions in Section 4 are designed to reveal.

The choice between these approaches is ultimately a decision about where accountability for fiduciary outcomes lives. For asset managers whose voting decisions carry legal and reputational accountability, that decision matters.

Section 4: Five Questions Every Asset Manager Should Ask

Evaluating any AI proxy voting solution requires understanding the full proxy voting spectrum – from governance research and policy development, through recommendation, client customization, and vote execution. Only a very small number of proxy voting providers have operated across that entire chain at global scale and at an enterprise, investment grade level. That full-spectrum experience is what determines whether the AI's designed scope is adequate for the accountability context that fiduciary obligations create.



The two approaches described in Section 3 produce outputs that can look similar on paper. Both generate voting recommendations. Both describe human involvement. Both claim data quality. The difference is in the overall vision, design, and architectural choices that produced them.

The questions that follow are designed to make the architecture visible. Each one targets a specific layer of the AI governance stack: the data foundation, the human production role, the quality standard, the designed scope of the AI, and the accountability structure for situations the system was not designed for. Together, they form a framework for evaluating the fitness of any AI proxy voting solution for the accountability context that fiduciary obligations create.

The questions are designed to move the conversation from philosophy to operations. They require a description of how the system actually works. The specificity of the answers will reveal the operational reality more clearly than any general description of the provider's approach.

Question 1: What Governs the Data the AI Is Operating On?

The output quality of any AI system is bounded by the quality of its inputs. In governance research and proxy voting, that means the data layer is the most consequential architectural choice a provider makes. Yet it is also the least visible to buyers.

The critical distinction is between governed data and extracted data. Governed data is produced under explicit, human-designed schemas: data points are defined before extraction begins, normalization rules specify how comparable concepts are treated across different regulatory regimes, and quality checks validate that the output meets a specified standard. Extracted data is produced by AI models that infer structure from source documents. It is faster to produce and cheaper to scale. It is also more likely to embed inconsistencies that are invisible at the output level but become liabilities when a specific vote is scrutinized.

The more revealing question is not whether a provider's data is accurate, but how accuracy is defined and enforced. Who designed the schema? How are governance concepts normalized across markets where they mean materially different things? How is source traceability maintained from disclosure document to final data point? What happens when the source data is ambiguous, incomplete, or structured differently from what the model expects?

SRD II, in force across EU member states, requires proxy advisors to disclose the primary information sources used and the procedures ensuring research quality, including staff qualifications. Glass Lewis publishes detailed methodology documentation, provides clients with details about the governance experts responsible for each research domain so they can conduct due diligence, and maintains source traceability from disclosure document to final data point, giving clients the specificity they need to satisfy their own regulatory obligations.¹⁵


Question 2: What Is the Human in Your Production Process Actually Doing?

This is the question that distinguishes Approach A from Approach B, and the answer will be critical to decision-making.

The question is a test for specificity. For a contested board election at a company in a non-U.S. market, describe what happens in your production process. At what point does an analyst engage? What methodology framework are they checking against? What source evidence are they expected to verify? What does the review generate in terms of documentation? What happens when the analyst's judgment differs from the AI recommendation?

A provider with governance embedded in their production process will answer these questions with specificity: names, decision types, methodology documents, and escalation procedures. The nature of the answer, whether it describes a governance process with defined checkpoints or an output review process, is itself informative.

The concern about automation bias is directly relevant here. The documented tendency of human reviewers to defer to AI outputs at volume is a known risk, and it is particularly acute in time-pressured environments like a proxy season, where thousands of ballots require attention in compressed windows. An oversight process that relies on humans catching errors in AI outputs is structurally different from a process where human expertise governs what the AI is permitted to do in the first place. The first process is a control. The second process is an architecture, with quality checks embedded proactively.

 **An oversight process that relies on humans catching errors in AI outputs is structurally different from a process where human expertise governs what the AI is permitted to do in the first place.**

Question 3: How Do You Define Investment-Grade, and What Specifically Does Your Data Meet?

Investment-grade governance data has become a common claim. It needs to be tested with a specific definition, not accepted as a credential.



Accurate



Consistent



Complete



Traceable



Valid



Timely

Functional Criteria for Buyers: Investment-Grade Governance Data is:

Accurate: Validated against source evidence rather than inferred from model confidence.

Consistent: Maintained through deterministic normalization rules rather than probabilistic interpretation across regulatory regimes.

Complete: With coverage gaps identified and documented rather than assumed from extraction completeness.

Traceable: With an unbroken, auditable chain from source document to final data point.

Valid: Meaning the degree to which governance data – such as CIK numbers or ownership percentages – conforms to specified formats, value ranges, and logical reporting rules.

Timely: Meaning the degree to which data – such as proxy filings or board changes – is available for analysis within the specific timeframe required to ensure the research remains relevant and actionable.

Ask the provider to describe how their data meets each of these six criteria specifically. What accuracy validation processes exist, and what standard do they test against? How is normalization governed across markets where the same governance concept has different regulatory definitions? What is the provider's documented approach to coverage gaps? How is traceability maintained at the individual data-point level? What checks are in place to ensure that data is valid and up to date?

Confidence scores and model accuracy rates measure an important dimension of AI performance, but they are an incomplete answer to these questions. They capture how often the AI appears to be right, not how the underlying data is governed. Investment-grade governance data requires both accurate outputs and an auditable architecture that produces them, so that when a specific voting decision is scrutinized, the reasoning can be reconstructed from source evidence. That is the standard Glass Lewis is building towards.

Question 4: What Is the Designed Scope of Your AI, and What Lies Outside It?

This question reframes the standard coverage conversation in a way that is specific to AI. Traditional governance research firms have always competed on breadth, measuring their universe by company count and market coverage. That is a useful data point. But for AI-powered systems, the more revealing question is about designed scope: what was the system built to handle, and what was not anticipated in its design?

Global institutional investors vote across markets with materially different governance norms, disclosure frameworks, and regulatory regimes. Building AI for that environment is a different engineering problem than building for a single regulatory context. And it requires methodology frameworks developed with the specific governance expertise of each market. Glass Lewis's AI architecture was designed from the outset for this complexity — the data models, normalization rules, and methodology frameworks reflect more than two decades of operational experience across more than 100 regulatory regimes and 40 languages.

The more meaningful measure is design intent: whether the AI architecture, the data models, the normalization rules, and the methodology frameworks were built for each market from the outset, with the governance expertise specific to that regulatory environment shaping the design. An architecture developed for one regulatory context and extended to cover others carries the assumptions of its original context, where those assumptions may not hold. Building for global complexity from the start is what allows AI outputs to reflect the analytical framework appropriate to the specific market in question.

 **Investment-grade governance data is a product of governed architecture.**

Ask the provider: for the markets that represent your clients' non-U.S. voting obligations, how were the data models and methodology frameworks for those markets developed? Who designed the normalization rules? When did coverage of those markets begin, and what governance expertise informed the design? The answers will indicate whether the system was built for global institutional voting or adapted for it.

A related question is whether the provider has governance research expertise physically present in the markets they cover. Global proxy voting requires local knowledge, not just data coverage, but analysts who understand the specific governance norms, disclosure practices, and regulatory context of each market. That presence cannot be replicated through data partnerships or technology integrations. It is built through years of operational commitment to those markets.

Question 5: What Is Your Accountability Structure When the System Reaches Its Limits?

Every AI system has a designed scope, and every proxy season will present situations that fall outside of it. Contested elections at companies with unusual ownership structures. Shareholder proposals on topics the training data did not anticipate. Governance events in markets where the system's coverage is thinner. What happens in these situations reveals whether the AI is designed for fiduciary accountability or for operational efficiency.

A system designed for fiduciary accountability has a defined escalation pathway for non-routine situations: a process for identifying when a situation has exceeded the AI's designed scope, a clear route to the governance expertise that applies to that situation, a documentation standard for what was considered and why, and a named accountability structure for the outcome. A system designed primarily for operational efficiency will have a generic fallback for exceptions, without the specificity that distinguishes genuine governance from output validation.

A strong answer will describe non-routine situations as designed for, rather than as an exception. In a globally diversified institutional portfolio, contested elections, novel proposals, and genuinely complex governance events occur multiple times per proxy season.

A system that treats them as exceptions rather than designed-for situations is a system calibrated for volume. The accountability structure around those situations is where the fiduciary rubber meets the road.

Ask the provider to walk through a specific non-routine scenario and describe exactly what their process looks like. The specificity of the answer, who is involved, what expertise they bring, what documentation is produced, and who bears accountability for the outcome, is the measure of whether the system is built for the accountability context that fiduciary obligations create.

Section 5: The Glass Lewis Approach

Glass Lewis has been in the governance research and proxy voting business for 23 years. In that time, the firm has developed a depth of expertise in what enterprise-level, investment-grade proxy research requires, as a practical matter of covering shareholder meetings across global markets, in more than 40 languages, within more than 100 distinct regulatory regimes.

That expertise, in which data points are material, how governance concepts translate across regulatory frameworks, where the edge cases are and why they matter, is the foundation upon which Glass Lewis' AI strategy has been built. Sophisticated AI pipelines are increasingly available across the governance industry. What determines their fitness for fiduciary purposes is the depth of institutional knowledge governing what those pipelines produce, and the years of operational presence across global markets that produced that knowledge.

Glass Lewis brings an equally deep understanding of what institutional asset managers and pension funds need at every stage of the proxy voting process: from governance research and policy development through recommendation, client customization, and vote execution. More than two decades of serving clients across that full chain has shaped how the firm defines quality, where it invests, and how it measures success.

The Operating Philosophy: Human-Centric AI

Glass Lewis' approach to AI is organized around a principle of Human-Centric AI (HCAI): a framework in which AI enhances institutional expertise rather than replacing it, and in which every output is designed to be traceable, auditable, and produced under human-governed rules. The framework has two operational dimensions.

Human expertise embedded in production: The governance analysts in Glass Lewis' production process are not positioned at the end of an automated pipeline to review outputs. They are integrally embedded in the process at defined checkpoints, where institutional knowledge shapes the analysis and informs the methodology the AI

produces. The methodology frameworks they work within were developed from the firm's accumulated experience in governance research.

Human oversight in this framework operates at three levels, each serving a distinct governance function.

1. **Output:** At the output level, high-stakes findings require analyst review and approval before reaching clients.
2. **Process:** At the process level, analysts monitor streams of AI outputs for anomalies and quality drift, recalibrating the system when patterns fall outside expected parameters.
3. **Architecture:** At the architecture level, governance methodologists and AI technologists determine whether a specific AI approach is fit for its intended purpose and retain the authority to halt or replace it if it is not. These three levels operate in parallel, not in sequence.

The result is oversight that covers individual decisions, process health, and architecture integrity simultaneously, a design directly aligned with the EU AI Act's Article 14 requirement that human oversight be built into the architecture.

Glass Lewis' new Climate Intelligence product, launching April 2026, demonstrates this approach in a live production environment. The system applies the firm's governance methodology to AI-powered climate risk assessment, with expert-defined analytical frameworks governing what the AI evaluates, analyst visibility into the evidence base and reasoning chain, and full override capability at every decision point. It is the most concrete current illustration of what the HCAI architecture produces in practice, and the standard against which Glass Lewis is building its AI implementations across research domains.

 **Glass Lewis's AI systems are being built to operate inside a structured methodology layer developed by governance domain experts.**

Rule-governed AI architecture: Glass Lewis's AI systems are being built to operate inside a structured methodology layer developed by governance domain experts. Rather than relying on AI models to extract and interpret governance data without constraint, this layer defines what the AI is permitted to conclude, how evidence must be evaluated, and what documentation must be produced for any output to be considered complete. The result is AI whose outputs are traceable to source evidence and explicit methodology logic – the standard that regulators are beginning to require and institutional clients are increasingly asking their providers to demonstrate.

The Data Ecosystem

Glass Lewis is building a governed, investment-grade data architecture designed to meet the six properties described in Section 4: accuracy validated against source evidence, consistency maintained through deterministic normalization rules, completeness verified against known coverage gaps, validity verified against specified formats and logical reporting rules, timeliness defined as the availability of each dataset within the window required for the research to remain actionable, and traceability preserved from source document to final output. This represents a current, active and continuing investment at Glass Lewis.

What makes this investment grade build defensible is what informs it. The data models, normalization rules, and governance frameworks that define what the platform must produce were derived from more than two decades of operational experience covering shareholder meetings across global markets, in multiple languages, within distinct regulatory regimes. The expertise came first. The architecture follows from it.

This sequencing matters. Organizations building governance AI platforms are also building data infrastructure. The difference is that Glass Lewis is building from a position of knowing what the data needs to look like to support defensible fiduciary decisions. That knowledge is the product of years of domain expertise and in-market analyst work that cannot be easily replicated.

Building a New Glass Lewis

Glass Lewis is in the process of becoming a different kind of company: a technology-enabled data and insights firm whose AI capability is grounded in institutional depth that takes decades to develop, and whose data architecture is governed by the principles that have made the firm's research credible to institutional clients across global markets.

Beginning in 2027, Glass Lewis will no longer issue a single benchmark recommendation.¹⁶ This is a commitment the firm has made publicly.¹⁷ Instead, Glass Lewis will provide multiple AI-powered research perspectives, calibrated to the specific investment philosophies and policies of individual clients. Delivering that at scale, with the analytical rigor and traceability that institutional clients require, ensures that when clients defend their voting decisions to their own stakeholders and regulators, the research behind those decisions can withstand scrutiny. It requires institutional expertise to know what the data needs to look like, governed architecture to produce it reliably, and AI capability to deliver it at the scale that global voting obligations demand.

 **Glass Lewis will provide multiple AI-powered perspectives, calibrated to the specific investment philosophies and policies of individual clients.**

The five questions in Section 4 of this paper are questions Glass Lewis can answer specifically. The firm invites any client, prospect, or evaluator to ask them.

What's Next

As a trusted corporate governance and proxy voting partner, Glass Lewis has always worked in concert with its clients to evolve with the market and guide clients toward sound governance decision-making.

Through this series of whitepapers, in combination with informational videos and planned discussion forums, we hope to help clients successfully navigate the rapidly changing, AI-enabled proxy voting landscape.

This paper is the first in a series. The papers that follow will go deeper into the data architecture, the Human-Centric AI framework, and the regulatory landscape that is reshaping what responsible AI in governance looks like.

Appendix: Key Terms

TERM	DEFINITION
AI Tools in Governance	In this paper, AI tools refer to machine learning systems, symbolic (rule-based) systems, large language models, and automated analytics applied to the processing, classification, and synthesis of governance data and the generation of proxy voting research and recommendations.
AI Pipeline	A structured, automated, and repeatable workflow that converts raw data into actionable intelligence, such as predictions or content generation.
Automation Bias	The documented tendency of human reviewers to defer to AI-generated outputs even when those outputs are incorrect. The EU AI Act's Article 14 human oversight requirements explicitly address automation bias as a risk that must be mitigated by design in high-risk AI systems.
EU AI Act	Regulation (EU) 2024/1689, the world's first comprehensive horizontal legal framework for AI. High-risk provisions including Article 14 human oversight requirements and Article 12 logging requirements apply from August 2, 2026. Establishes requirements for traceability, explainability, human oversight by design, and data governance for AI systems in high-risk contexts.
Explainability	The capacity of an AI system to produce outputs that can be traced to their inputs and explained in terms that a human can understand and defend. In proxy voting AI, explainability means being able to specify, for any recommendation, the data points that informed it, the methodology that evaluated them, and the reasoning that produced the output.

Fiduciary Accountability	<p>The legal and ethical obligation of an asset manager to act in the best interests of beneficiaries in exercising voting rights. Fiduciary accountability is non-delegable: the asset manager, not the AI system or the advisory provider, is legally responsible for voting decisions and must be able to defend them with reference to a documented process and rationale.</p>
Governed Data Architecture	<p>A data infrastructure in which schemas, normalization rules, and quality standards are defined by domain experts and maintained through deterministic processes. Contrasted with probabilistic data extraction, in which AI systems infer data structure from source documents without a pre-defined governance layer.</p>
Human-Centric AI (HCAI)	<p>A framework for AI implementation in which human expertise and institutional knowledge govern the production process rather than reviewing its outputs. Under HCAI, AI enhances human capability rather than replacing it, and the methodology framework governing AI behavior is designed by domain experts before the AI is trained. HCAI is a production philosophy, distinct from HITL in the narrow sense of output review.</p>
Human-in-the-Loop (HITL)	<p>A system design in which human judgment is integrated into an AI production process. In proxy voting AI, HITL describes a wide range of implementations, from human experts whose methodology governs AI behavior throughout production, to human reviewers who evaluate AI-generated outputs before finalization. The nature and location of human involvement determines its governance value.</p>
Investment-Grade Data	<p>Governance data that meets institutional standards for accuracy, consistency, normalization, and traceability. Investment-grade governance data is validated against source evidence, normalized through deterministic rules across regulatory regimes, and produced under human-governed processes rather than probabilistic AI extraction.</p>
SRD II (Shareholder Rights Directive II)	<p>EU directive in force across member states requiring institutional investors and proxy advisors to disclose their engagement policies, voting behavior, and methodology. Proxy advisors must annually disclose the essential features of the methodologies and models applied, primary information sources, and procedures ensuring research quality.</p>

Notes and references

1. Financial Reporting Council. UK Stewardship Code 2020. Principle 12. December 2019. <https://www.frc.org.uk/library/standards-codes-policy/stewardship/uk-stewardship-code/>
2. Financial Services Agency of Japan. Principles for Responsible Institutional Investors: Japan's Stewardship Code. Principle 5. June 2025. <https://www.fsa.go.jp/en/refer/councils/stewardship/20250626/01.pdf>
3. European Union. Shareholder Rights Directive II (SRD II), Article 3j. Directive (EU) 2017/828 of the European Parliament and of the Council. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32017L0828>
4. Daly, B. "(Re)Empowering Fiduciaries in Proxy Voting." U.S. Securities and Exchange Commission. January 8, 2026. <https://www.sec.gov/newsroom/speeches-statements/daly-remarks-nycba-proxy-010826>
5. European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council (EU AI Act), Article 14 (Human Oversight). August 1, 2024. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
6. Ibid.
7. The example is a real-world occurrence for demonstrative purposes, with the company anonymized.
8. Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human-Computer Interaction, 36(6), pp. 495–504 <https://www.tandfonline.com/doi/abs/10.1080/10447318.2020.1741118>
9. European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council (EU AI Act), Article 14 (Human Oversight). August 1, 2024. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
10. International Organization for Standardization. ISO/IEC 42001:2023. Information technology – Artificial intelligence – Management system. 2023. <https://www.iso.org/standard/81230.html>
11. Financial Conduct Authority. AI Update: How Existing FCA Rules Apply to Artificial Intelligence. 2024. <https://www.fca.org.uk/publication/corporate/ai-update.pdf>
12. Financial Services Agency of Japan. Principles for Responsible Institutional Investors: Japan's Stewardship Code. Principle 5. June 2025. <https://www.fsa.go.jp/en/refer/councils/stewardship/20250626/01.pdf>
13. Daly, B. "(Re)Empowering Fiduciaries in Proxy Voting." U.S. Securities and Exchange Commission. January 8, 2026. <https://www.sec.gov/newsroom/speeches-statements/daly-remarks-nycba-proxy-010826>
14. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1, January 2023. <https://doi.org/10.6028/NIST.AI.100-1>
15. European Union. Shareholder Rights Directive II (SRD II), Article 3j. Directive (EU) 2017/828 of the European Parliament and of the Council. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32017L0828>
16. Glass Lewis. "Glass Lewis Leads Change in Proxy Voting Practices." October 15, 2025. <https://www.glasslewis.com/news-release/glass-lewis-leads-change-in-proxy-voting-practices>
17. Mann, B. "Glass Lewis Is Adapting to Customer Needs." The Wall Street Journal. November 25, 2025. <https://www.wsj.com/opinion/glass-lewis-is-poised-to-change-its-old-ways-proxy-voting-7f6b6d0c>

Disclaimer

© 2026 Glass, Lewis & Co., and/or its affiliates. All Rights Reserved.

This document is for informational purposes only. Glass Lewis shall not be liable for any losses or damages of any kind arising from or in connection with any of the information presented through this document. Under no circumstance shall this document be construed as providing consulting or advisory services to any publicly traded company, its board of directors or any advisor thereof. Furthermore, none of the statements made in this document shall be deemed to express an opinion or conclusion, nor be deemed to provide any assurance, regarding the outcome of any proxy vote recommendations issued by Glass Lewis. Glass Lewis' research reports are based solely on publicly available information and all issues are analyzed on a case-by-case basis. Glass Lewis' market-specific policy guidelines are available at <https://www.glasslewis.com/guidelines/>.

Glass Lewis provides its clients with research, data, and analysis of proxy voting issues, and makes recommendations as to how institutional shareholders should vote their proxies, without commenting on the investment merits of the securities issued by the subject companies. Therefore, none of Glass Lewis' proxy vote recommendations should be construed as a recommendation to invest in, purchase, or sell any securities or other property. Moreover, Glass Lewis' proxy vote recommendations are solely statements of opinion, and not statements of fact, on matters that are, by their nature, judgmental. Glass Lewis research, analyses, and recommendations are made as of a certain point in time and may be revised based on additional information or for any other reason at any time.

The information contained in this report is based on publicly available information. While Glass Lewis exercises reasonable care to ensure that all information included in this report is accurate and is obtained from sources believed to be reliable, no representations or warranties express or implied, are made as to the accuracy or completeness of any information included herein. Such information may differ from public disclosures made by the subject company. In addition, third-party content attributed to another source, including, but not limited to, content provided by a vendor or partner with whom Glass Lewis has a business relationship, as well as any Report Feedback Statement, are the statements of those parties and shall not be attributed to Glass Lewis. Neither Glass Lewis nor any of its affiliates or third-party content providers shall be liable for any losses or damages arising from or in connection with the information contained herein or the use or inability to use any such information.

Glass Lewis' reports are intended to serve as a complementary source of information and analysis for subscribers in making their own voting decisions and therefore should not be relied on by subscribers as their sole determinant in making voting decisions. Glass Lewis expects its subscribers to possess sufficient experience and knowledge to make their own decisions entirely independent of any information contained in its reports, and subscribers are ultimately and solely responsible for making their own voting decisions, including, but not limited to, ensuring that such decisions comply with all agreements, codes, duties, laws, ordinances, regulations, and other obligations applicable to such subscriber.

All information contained in this document is protected by law, including but not limited to, copyright law, and none of such information may be copied or otherwise reproduced, repackaged, further transmitted, transferred, disseminated, redistributed or resold, or stored for subsequent use for any such purpose, in whole or in part, in any form or manner or by any means whatsoever, by any person without Glass Lewis' express prior written consent. The foregoing includes, but is not limited to, using this document, in any manner and in whole or in part, in connection with any training, self-improving, or machine learning software, algorithms, hardware, or other artificial intelligence tools or aids of any kind, including, without limitation, large language models or other generative artificial intelligence platforms or services, whether proprietary to you or a third party, or generally available (collectively, "AI") as well as any services, products, data, writings, works of authorship, graphics, pictures, recordings, any electronic or other information, text or numerals, audio or visual content, or materials of any nature or description generated or derived by or using, in whole or in part, AI.

This document should be read and understood in the context of other information Glass Lewis makes available concerning, among other things, its research philosophy, approach, policy guidelines, supplementary guidance and methodologies, sources of information, and conflict management, avoidance and disclosure policies and procedures, which information is incorporated herein by reference. Glass Lewis recommends all clients and any other consumer of this document carefully and periodically evaluate such information, which is available at: <http://www.glasslewis.com>.



AI and the Fiduciary Test

Glass Lewis is a leading independent governance research and proxy advisory firm serving more than 1,300 institutional investors globally.

Connect With Glass Lewis

Corporate Website | [glasslewis.com](https://www.glasslewis.com)

Email | info@glasslewis.com

Social | [!\[\]\(c428d2e50ffe0e934b3cfcdbc1ca5f24_img.jpg\)](#)