Item 14. Development of AI Governance Principles

CalSTRS Teachers’ Retirement Board | July Offsite

Mosaic Governance Advisors | Dr. Param Vir Singh, Carnegie Mellon University

MOSAIC
Governance Advisors

Agenda

- Introduction
- Foundational Beliefs
- AI Governance Principles
- Next Steps

Today’s Discussion

- Build on initial alignment from survey results
- Stress-test draft governance principles
- Concepts, not wordsmithing

Governing in the Age of AI:

Defining Fiduciary Roles & Establishing Board-Level Principles

THE INTELLIGENT FUTURE

Param Vir Singh
Associate Dean for Research
Carnegie Bosch Professor of Business Technologies

Carnegie Mellon University
Tepper School of Business

What do we mean by AI?

AI systems perform tasks historically requiring human intelligence: **reasoning**, **language processing**, and **decision support**.

Traditional AI

PREDICTS

Processes historical data to forecast trends and flag anomalies.

Generative AI

CREATES

Synthesizes unstructured data to draft text, summaries, and media.

Agentic AI

ACTS

Executes multi-step workflows and acts autonomously within bounded parameters.

Member Services: Generative vs Agentic AI

GENERATIVE AI

Draft the Communications

"Draft a clear response explaining how a complex, recent legislative update affects this specific teacher's retirement estimate."

Output: A synthesized draft response for a human staff member to verify, modify, and send.

AGENTIC AI

Execute the Process

"Identify missing verification files, safely query our secure database, contact the employer for back-records, and assemble the file."

Output: Autonomous execution of workflows that actively interacts with multiple systems.

Governance should cover systems that **CREATE** and that **ACT.**

Fiduciary Oversight in the Age of AI

- Unlike traditional software, AI is **probabilistic**—it learns from patterns rather than strict, hardcoded rules.
- Local operational errors can instantly scale into systemic enterprise liabilities, affecting capital allocation and member rights.

Four Core Questions for the Board

- Where is AI influencing material decisions?
- What authority is being delegated?
- Who remains ultimately accountable?
- What oversight is proportional to the risk?

“Noses In, Fingers Out”

Management Decides

Execution & Selection

- Selecting tools and technical vendors
- Managing operational deployment
- Enforcing technical controls
- Ensuring day-to-day data hygiene

The Board Decides

Guardrails & Accountability

- Governance expectations & values
- Acceptable enterprise risk tolerance
- Final accountability frameworks
- Oversight reporting thresholds

The Governance Architecture

GOVERNANCE LAYER	PRIMARY FOCUS	RESPONSIBILITY
Board Level (Governance Principles)	**Oversight & Accountability**	Establishing fiduciary expectations, defining risk boundaries, and protecting institutional trust.
Management Level (Staff AI Policy)	**Operating Principles & Execution**	Implementing day-to-day controls, vendor validation, output testing, and operational deployment.

VALUE OF FOUNDATIONAL BELIEFS

What are Foundational Beliefs?

Foundational Beliefs: The Board's Enduring Worldview

Foundational beliefs are short, durable statements that capture how the Board understands the world at the time they are adopted.

Beliefs answer: "What do we believe?"

Principles answer: "How should we govern?"

Why do Boards Document Beliefs?

Continuity

- Preserves the Board's unified governance perspective across standard trustee turnover and long-term leadership successions.

Interpretive Lens

- Helps future Boards understand the original strategic intent behind core principles, ensuring consistent interpretation as new AI applications arrive.

Durability

- AI capabilities accelerate rapidly.
- Documented beliefs provide a permanent strategic anchor even when implementation parameters shift.

Belief: AI will have a material impact on how CalSTRS operates.

Implication: Board oversight must remain durable, principles-based, and proportional to risk.

From Beliefs to Governance Principles

Foundational Belief	Governance Principle
AI oversight is grounded in fiduciary responsibility.	Fiduciary Alignment
Humans remain responsible for consequential decisions.	Human Accountability
Public trust must be preserved.	Trust
Human judgment remains primary.	Human Judgment
AI should be governed proportionately and fairly.	Fairness

Foundational Belief Statements

Draft for board consideration | Numbered for discussion purposes only

AI and Fiduciary Responsibility

1. Artificial Intelligence (AI) is a distinct category of technology that requires Board-level oversight.
2. We ground our oversight of AI usage in our fiduciary responsibility to ensure CalSTRS meets the mission of securing the financial future and sustaining the trust of California's educators.

AI's Long-Term Impact

3. There are both opportunities and risks in adopting AI.
4. AI will have a material impact on how CalSTRS operates.

Governance and Human Judgment

5. Human judgment should remain primary for decision-making processes affecting members and those with material investment implications.
6. Board-level oversight of AI should be in proportion to the risk and impact of individual use cases.

Innovation with Guardrails

7. The governance principles are intended to provide clear guardrails that enable CalSTRS staff to deliberately, intentionally, and proactively pursue innovation.

Discussion

- Do these foundational beliefs appropriately reflect the Board's consensus view?

Board-Level AI Governance Principles

AI Governance Principles

Draft for board consideration | Numbered for discussion purposes only

- **Fiduciary Alignment:** AI is only adopted when it strengthens CalSTRS' ability to meet its fiduciary obligations and long-term responsibilities to members.
- **Human Accountability:** CalSTRS' people remain fully responsible for decisions and outcomes influenced by AI and must be prepared to explain and stand behind those decisions.
- **Trust:** AI use upholds CalSTRS' role as a trusted public institution responsible for the retirement security of California's educators.
- **Human Judgment:** AI use builds CalSTRS' institutional capability by enhancing the human judgment, collaboration, and discretion required to do CalSTRS' work well and to continue to grow and develop the CalSTRS team.
- **Fairness:** CalSTRS seeks to avoid unfair or discriminatory outcomes in its use of AI.

Discussion Context & Objective

- Principles, once Board-approved, will be public-facing
- They can be easy to agree with in isolation
- Governance happens when principles interact with each other in the face of operating realities
- Engage in scenario thinking to align on the appropriateness of the principles and identify where change might yield improvements
- Focus on concepts, not wordsmithing

MOSAIC Governance Advisors

Small Group Exercise

- Break into small groups
- Each group has a worksheet aligned to one principle
- Groups have 30 minutes for discussion
- Appoint a scribe for note-taking

Small Groups

Fiduciary Alignment

- Denise Bradford
- Ken Tang
- Steve Juarez
- Brian Bartow
- Cassandra Lichnock

Human Accountability

- Office of the State Treasurer
- Office of the State Superintendent
- Lisa Blatnick
- Ashish Jain

Trust

- Office of the State Controller
- Michael Gunning
- Teresa Schilling
- Jeff Zimmer

Human Judgment

- Sharon Hendricks
- Harry Keiley
- Scott Chan
- Melissa Norcia

Fairness

- Office of the Director of Finance
- Karen Yamamoto
- David Lamoureux
- Julie Underwood

Discussion

MOSAIC
Governance Advisors

Discussion Focus

- Are you comfortable with the way the principle is currently worded?
- If it needs more work, what concepts are missing or need more clarity?

Scenario 1: Fiduciary Alignment

CalSTRS provides the Board with access to an AI tool that helps summarize Board materials and generate questions. The tool also includes a chatbot feature that allows Board members to query past meeting materials. The tool is intended to improve efficiency and effectiveness and to save time.

Scenario 2: Human Accountability

CalSTRS staff uses an AI tool to review member benefit calculations and flag potential inconsistencies for staff to review before final decisions are made.

Scenario 3: Trust

CalSTRS uses an AI chatbot to respond to member questions about benefits, timelines, and account information, alongside existing call center and staff support.

Scenario 4: Human Judgment

CalSTRS staff uses an AI tool to conduct and summarize research, highlight key risks and considerations, and provide an initial recommendation prior to engaging in discussion about investment decisions.

Scenario 5: Fairness

CalSTRS implements an AI tool to prioritize incoming member service requests based on factors such as urgency and complexity.

AI Governance Principles

Draft for board consideration | Numbered for discussion purposes only

1. **Fiduciary Alignment:** AI is only adopted when it strengthens CalSTRS' ability to meet its fiduciary obligations and long-term responsibilities to members.
2. **Human Accountability:** CalSTRS' people remain fully responsible for decisions and outcomes influenced by AI and must be prepared to explain and stand behind those decisions.
3. **Trust:** AI use upholds CalSTRS' role as a trusted public institution responsible for the retirement security of California's educators.
4. **Human Judgment:** AI use builds CalSTRS' institutional capability by enhancing the human judgment, collaboration, and discretion required to do CalSTRS' work well and to continue to grow and develop the CalSTRS team.
5. **Fairness:** CalSTRS seeks to avoid unfair or discriminatory outcomes in its use of AI.

Next Steps

Next Steps

- Refine draft foundational beliefs and governance principles for further consideration at September's Board meeting
- Additional Next Steps
  - September ARM Committee
  - Integration discussions with the CEO to implement Working Group's recommendations
  - AI Use Case Overview and Readiness Assessment

Thank You

MOSAIC
Governance Advisors