AI Agents in Finance:
The Risks You Can't Afford to Ignore

Mark McDonald | April 2026

Accelerating AI for Finance and Accounting

AI Agents Are On the Rise

AI agents have captured the attention of finance and accounting teams

59%

of finance leaders report using AI in their finance function

Gartner AI in Finance Survey, 2025

80%

of finance professionals expect AI agents to become standard tools within 5 years

Deloitte Center for Controllership, 2025

79%

of senior executives say AI agents are already being adopted at their companies

PwC, 2025

AI agents present new opportunities for a function in need of improvements

But Adoption Is Outpacing Governance

The speed of adoption creates new risks in finance and accounting

21% cite trust as the #1 barrier to AI agent adoption in finance

Deloitte Center for Controllership, 2025

73% of governance executives say AI has exposed gaps in oversight

OneTrust Governance Survey, 2025

29% of employees have turned to unsanctioned AI agents for work

Microsoft Cyber Pulse Report, 2025

Agents are here

Do you know your risks?

Do we know what agents do?

How do we control them?

1. Defining Agents
2. Identifying Risks and Controls
3. Moving Forward
4. Q&A

AUDIENCE POLL

What is your biggest concern with AI agents in finance?

A Data security and privacy

B Accuracy and hallucinations

C Lack of regulatory guidance

D Unsanctioned use by employees

Select your answer in the poll

1 **Defining Agents**

2 Identifying Risks and Controls

3 Moving Forward

4 Q&A

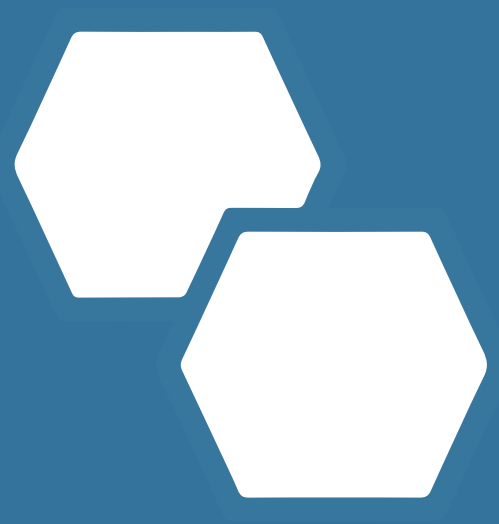

What is an AI agent?

Classical AI is Great With Numbers

Classical AI

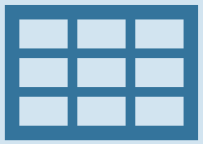

Numbers

Generative AI is Great with Words

Classical AI

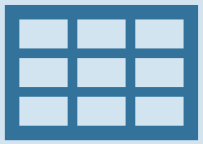

Numbers

Generative AI

Text

Recent Developments Address Shortcomings

Classical AI

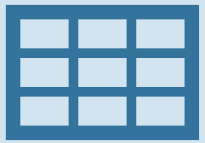

Numbers

Generative AI

Text

Agentic AI Solves Previous Problems

Classical AI

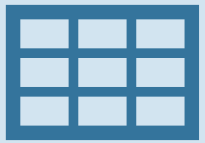

Numbers

Agentic AI

Generative AI

Text

LLMs Match Text Patterns

Chatbot

My heart aches. What should I do?

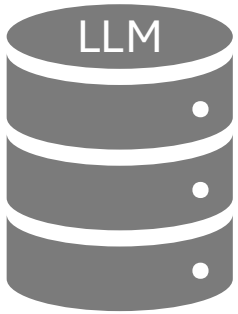

LLMs Match Text Patterns

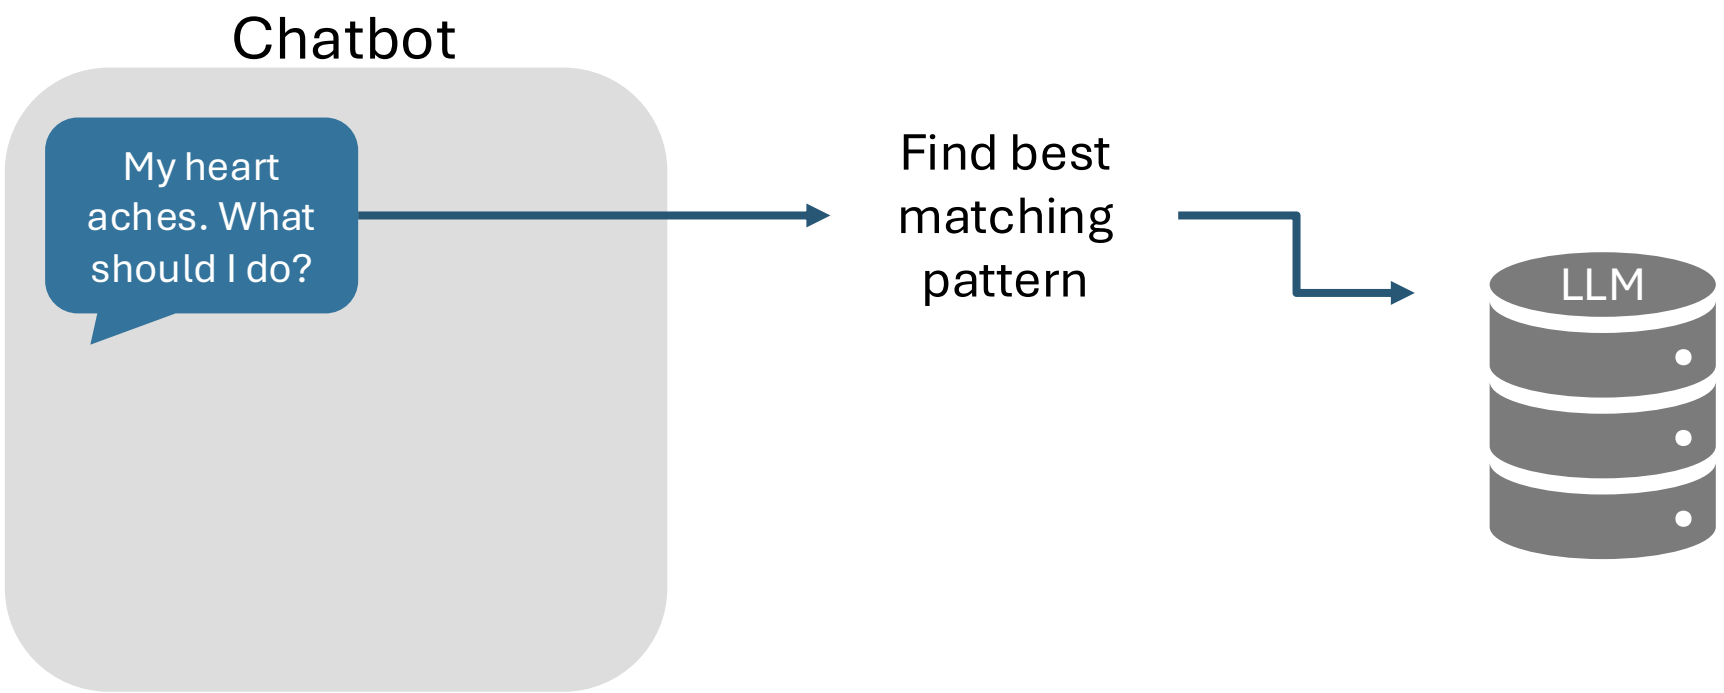

LLMs Match Text Patterns

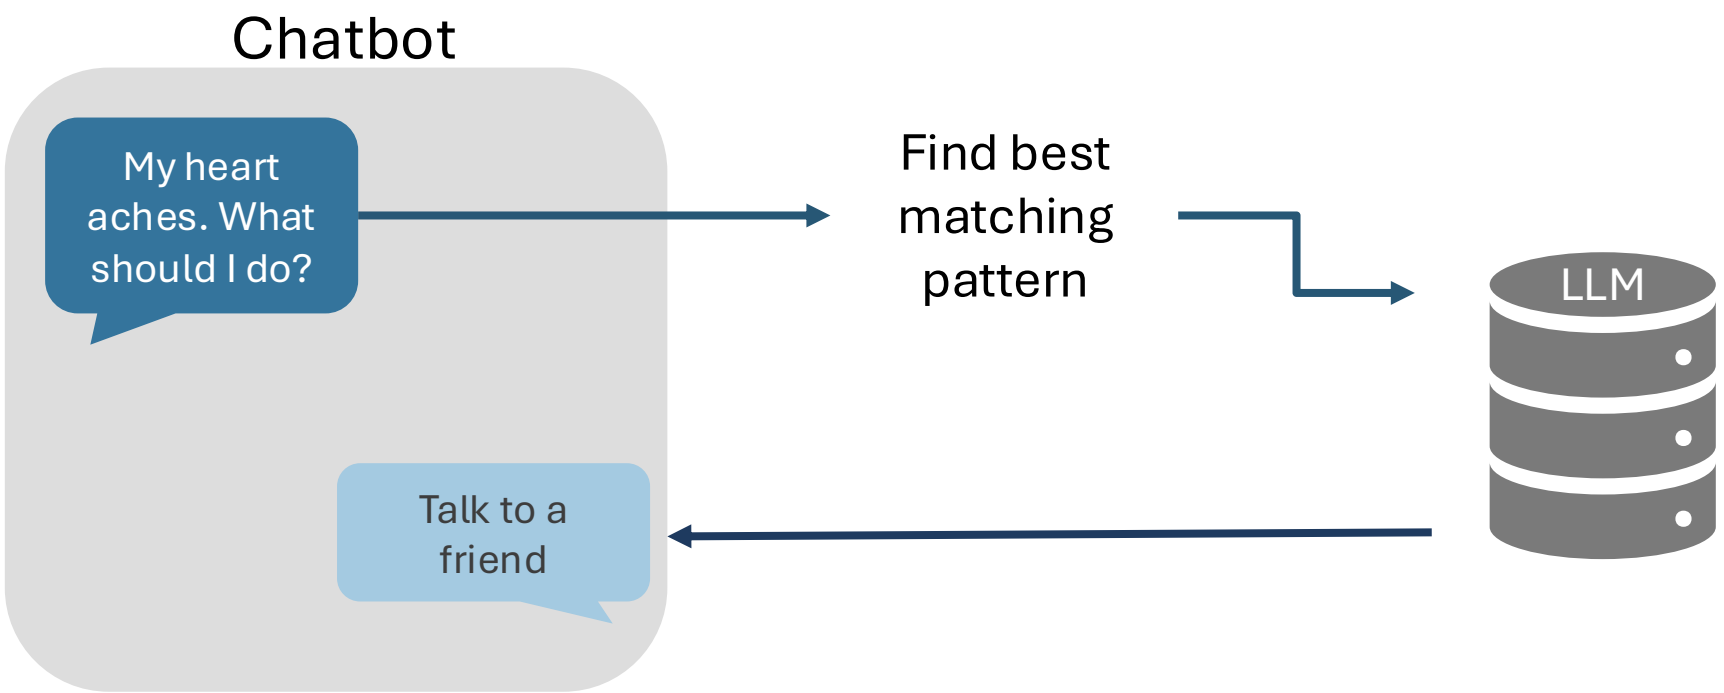

Agents Handle The Process Differently

Chatbot

What is balance of account #123?

Chatbot Looks For Help Responding

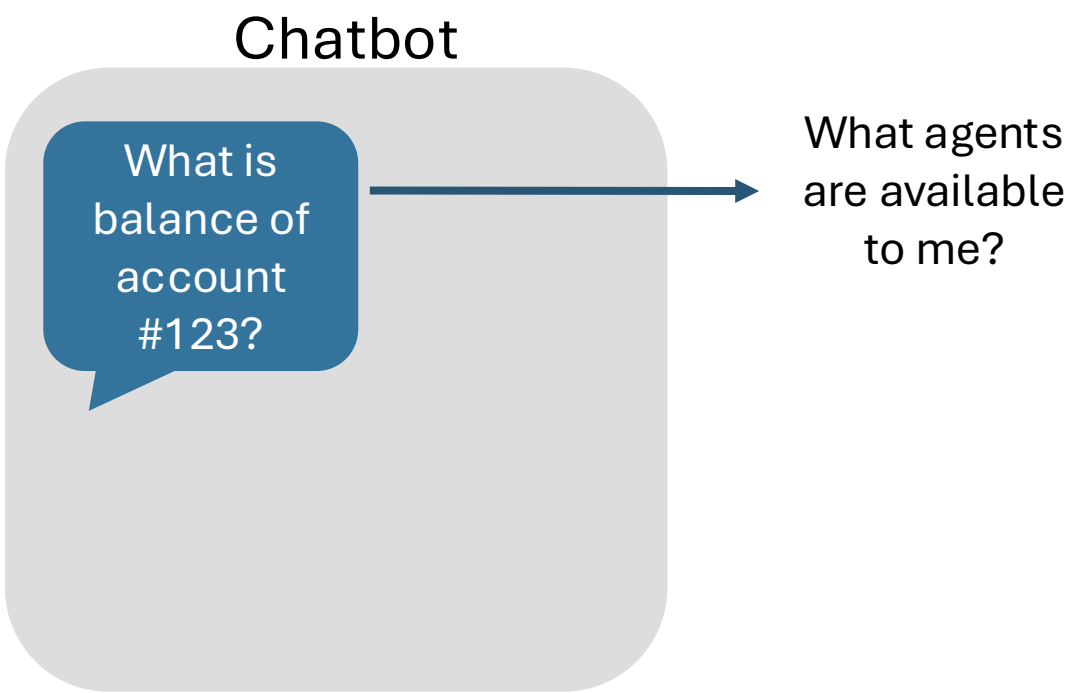

Chatbots Are Connected to Agents

Chatbot Looks For Best Agent

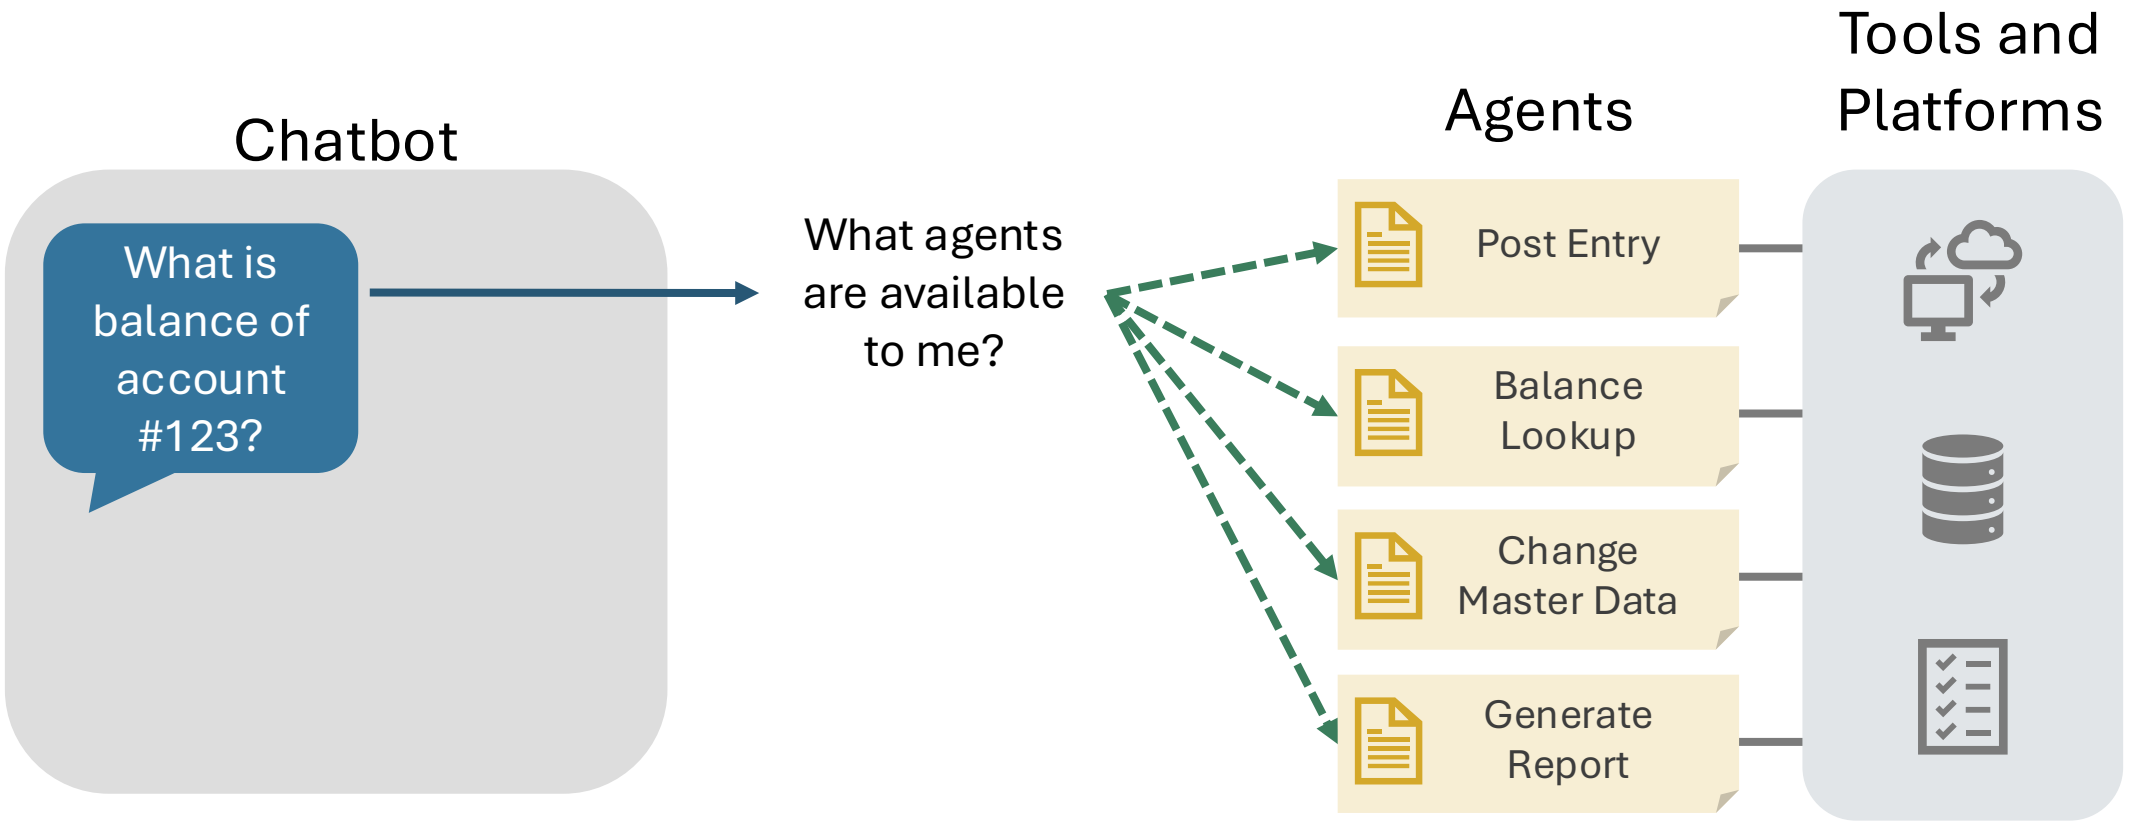

Chatbot Delegates Request

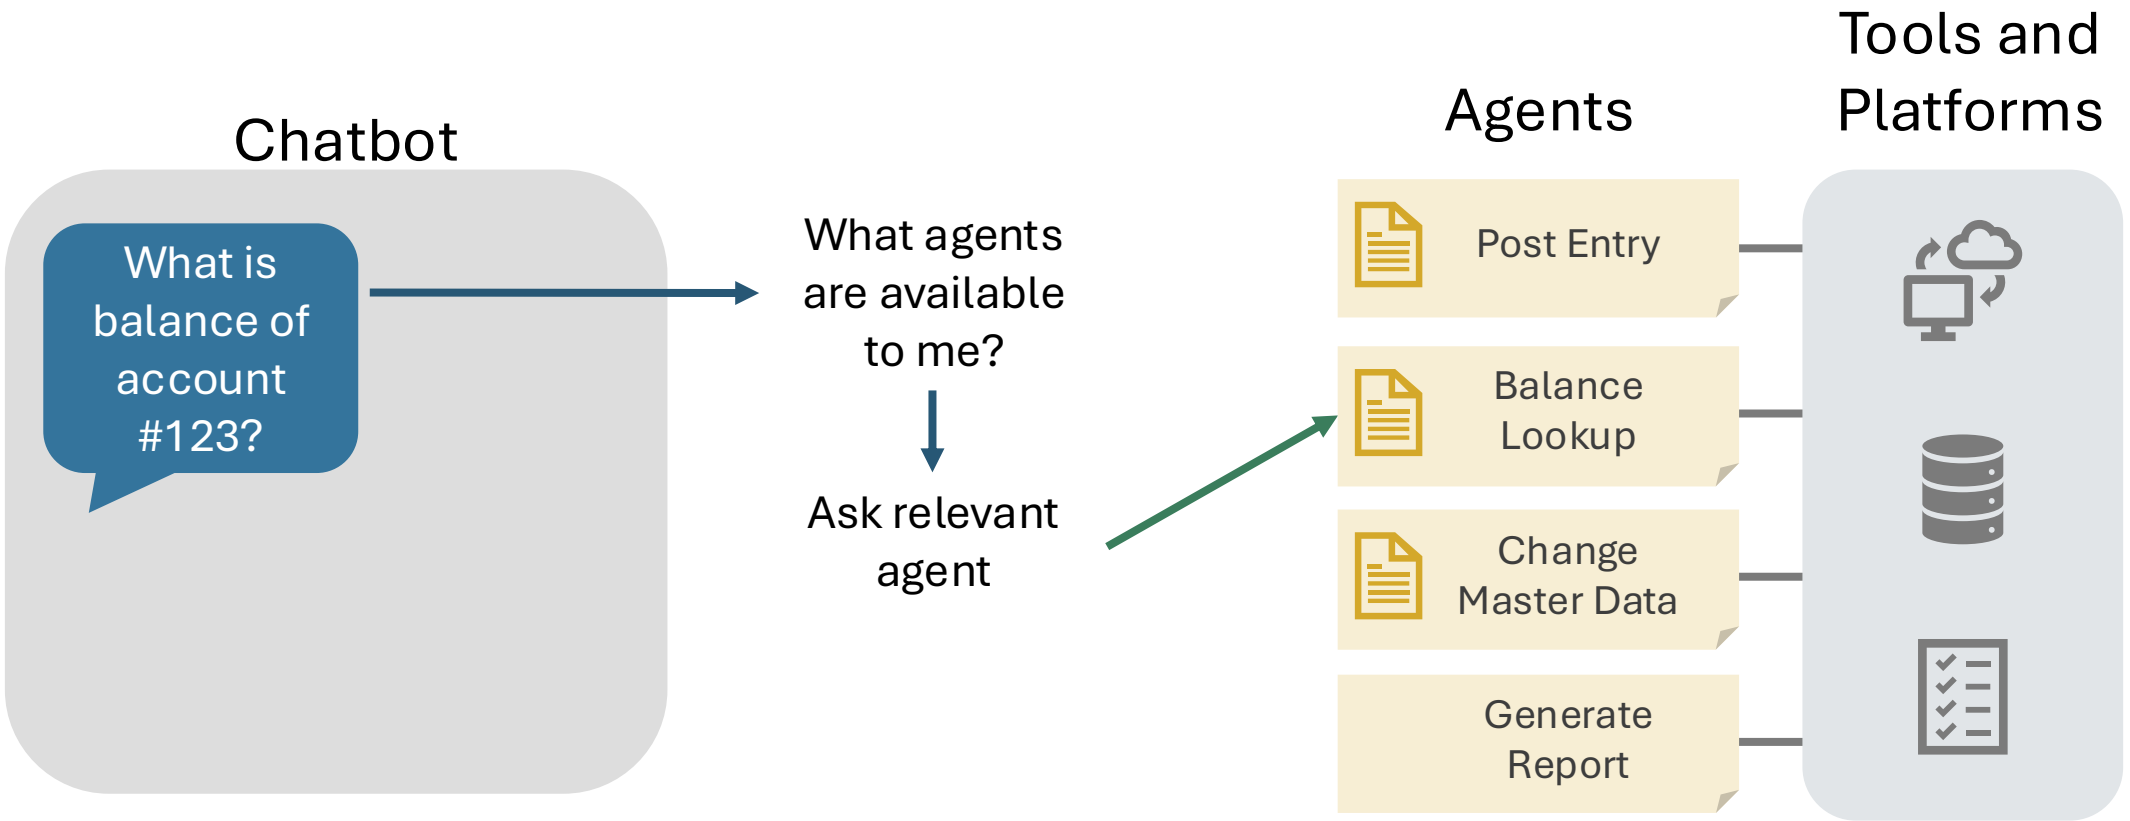

Agent Performs Promised Task

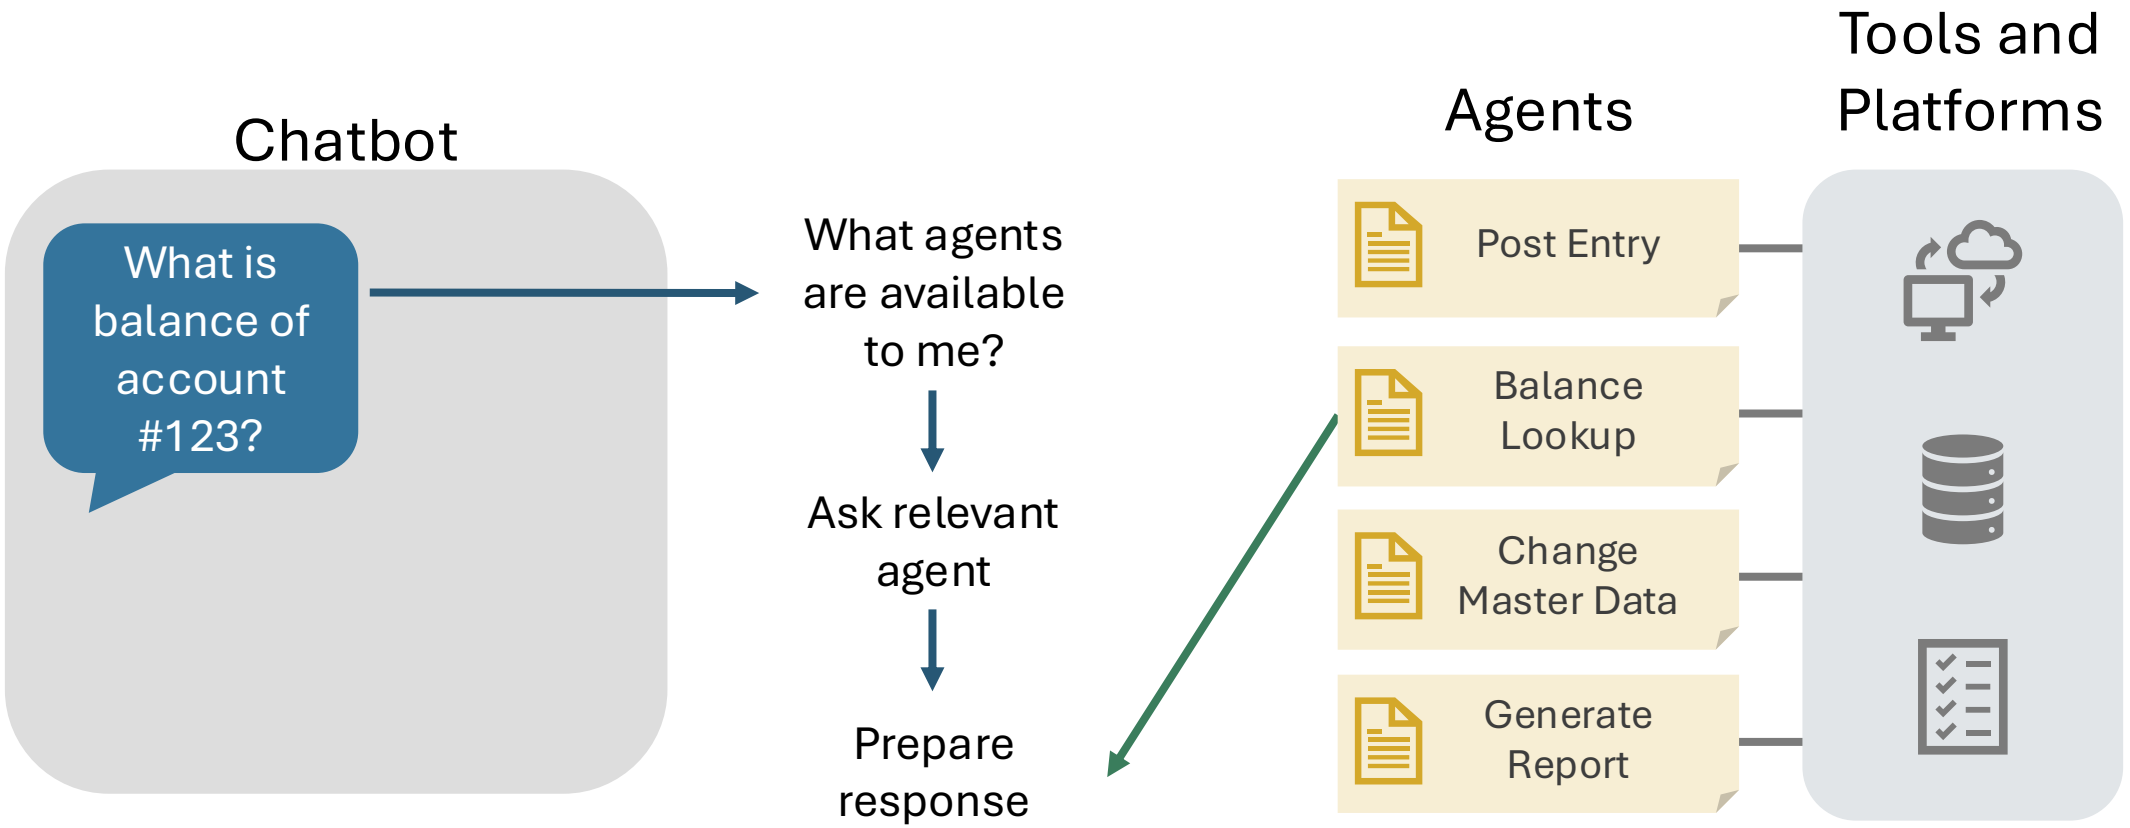

Chatbot Delivers Formatted Results

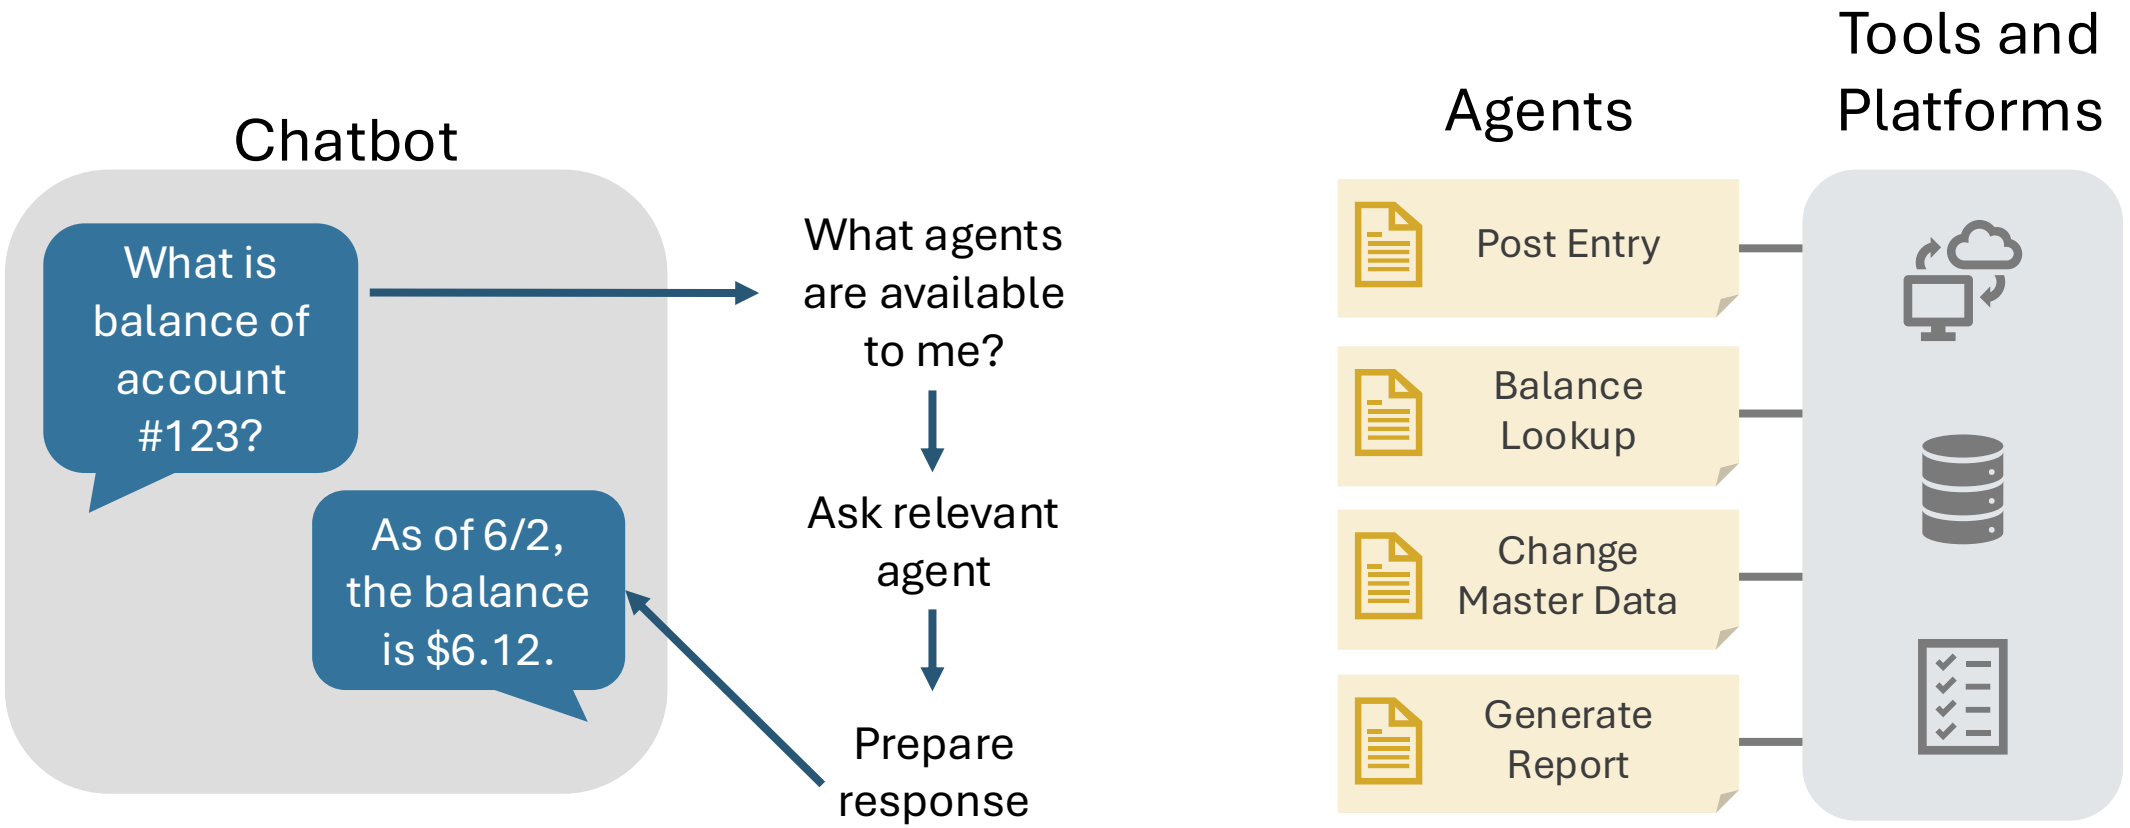

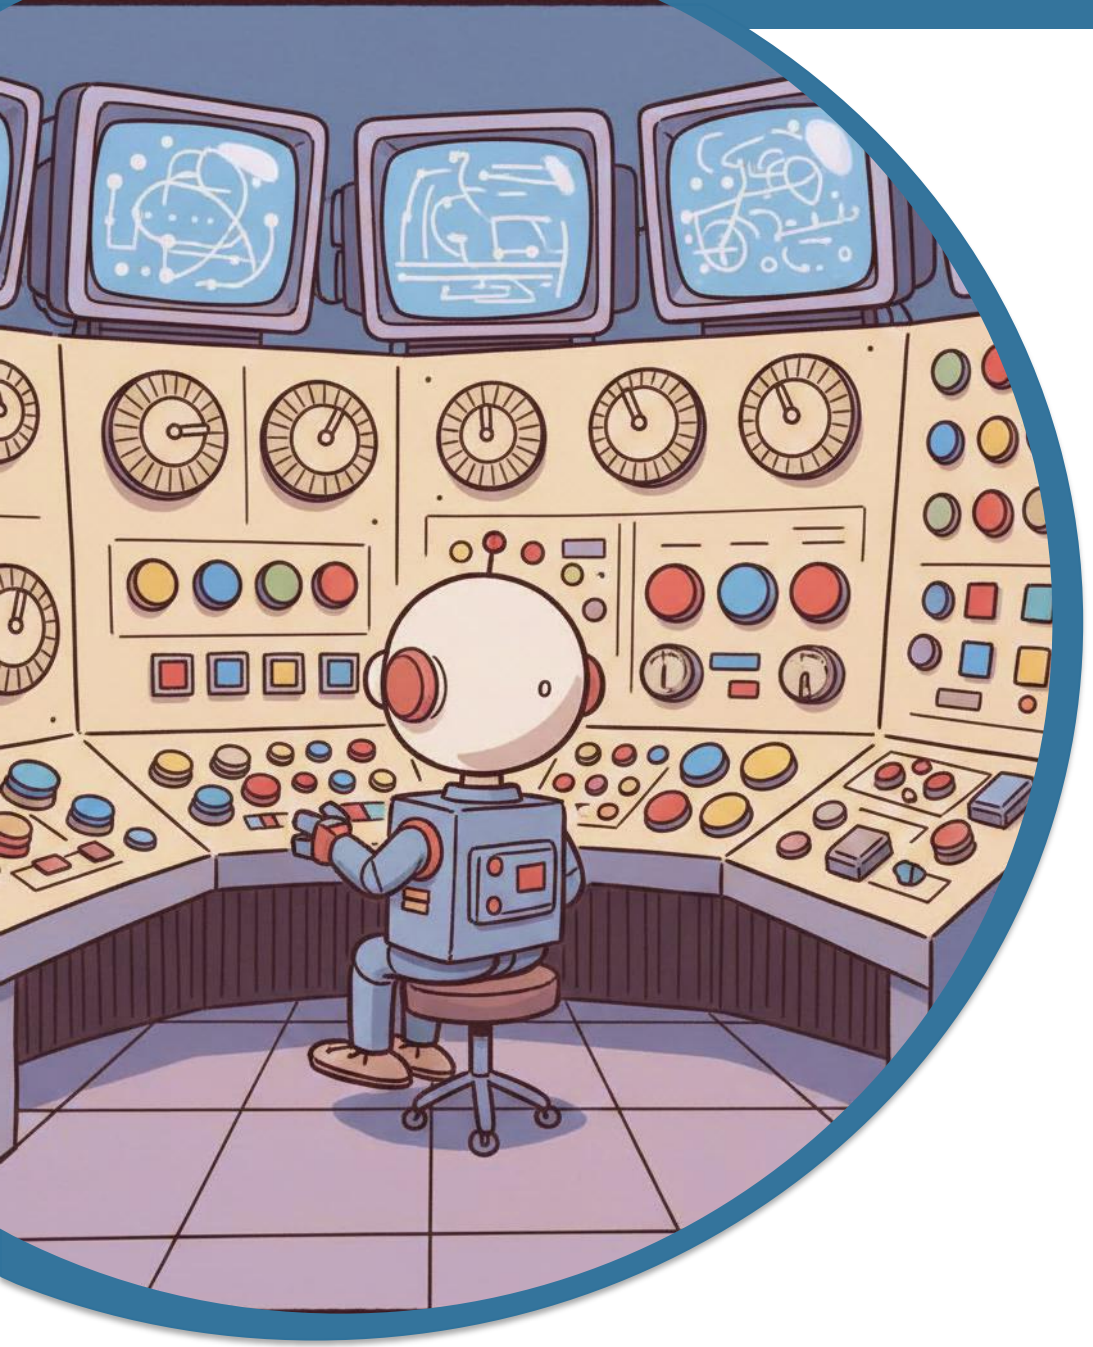

The Power of Agents

- Orchestrate workflows
- Access enterprise data
- Leverage legacy environment
- Enforce security and privacy
- Allow customization and control

AUDIENCE POLL

How much do you trust AI agents to make decisions in your finance workflows?

A Fully. Let them run autonomously.

B Only within defined guardrails

C Only with human approval each time

D Not at all. Too risky right now.

Select your answer in the poll

Agents also have limitations

1 Defining Agents

2 Identifying Risks and Controls

3 Moving Forward

4 Q&A

Regulatory guidance varies

Limited AI Guidance in the US

Bodies rely on existing frameworks rather than issuing AI-specific rules

SEC
Securities & Exchange Commission

PCAOB
Public Company Accounting Oversight Board

AICPA
American Institute of CPAs

FASB
Financial Accounting Standards Board

- Regulatory bodies offer little conclusive guidance
- Most align AI to existing frameworks
- No directives that offer concrete controls

Some published opinions from bodies point to future guidance

EU AI Governance is More Specific

Unlike the US, the EU has enacted AI legislation with specific compliance obligations, mostly to protect individuals

Main EU Directives

- AI literacy required for users of AI systems
- Human oversight of AI-supported decisions that impact individuals
- Lawful basis for use of personal data
- Technical documentation and transparency for high-risk AI systems

EU AI Act
Regulation (EU) 2024/1689

GDPR
General Data Protection Regulation

DORA
Digital Operational Resilience Act

ESAs
EBA / ESMA / EIOPA

EU rules application based on where AI systems produce effects, not where the organization is incorporated

EU + CA AI Governance is More Specific

Unlike the US, the EU has enacted AI legislation with specific compliance obligations, mostly to protect individuals

Main EU + CA Directives

- AI literacy required for users of AI systems
- Human oversight and control of AI-supported decisions that impact individuals
- Lawful basis for use of personal data
- Technical documentation and transparency for high-risk AI systems

EU AI Act
Regulation (EU) 2024/1689

GDPR
General Data Protection Regulation

DORA
Digital Operational Resilience Act

ESAs
EBA / ESMA / EIOPA

+

CA rules are applied after reaching thresholds in revenue ($26.6M), number of CA residents impacted (≥100k) and portion of revenue from that data (≥50%).

What do we do with this diverse set of guidelines?

AUDIENCE POLL

How prepared is your finance organization to manage AI agent risks today?

A We have a framework in place

B We're working on it

C We know we need to but haven't started

D It's not on our radar yet

Select your answer in the poll

AI Agent Risk Control Framework

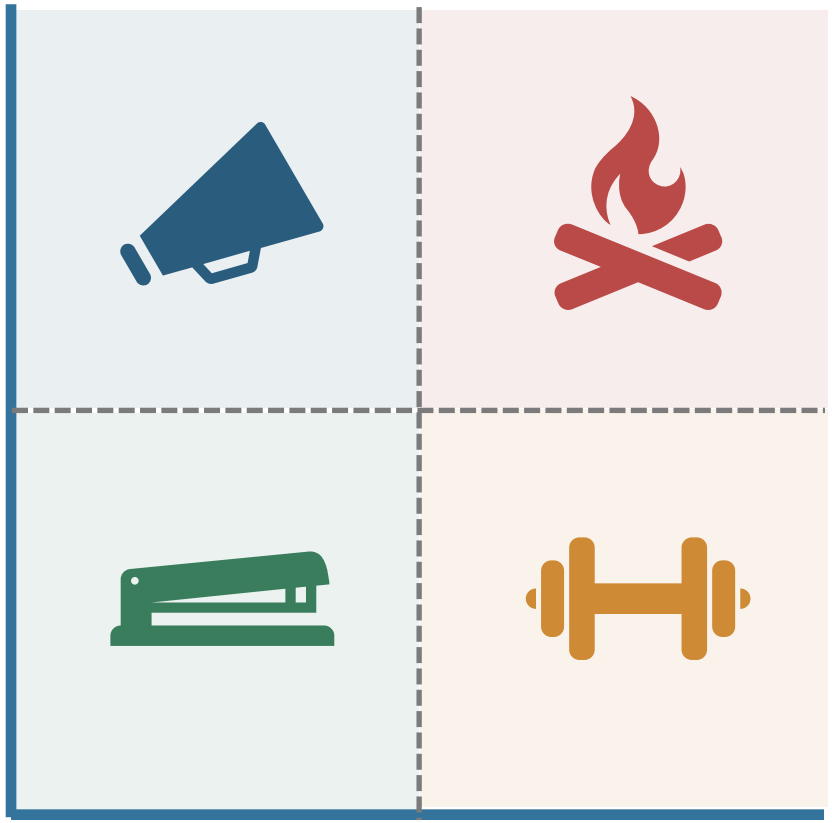

Framework Objectives

- Document agent usage
- Expose, measure and mitigate risks
- Covers existing and anticipated guidelines
- Assign controls to agents based on their risk profile
- Communicate to stakeholders

Dimension 1 of 2: Task Complexity

Task Complexity
Complexity of task supported by the agent

Dimension 1 of 2: Task Complexity

Task Complexity
Complexity of task supported by the agent

Complexity Characteristics

- Subjectivity of output
- Verification effort
- Number of sources and supporting agents

Dimension 2 of 2: Risk Exposure

Task Complexity
Complexity of task supported by the agent

Risk Exposure
Potential magnitude of harm in the event of failure

Dimension 2 of 2: Risk Exposure

Risk Components

- Financial risk
- Operational risk
- Security risk

Task Complexity
Complexity of task supported by the agent

Risk Exposure
Potential magnitude of harm in the event of failure

Assign Dimensions to Axes

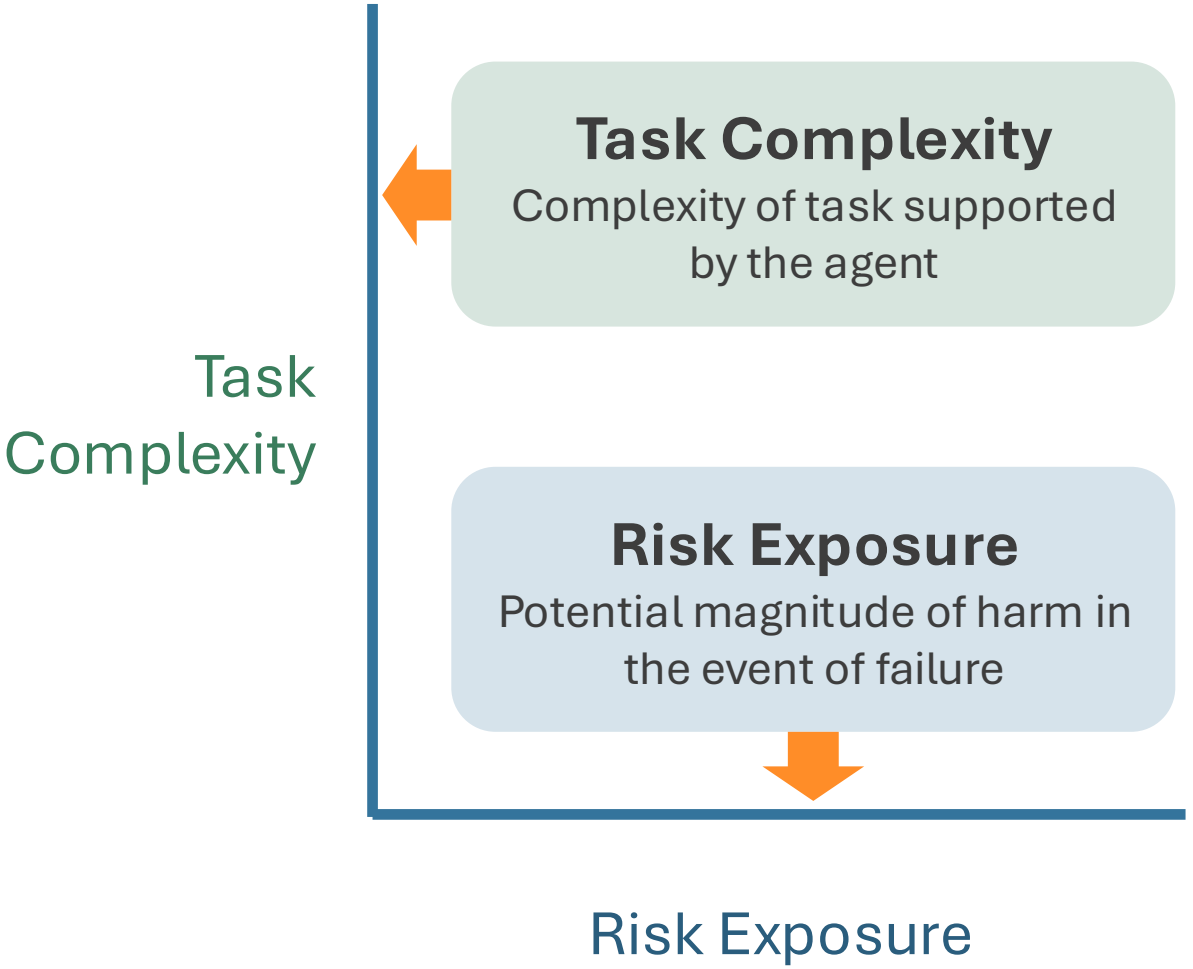

AI Agent Risk Control Framework

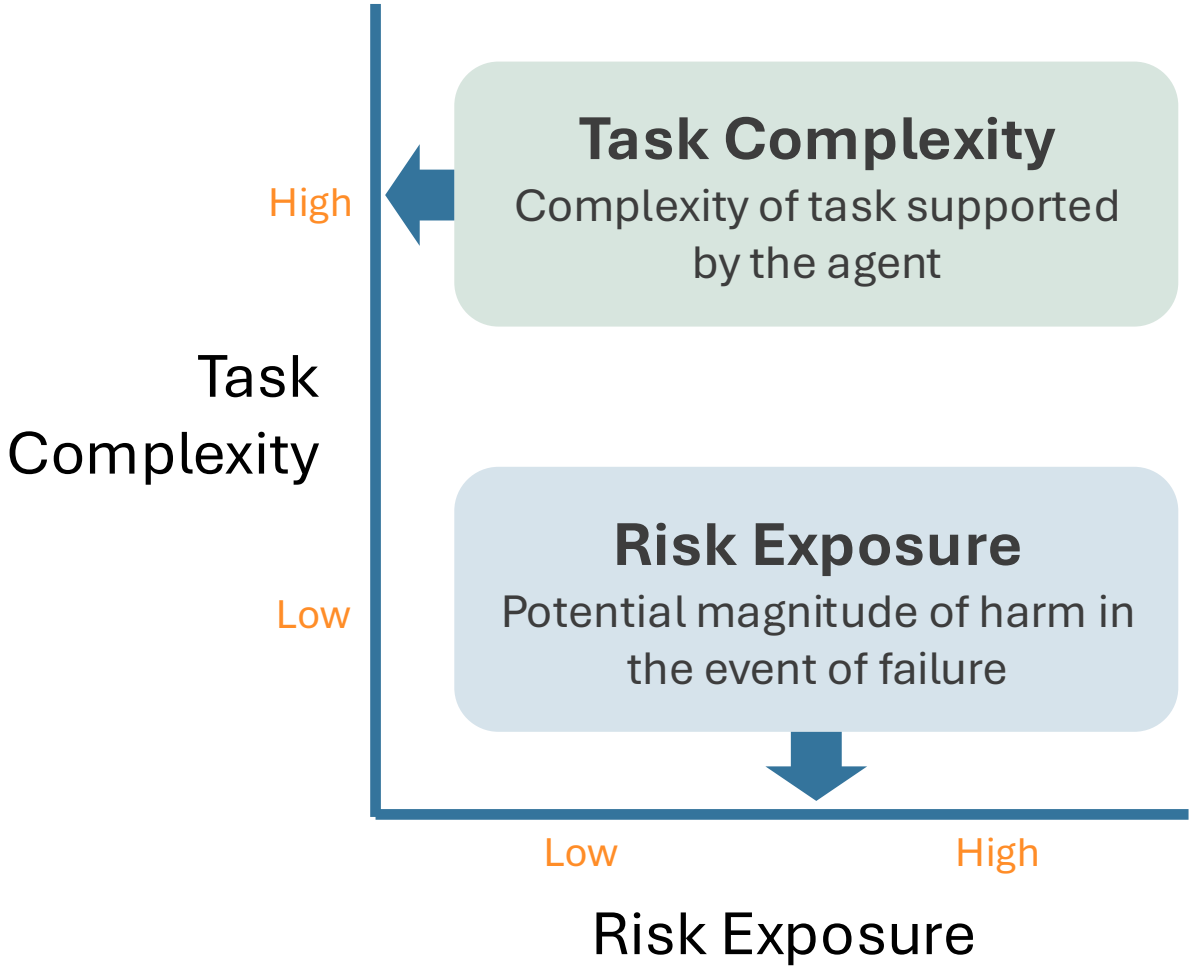

Agents are aligned to 4 Quadrants

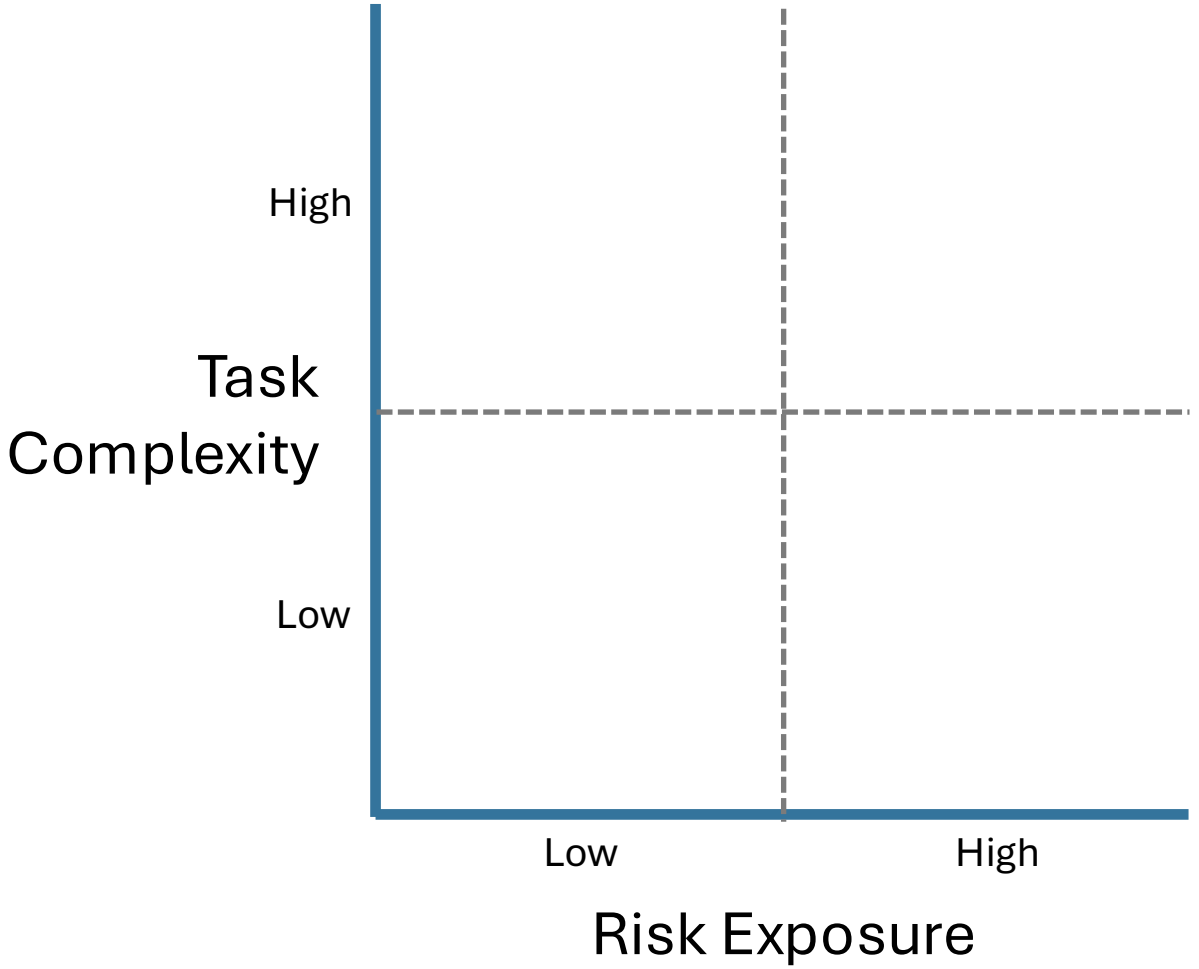

Agents are aligned to 4 Quadrants

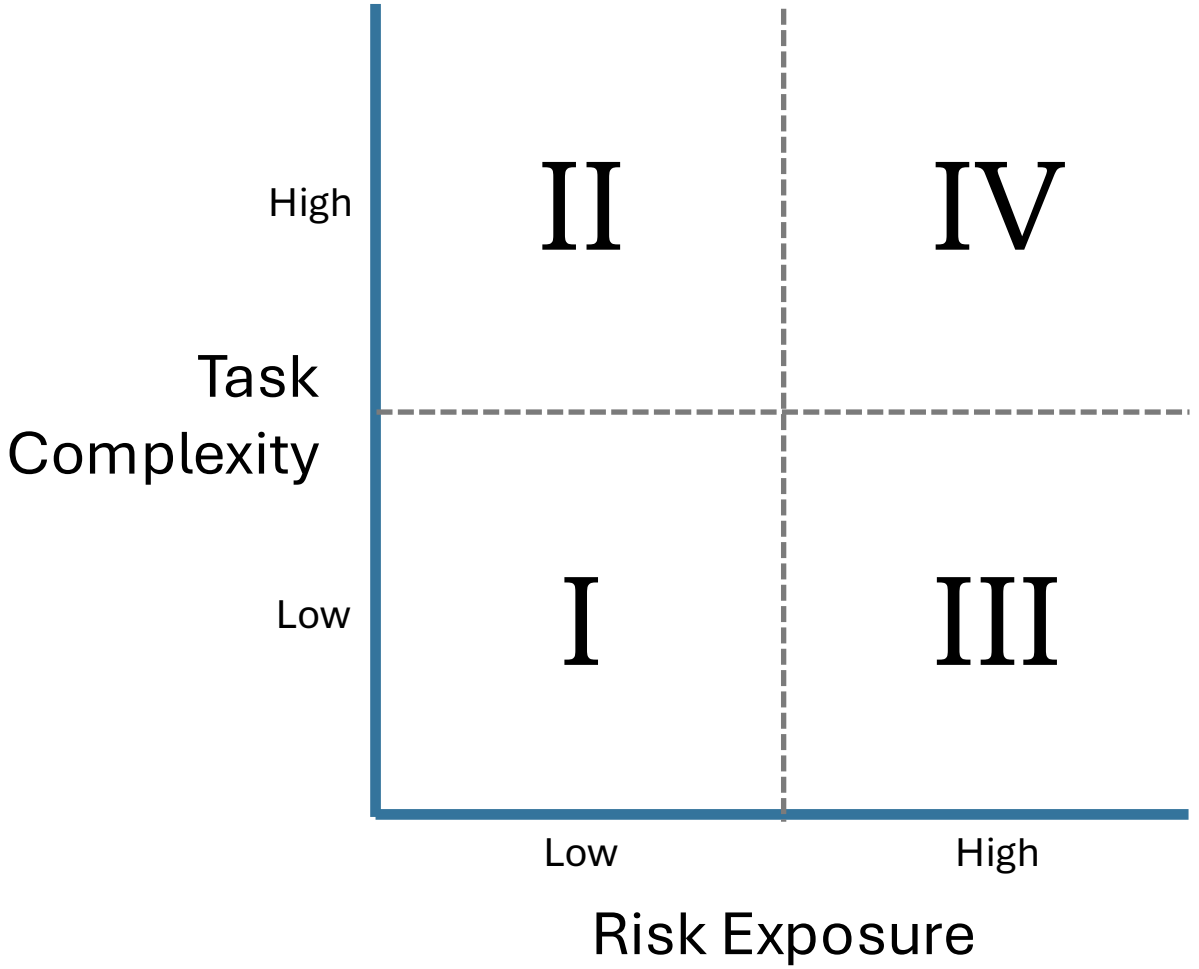

Increasing Risk and Controls

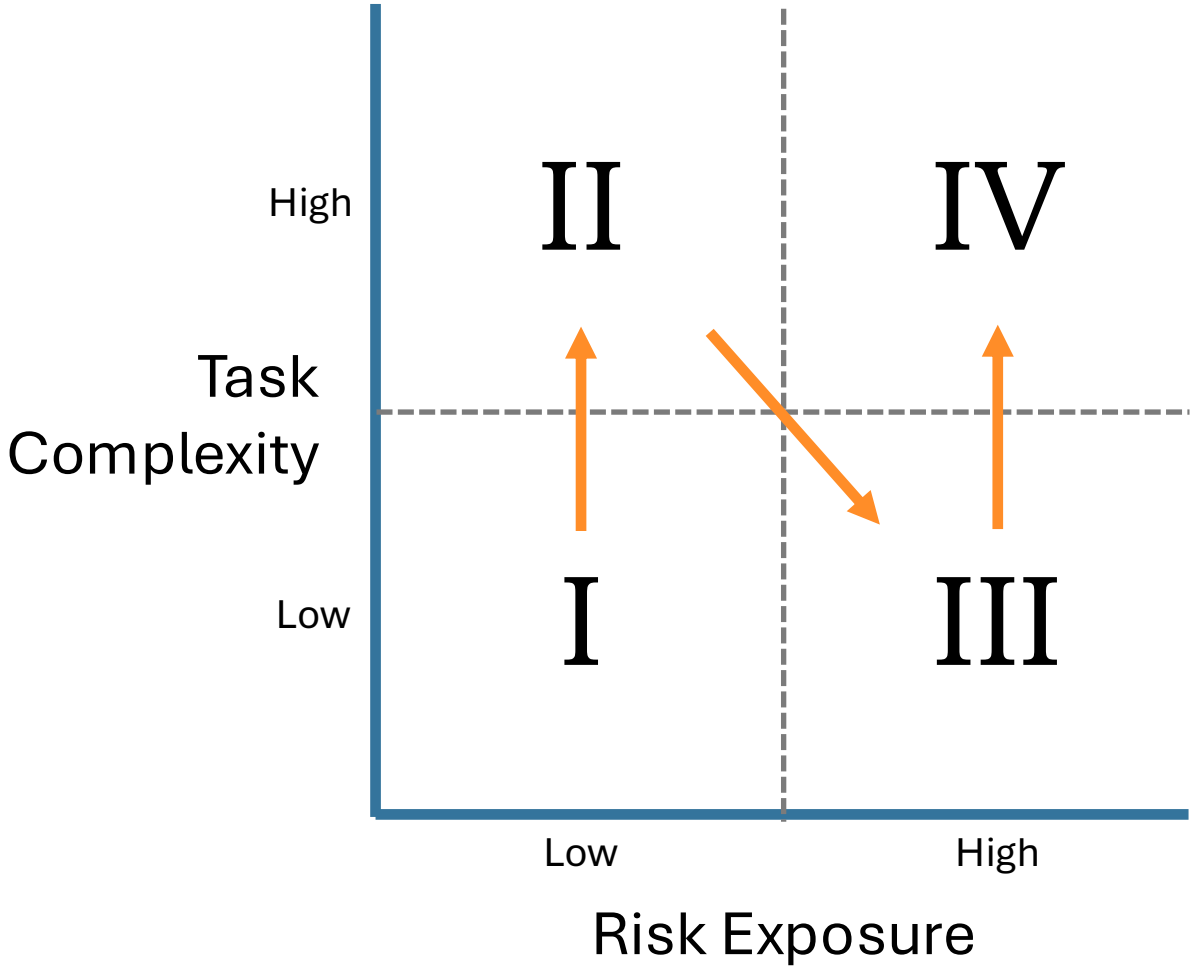

I - Clerical Agents

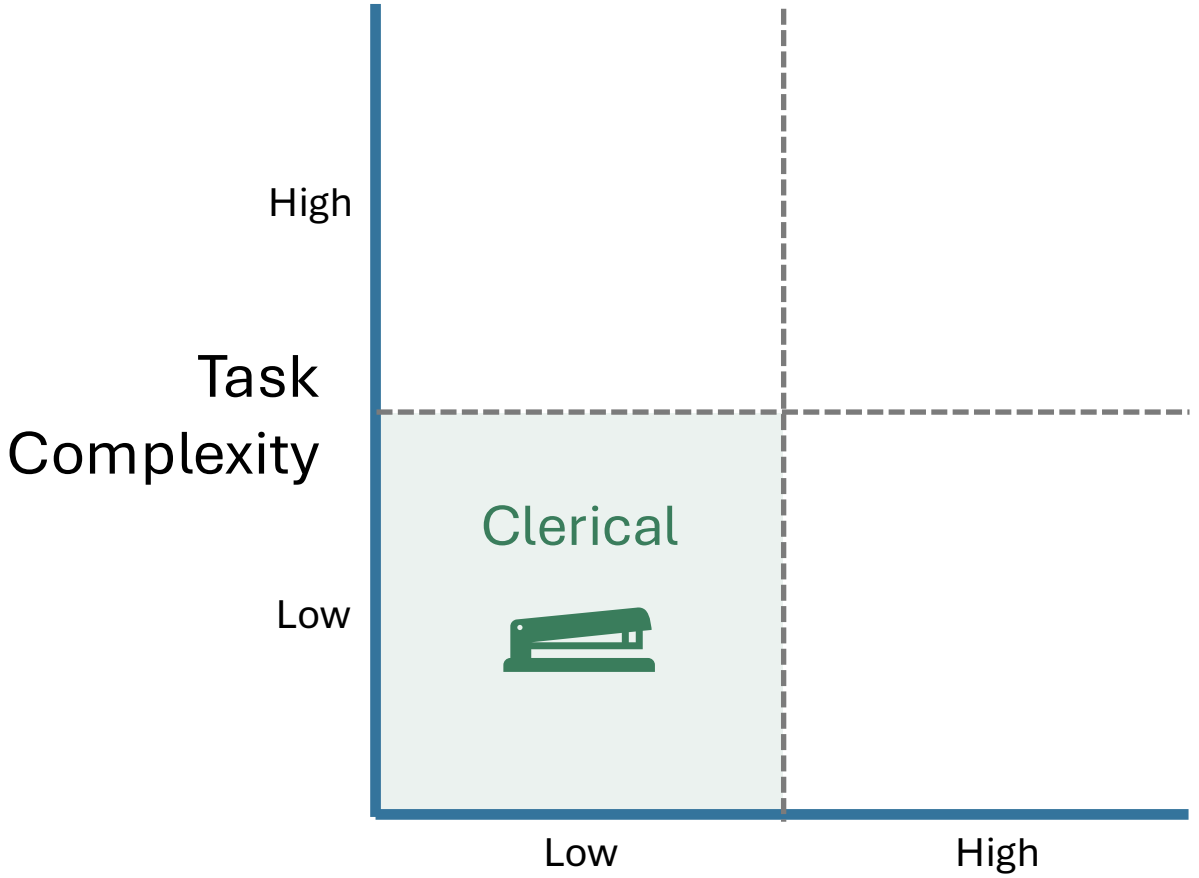

Clerical Agents

Examples

- Extract invoice data from PDFs
- Draft payment reminders
- Reconciliations
- Retrieve data
- Generate internal reports

II - Advisory Agents

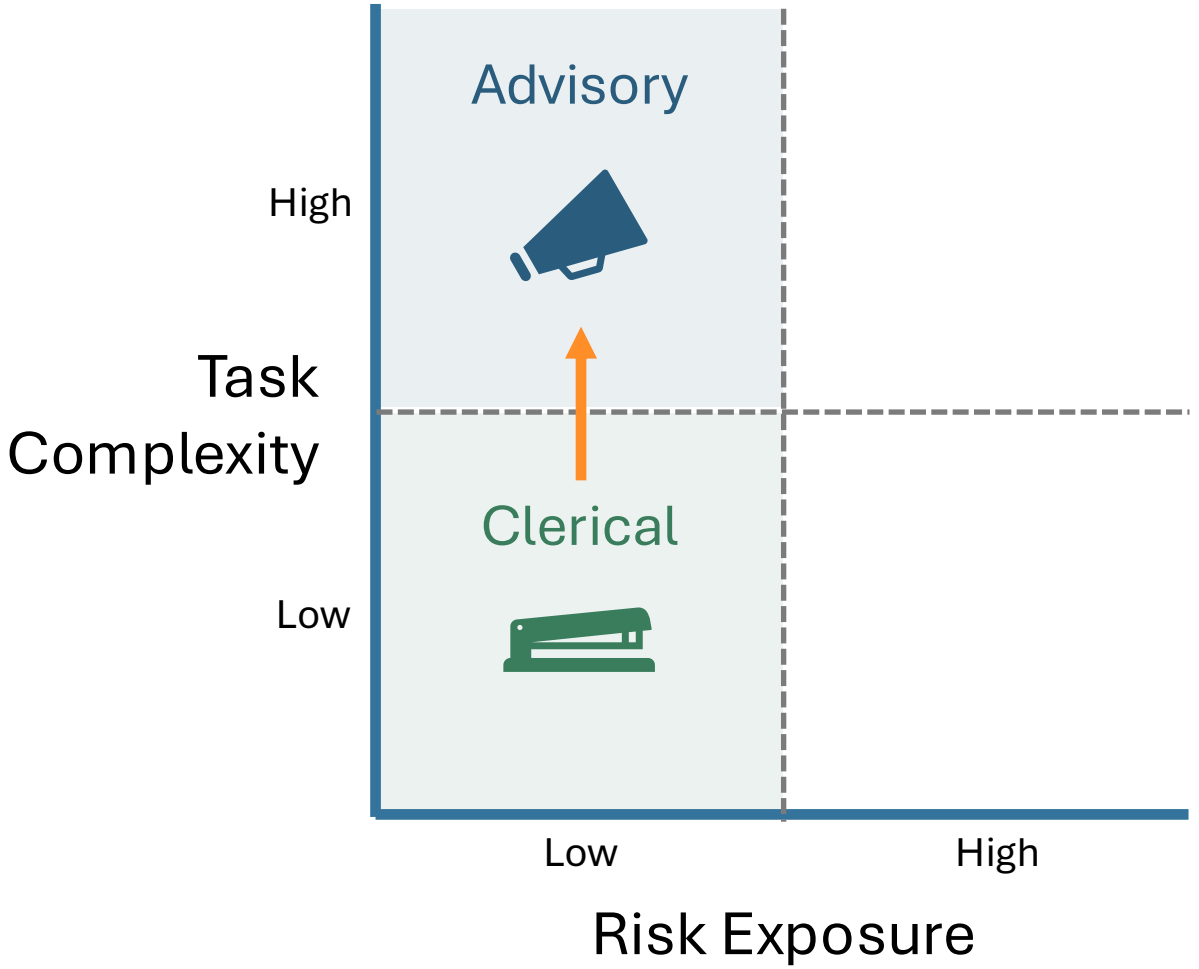

Advisory Agents

Examples

- Variance commentary
- Summarize contract terms
- Benchmark profitability
- Review expense reports
- Error, fraud and anomaly detection

III - Accountable Agents

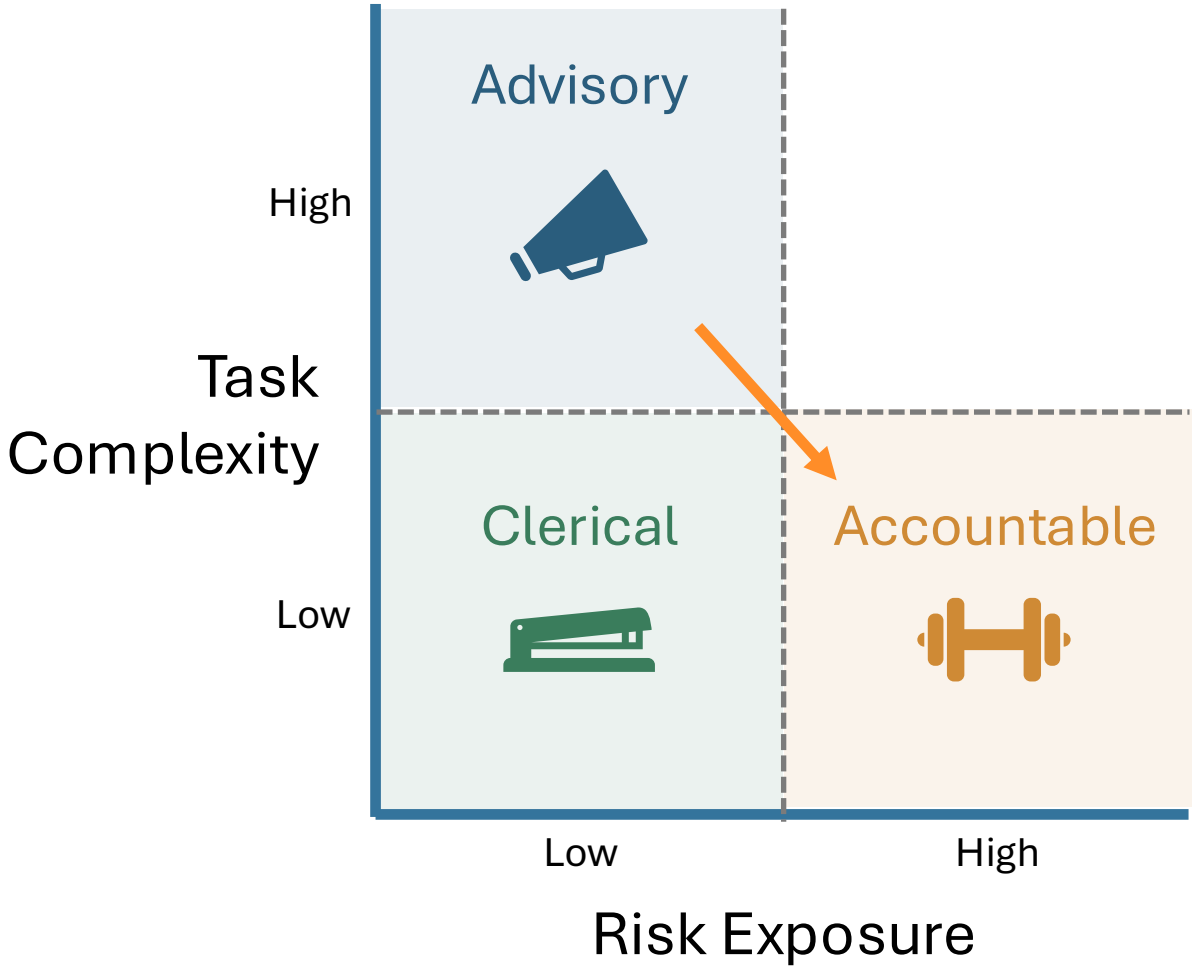

Accountable Agents

Examples

- Post journal entries
- Update master data
- File/calculate tax returns
- Reclassify GL transactions
- Process payroll

CAUTION!

Simplicity ≠ Low Risk

IV - Consequential Agents

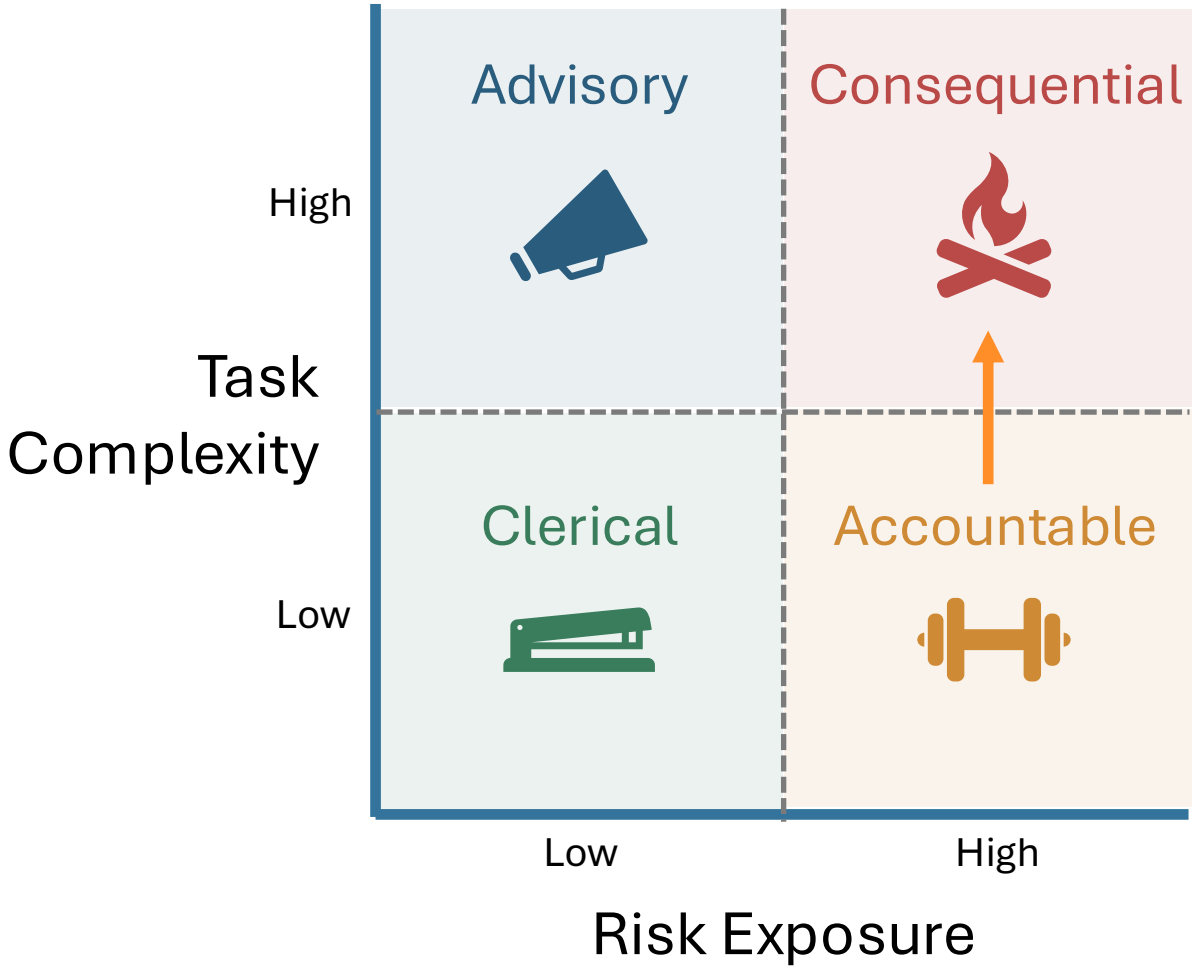

Consequential Agents

Examples

- Close automation
- Provision calculations
- Transfer pricing
- Goodwill impairment
- Revenue recognition memos
- Tax provision modeling
- M&A Support

How to Assign Agents to a Quadrant?

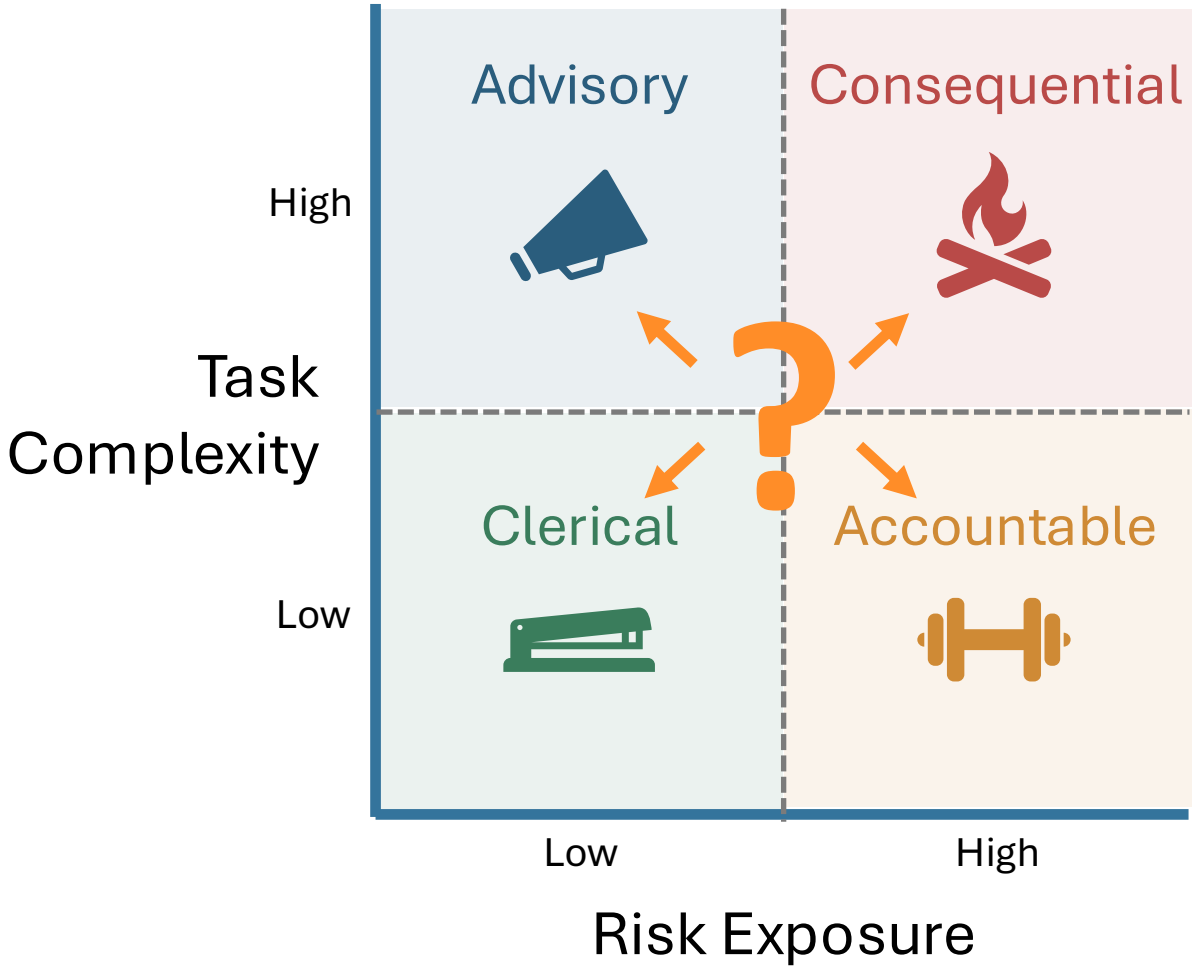

How to Assign Agents to a Quadrant?

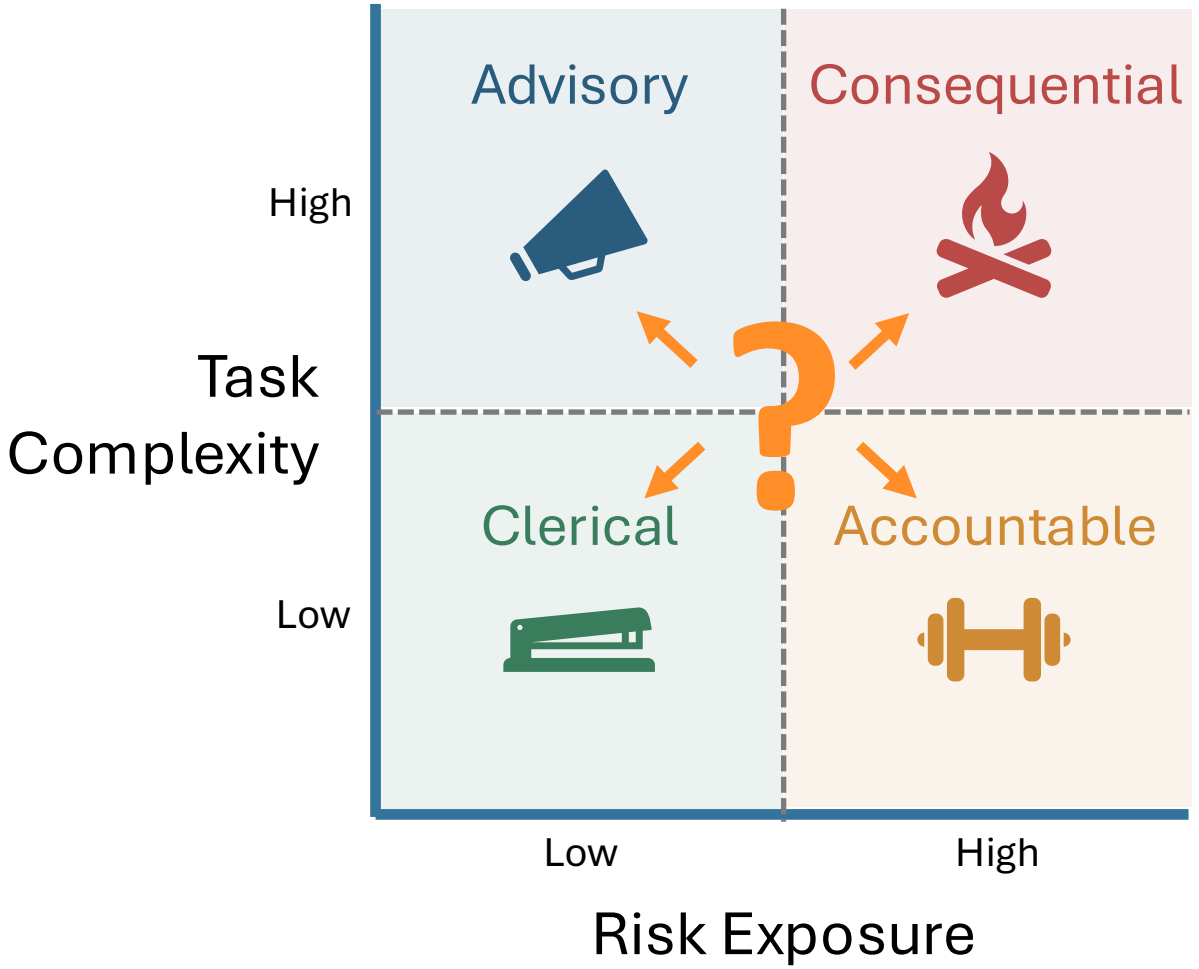

Two Part Assignment Process:

1. Identify risk exposure
2. Identify task complexity

Hard Criteria for Low Risk Exposure

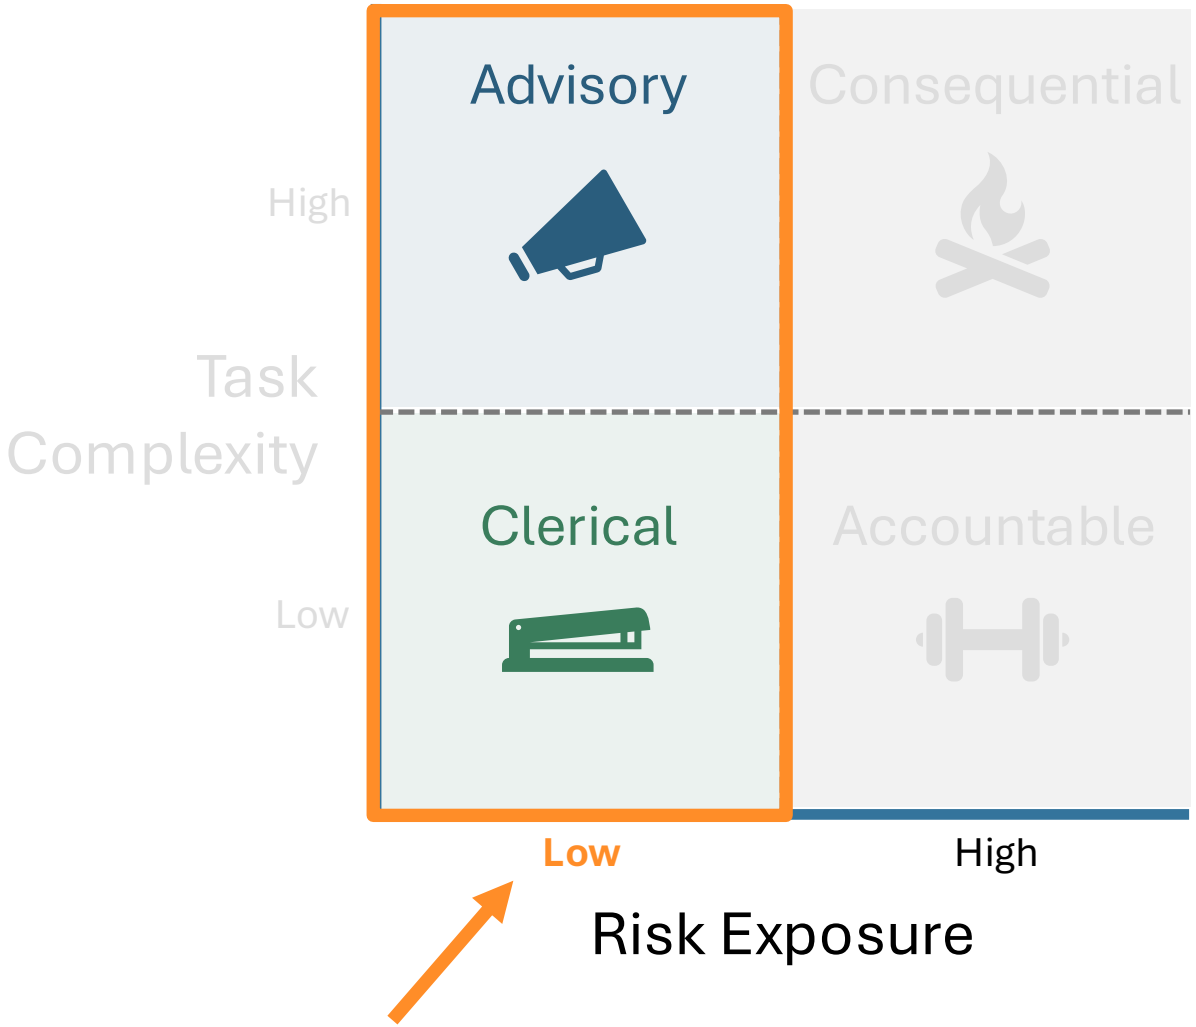

Low Exposure Criteria

ALL items are met:

1) Read only access to data
2) No confidential information
3) Shares results internally
4) Doesn't impact individuals

Hard Criteria for High Risk Exposure

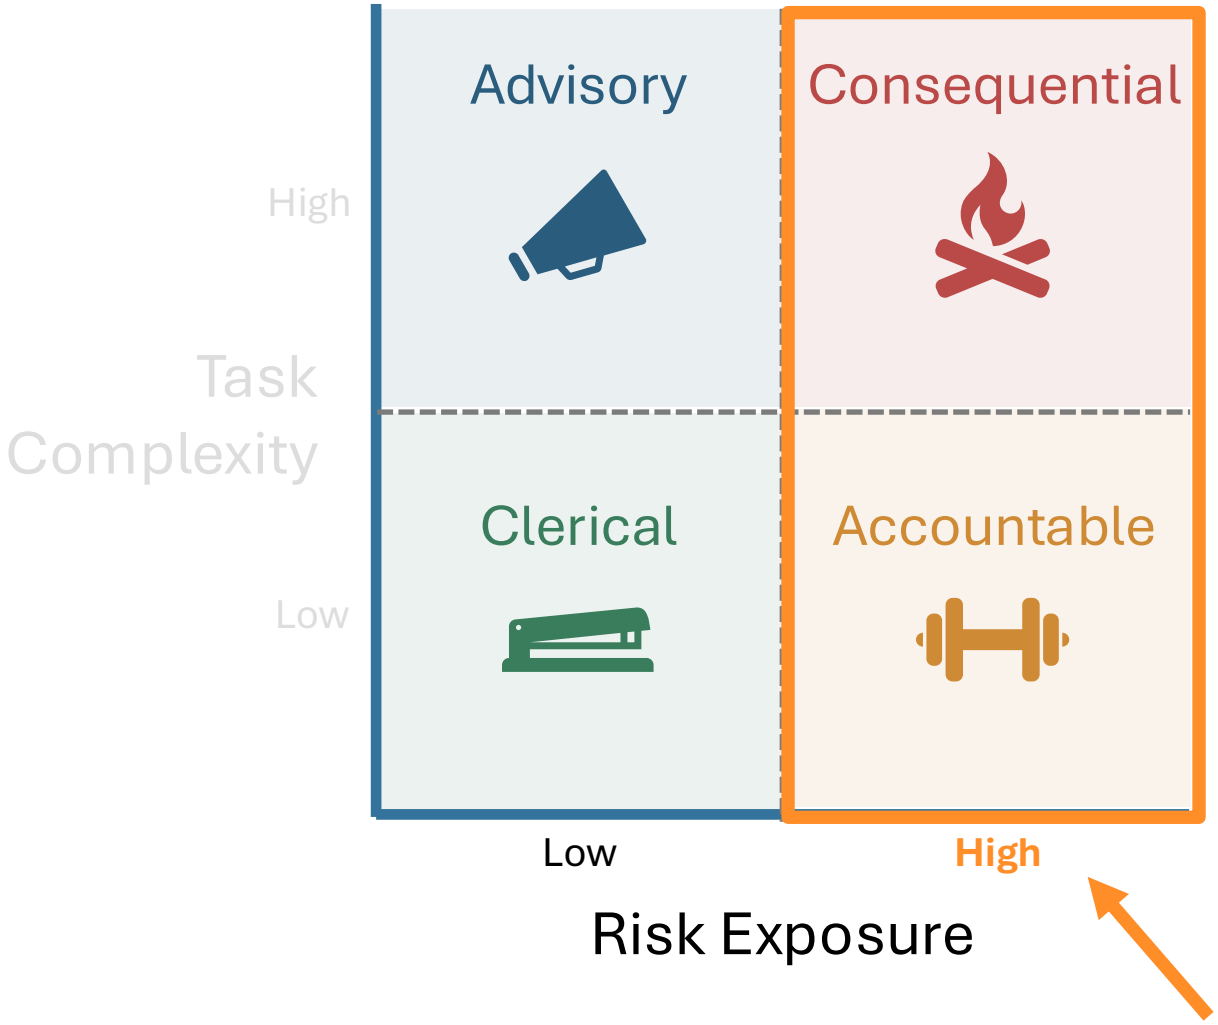

High Risk Criteria

ANY item is met:

1) Write access
2) Confidential information (e.g. – PII, HIPPA, etc.)
3) Shares results externally
4) Impacts individual person

Task Complexity Requires Judgement

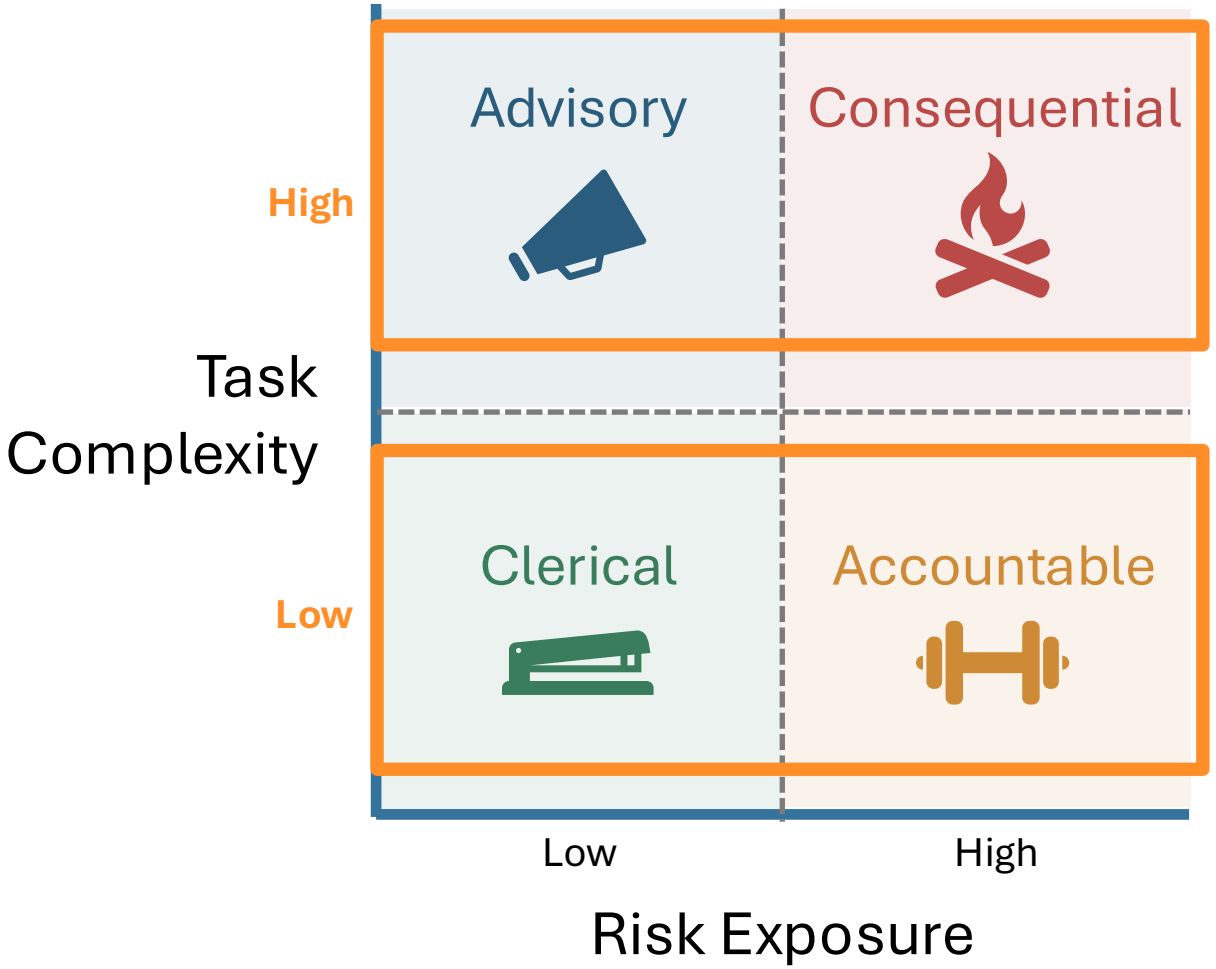

Assessing task complexity requires interpretation based on your business's circumstances

Are outputs factual or subjective?

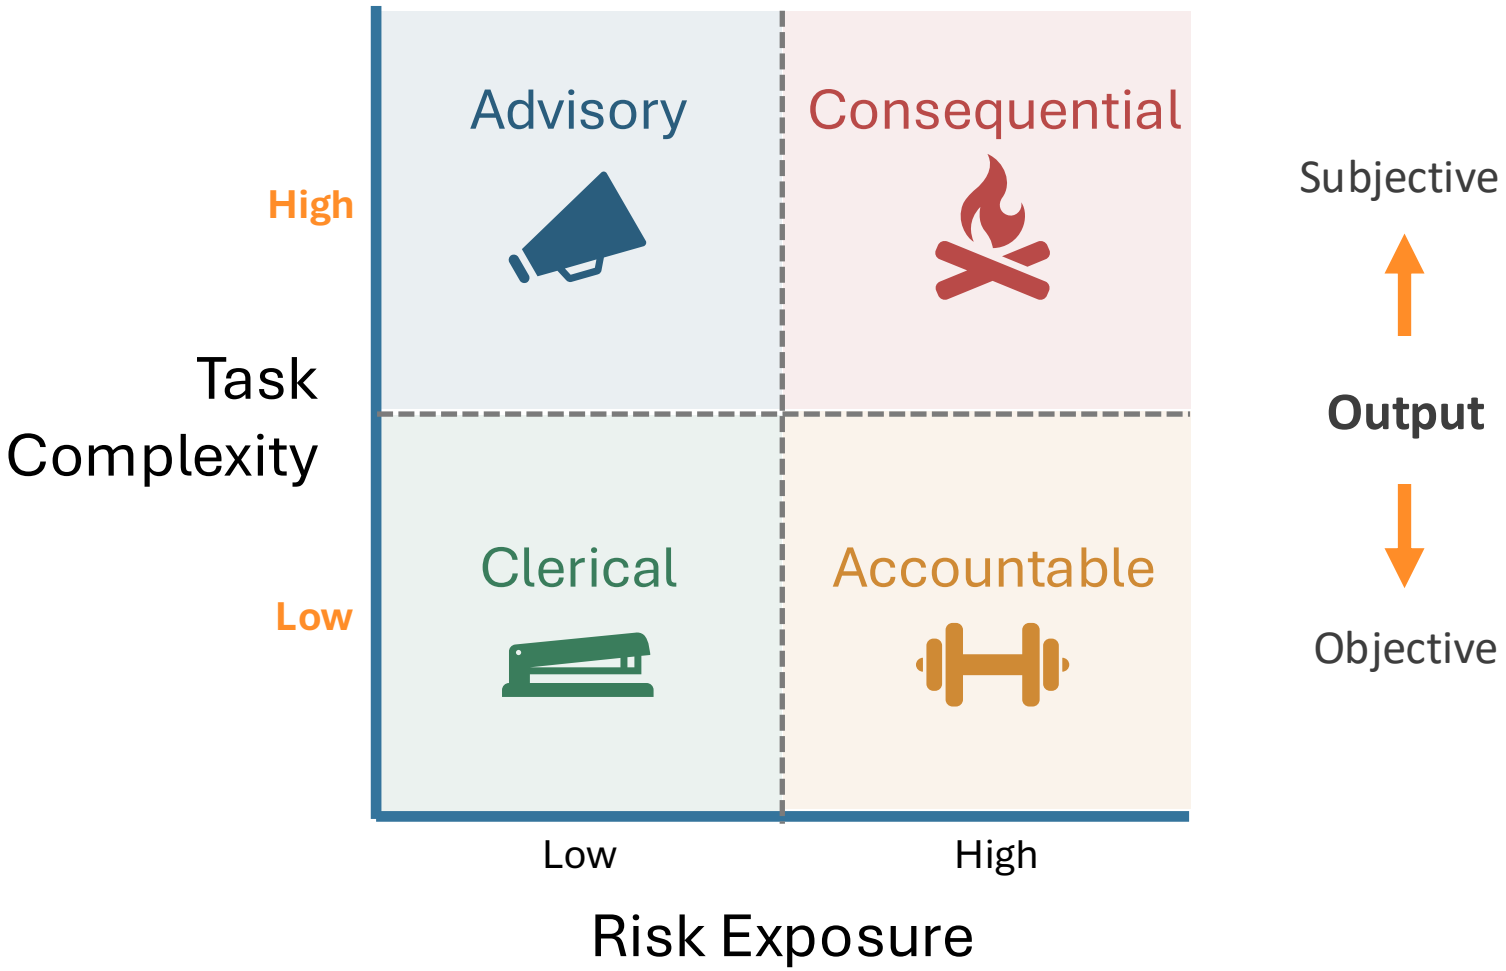

How difficult are results to verify?

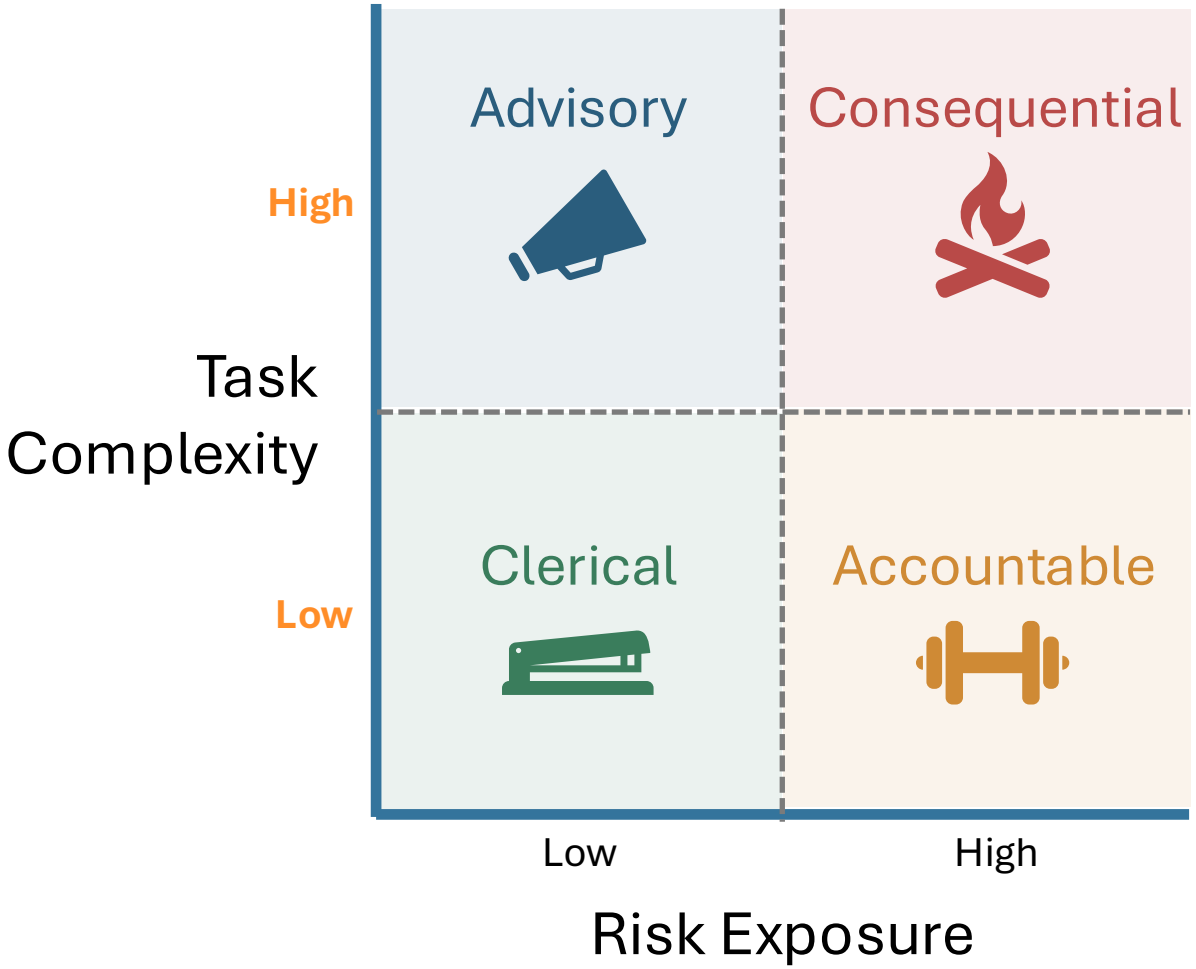

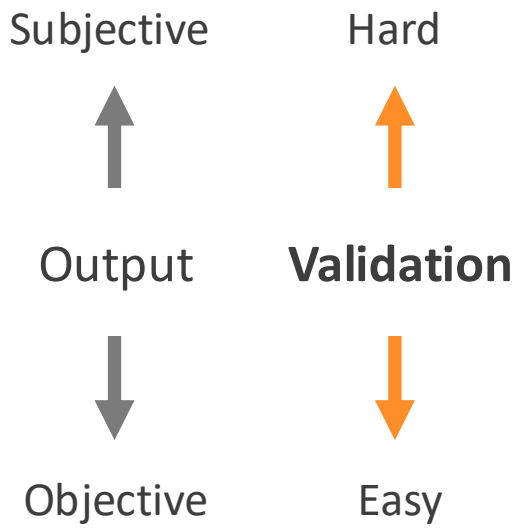

How many data sources are used?

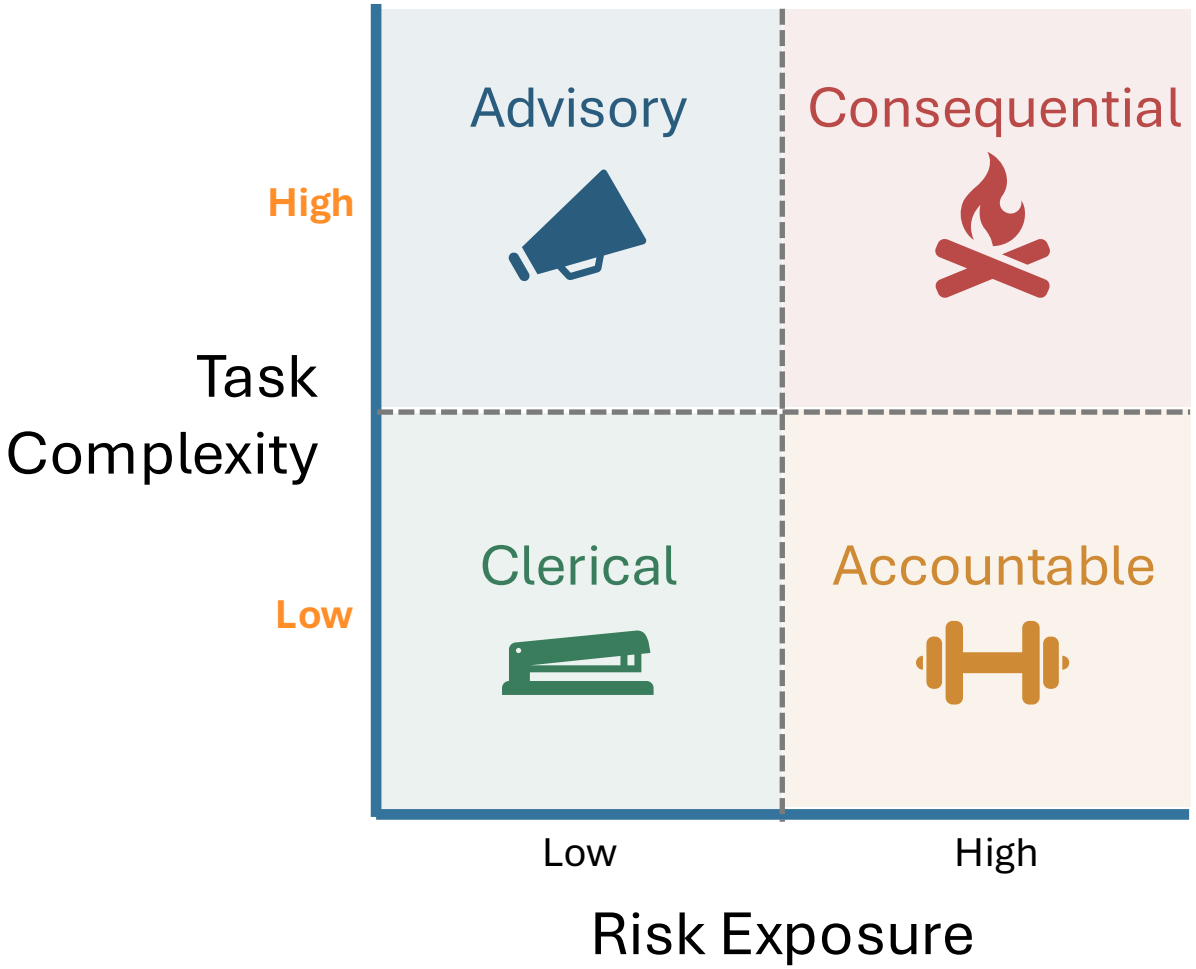

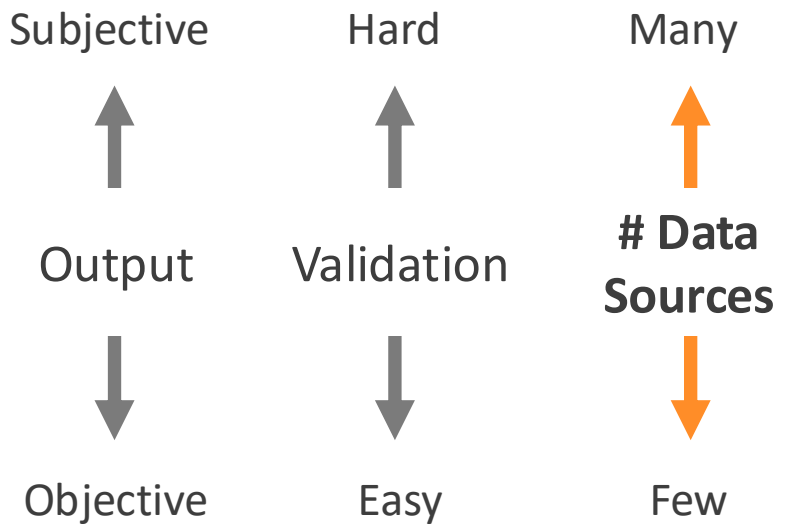

How many agents are used?

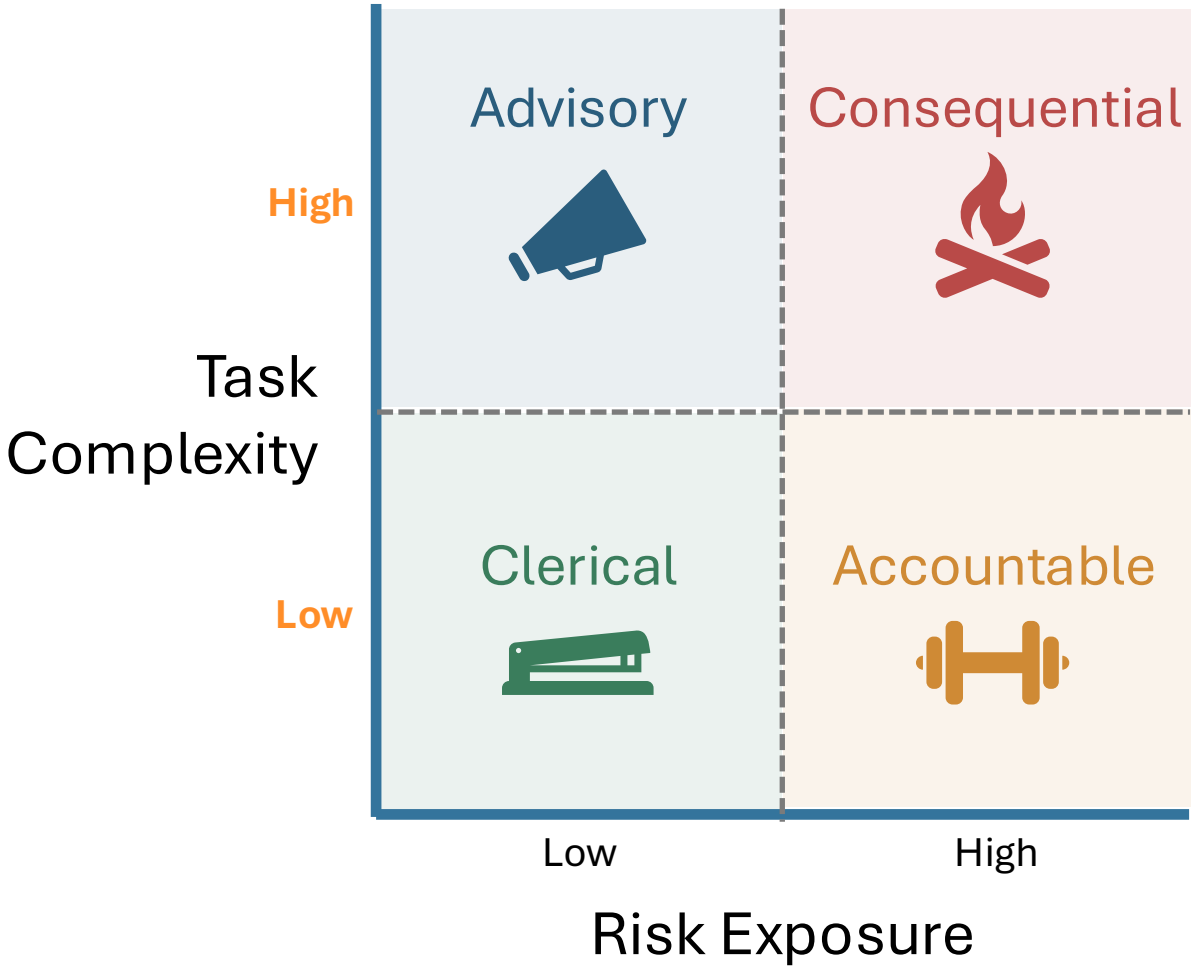

Subjective	Hard	Many	Many
↑	↑	↑	↑
Output	Validation	# Data Sources	**# of Agents**
↓	↓	↓	↓
Objective	Easy	Few	Few

AI Agent Risk Control Framework

	I – Clerical	II – Advisory	III – Accountable	IV - Consequential
RISK - Hard Criteria				
Access	Read Only		Write	
Confidential Info	No		Yes	
Data Transmission	Internal Only		External	
Individual Impact	No individuals impacted		Impact Individual Person	
COMPLEXITY - Subjective Assessment				
Output	Objective	Subjective	Objective	Subjective Interpretation
Validation	Easy	Moderate	Easy	Difficult
# Data Sources	1-2	3+	1-2	3+
# Agents	1-2	3+	1-2	3+

How do I assign controls to agents in each quadrant?

Escalating Control Requirements

I – Clerical	II – Advisory	III – Accountable	IV - Consequential
			

Controls

- Periodic sampling
- Documented system processes
- Exception-based monitoring
- Automated validation and review
- **AI literacy training**

Escalating Control Requirements

I – Clerical	II – Advisory	III – Accountable	IV - Consequential
			
Controls	+		

- Periodic sampling
- Documented system processes
- Exception-based monitoring
- Automated validation and review
- AI literacy training

- **Human review before output is used**
- Version-controlled prompts with change tracking
- Audit trail of prompts
- Data classification review to prevent inadvertently leaking confidential data
- **GenAI disclosure**

Escalating Control Requirements

I – Clerical	II – Advisory	III – Accountable	IV - Consequential
			

Controls

- Periodic sampling
- Documented system processes
- Exception-based monitoring
- Automated validation and review
- AI literacy training

+

- Human review before output is used
- Version-controlled prompts with change tracking
- Audit trail of prompts
- Data classification review to prevent inadvertently leaking confidential data
- GenAI disclosure

+

- Pre-execution human approval of prompts
- **Meaningful human review**
- **Segregation of duties between developers**
- Full audit trail
- Least-privilege access restrictions
- PII handling protocols

Escalating Control Requirements

I – Clerical	II – Advisory	III – Accountable	IV - Consequential
			
• Periodic sampling • Documented system processes • Exception-based monitoring • Automated validation and review • AI literacy training	• Human review before output is used • Version-controlled prompts with change tracking • Audit trail of prompts • Data classification review to prevent inadvertently leaking confidential data • GenAI disclosure	• Pre-execution human approval of prompts • Meaningful human review • Segregation of duties between developers • Full audit trail • Least-privilege access restrictions • PII handling protocols	• Multi-step human-in-the-loop review • **Explainability and fairness review** • Reproducibility testing with varied inputs • Secondary human expert validation • **Failover and model exit strategy**

Controls

1 Defining Agents

2 Identifying Risks and Controls

3 Moving Forward

4 Q&A

Key Takeaways

1

AI Agents Are Not Inherently High Risk

Media claims explain what agents can do, not what they are. Agents are observable and controllable.

Agent risk dependent on an agent's purpose

2

AI Agents Are Already Here

Your business users are already using agents whether managed or not. Ignoring them creates invisible risk.

Controls are necessary now

3

Simplicity does not equal safety

The agents most likely to cause harm are often the ones performing the simplest tasks in sensitive areas with privileged access.

Employ discipline

Next Steps

- Inventory AI agents currently in use
- Map agents to risk framework
- Establish an agent registration and approval policy
- Prioritize decisions about higher risk agents already in use
- Use pre-packaged platforms to support control automation
- Use framework to communicate agent exposure and governance gaps

1 Defining Agents

2 Identifying Risks and Controls

3 Moving Forward

4 Q&A

Mark D McDonald

Founder, Finance-Next LLC

Q&A

Accelerating AI for Finance and Accounting

www.finance-next.com

Mark D McDonald

Founder, Finance Next LLC

Thank You!

Accelerating AI for Finance and Accounting

www.finance-next.com