RESEARCH PRESENTATION

A Balanced Enterprise Architecture Framework for Secure and Ethical AI Adoption

A framework for governable AI adoption in enterprise settings

The problem

AI pilots succeed locally, but enterprise adoption stalls between pilot and production.

The framework

Balanced layers linking business value, technology readiness, developer enablement and user trust.

The contribution

Security, ethics and governance embedded by design, with guidance for research and practice.

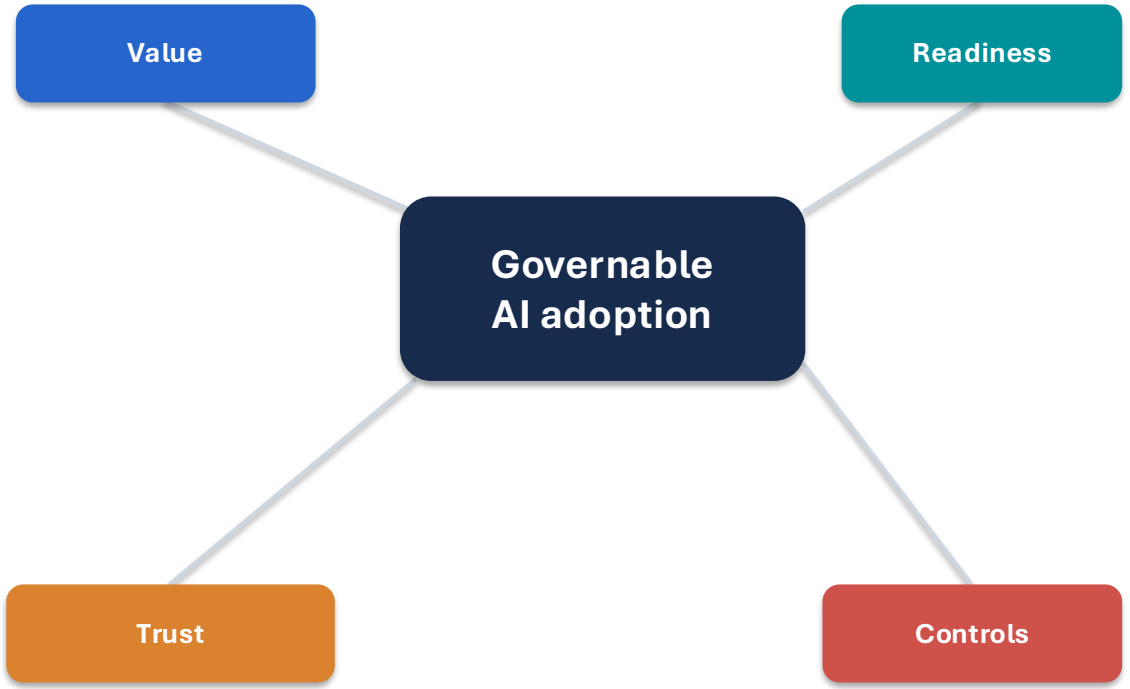

Bhagyeshkumar Chokhawala
Conference presentation

Integrating business value, technology readiness, developer enablement, user trust, security, and ethics.

Research motivation and adoption gap

AI pilots often succeed locally, while enterprise-scale adoption remains fragmented, costly, and risky.

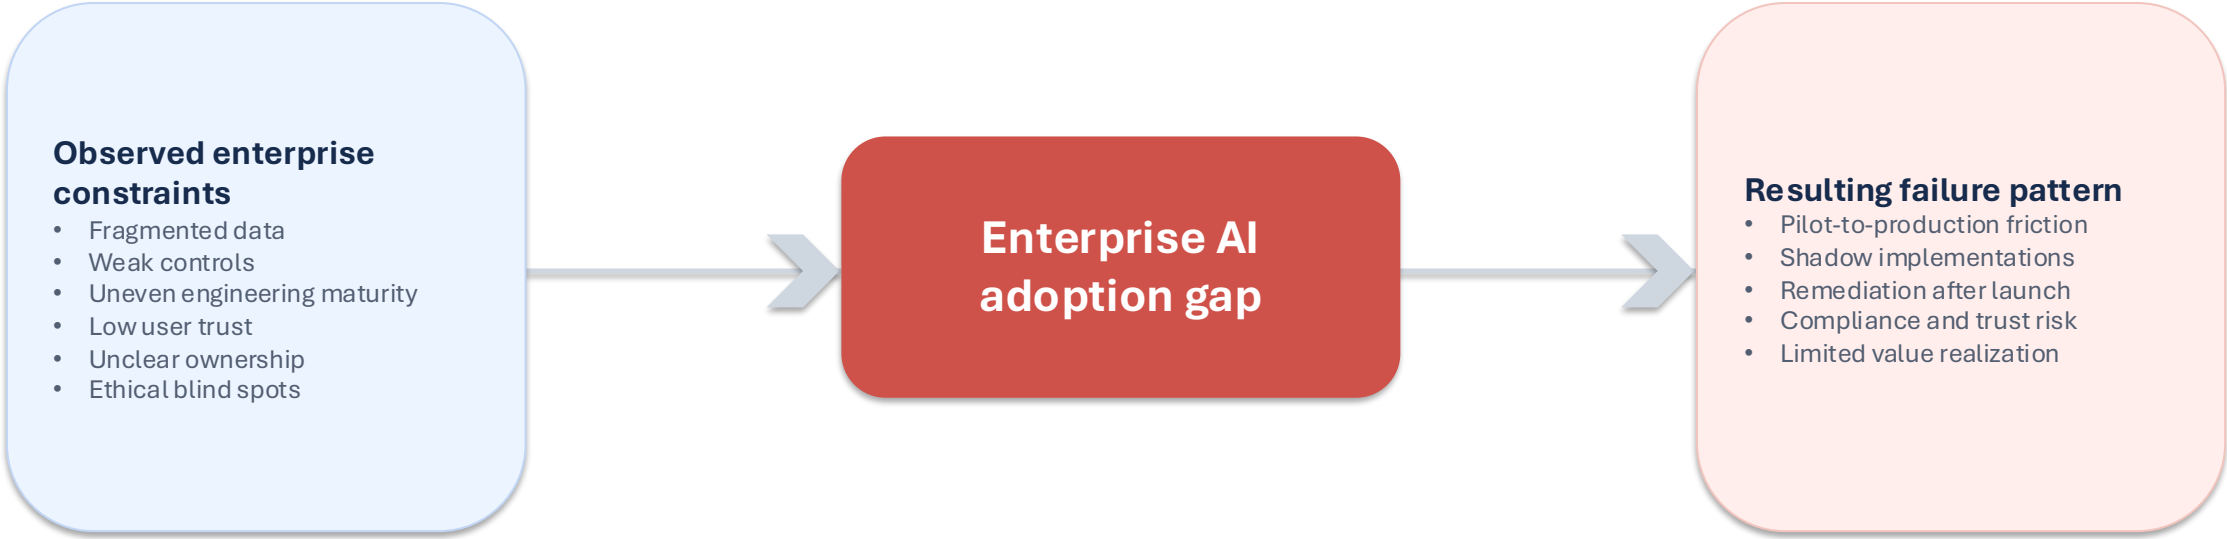

Core premise: adoption maturity depends on architectural balance, not model availability alone.

RESEARCH OBJECTIVE

Guiding research question and design objective

The paper develops a conceptual architecture artifact for structured, secure, and ethical AI adoption.

How can enterprises architect AI adoption so business value, technology readiness, developer enablement, user trust, security, and ethics scale together?

1. **Synthesize architecture and AI governance literature**
2. **Define a balanced enterprise AI framework**
3. **Operationalize security and ethics as controls**
4. **Provide execution artifacts and roadmap**

Methodological positioning: conceptual synthesis + design-oriented framework development

Theoretical and practice foundations

The framework integrates enterprise architecture, AI risk governance, MLOps, human-AI interaction, DevEx, and stakeholder theory.

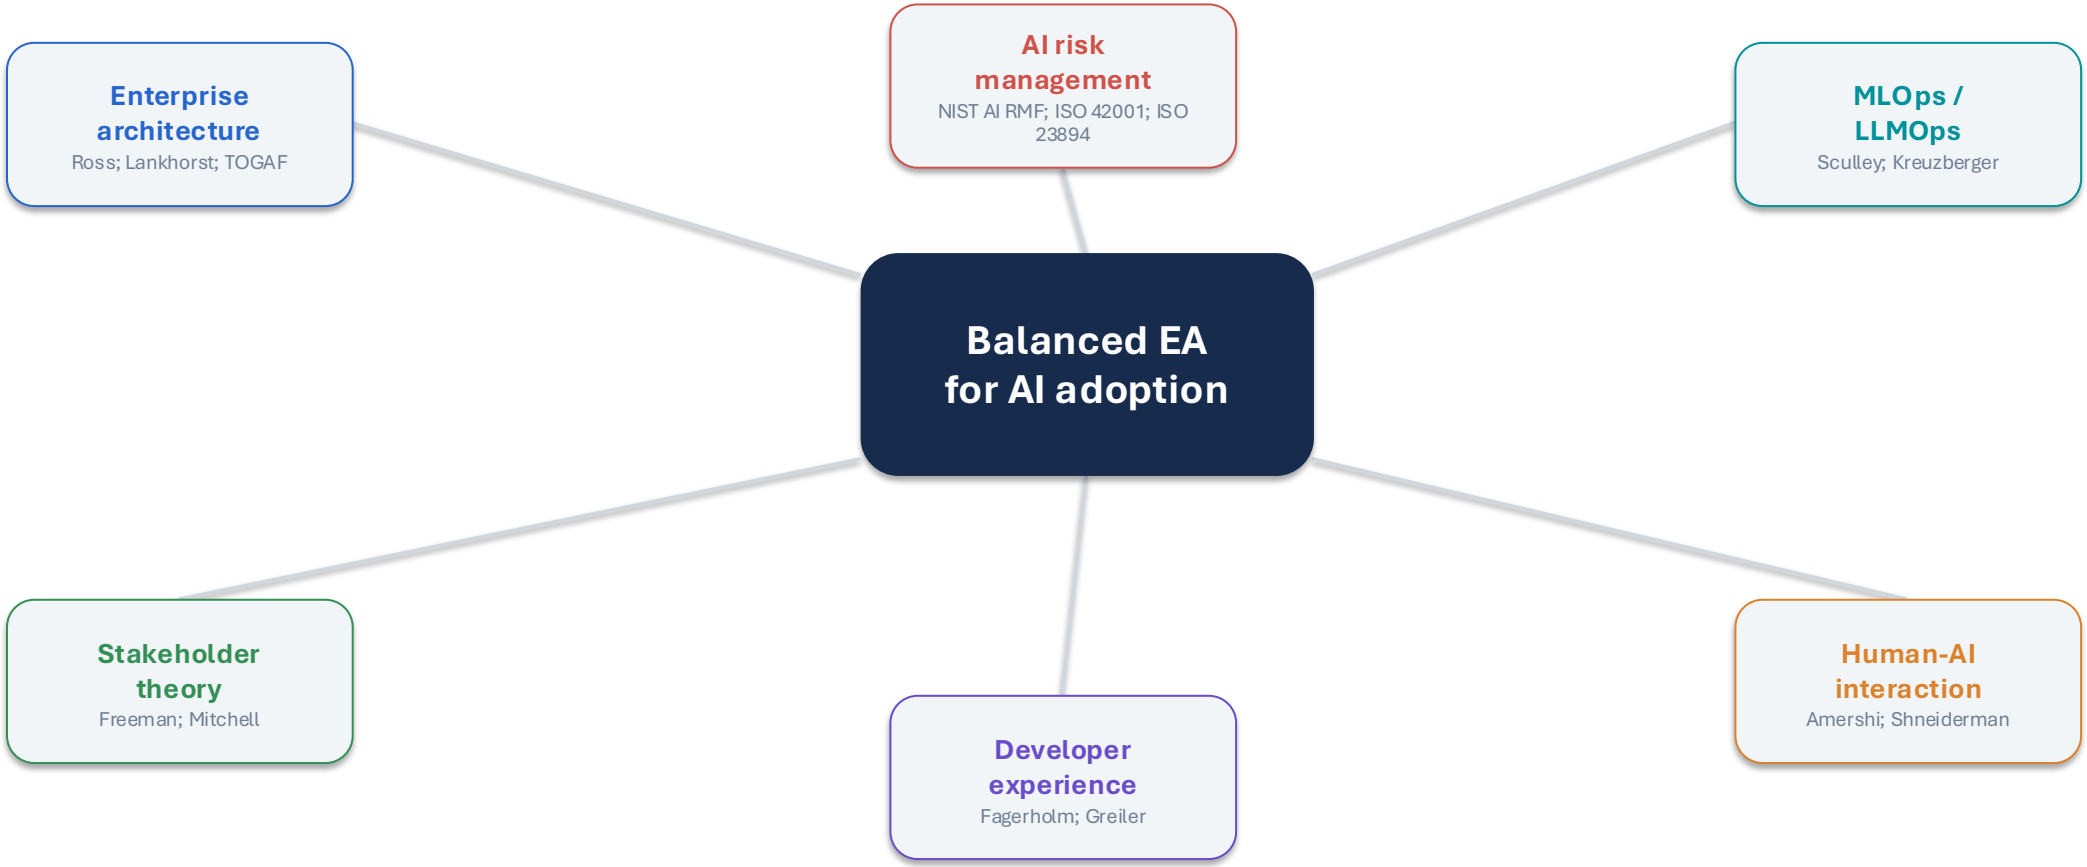

Problem model

AI adoption outcomes depend on the interaction between enterprise context, architectural balance, and governance evidence.

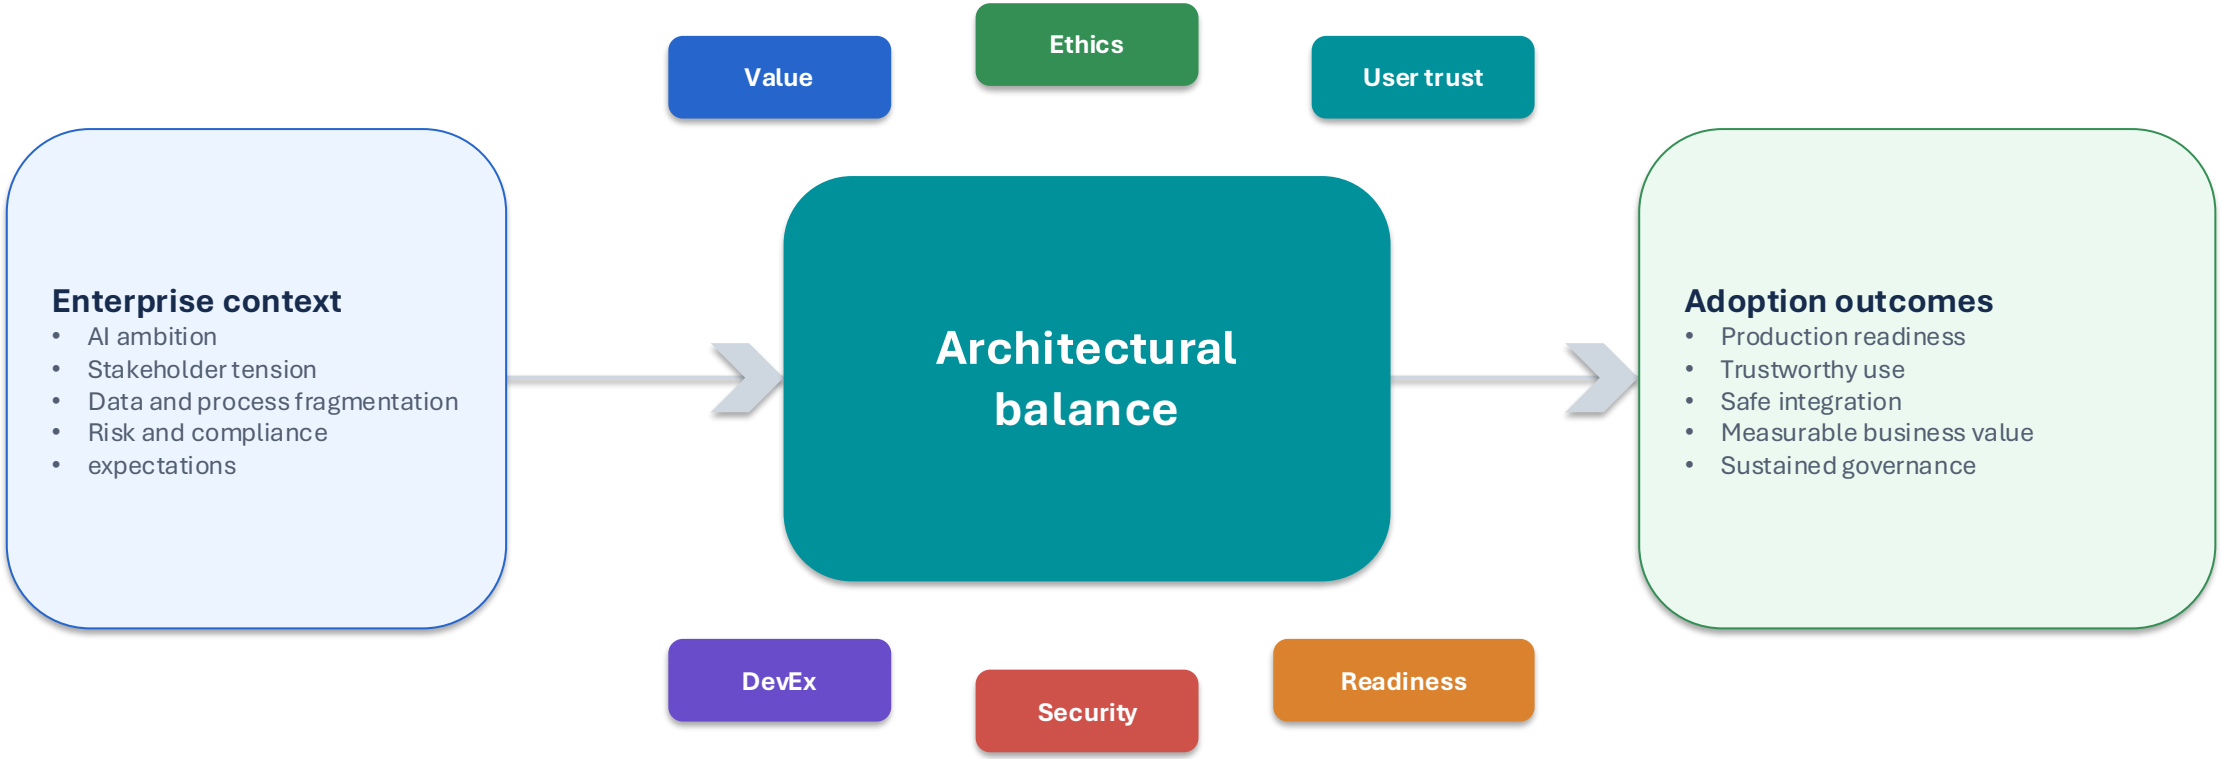

The proposed 4+2 balanced architecture framework

Four adoption perspectives are coordinated by governance and operating model, with security and ethics embedded across both.

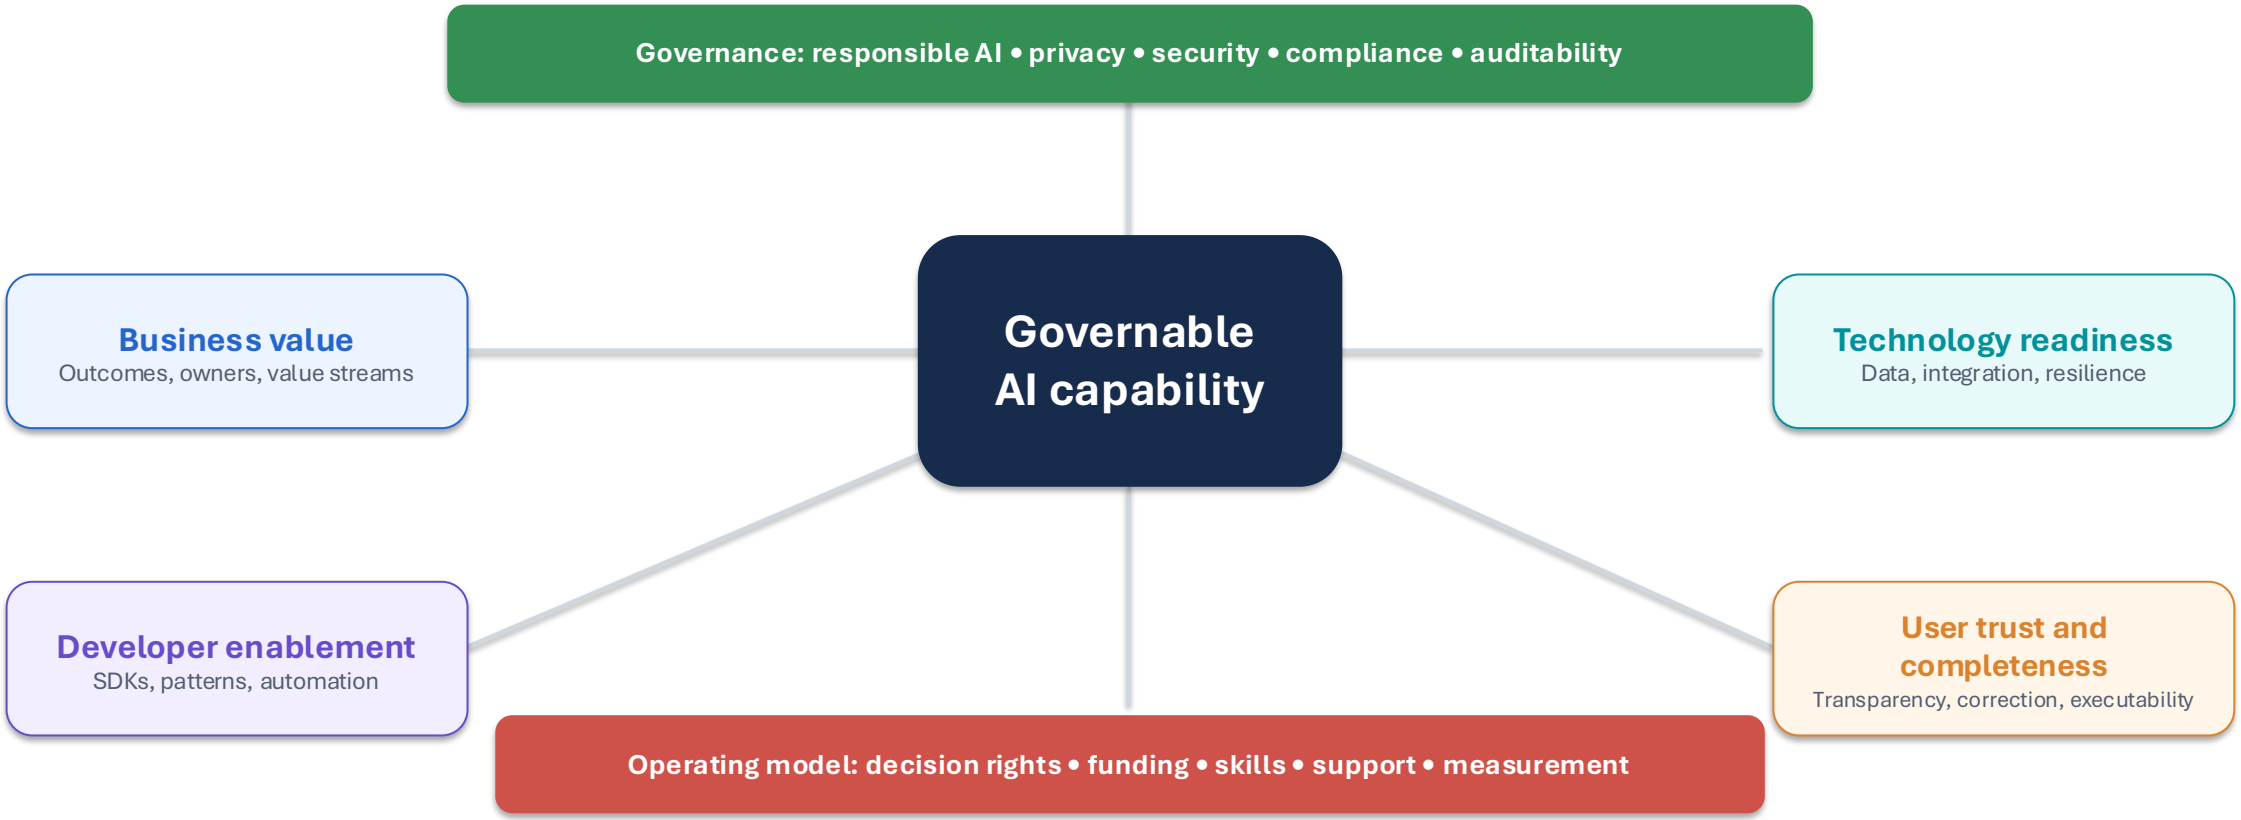

Framework construct definitions

Research-style interpretation of the framework as a set of architecture constructs and observable indicators.

Construct	Definition	Observable indicators
Business value	Degree to which AI use cases align with measurable enterprise outcomes	Capability owner; value thesis; KPI linkage
Technology readiness	Availability of data, integration, platform, and operational foundations	Trusted data; model access; observability; resilience
Developer enablement	Ability of delivery teams to implement governed AI patterns repeatedly	SDKs; templates; tests; reference patterns
User trust and completeness	Ability of AI-enabled journeys to remain understandable, correctable, and executable	Transparency; override; feedback; downstream integrity
Security and ethics	Controls that prevent misuse, harm, privacy leakage, or unfair outcomes	Least privilege; guardrails; bias review; accountability

Security and ethics as embedded design constructs

The paper treats security and ethics as architectural qualities that must be designed before production scale.

Security controls

Identity and access
Least privilege, model gateway permissions

Data protection
Classification, privacy, encryption, retention

Interaction defense
Prompt injection controls, retrieval guardrails

Action safety
Tool authorization, policy checks, rollback

Operational assurance
Logging, monitoring, incident response

Ethical AI controls

Fairness
Bias assessment and impact review

Transparency
Explainability, capability limits, provenance

Human agency
Oversight, escalation, appeal

Inclusion
Accessibility and responsible design

Accountability
Owner, audit trail, redress pathway

Trustworthy enterprise AI

Architecture design principles derived from the framework

The principles convert the conceptual model into rules that guide design, review, and implementation.

1	Anchor AI in business capabilities	Avoid technology-led pilots without accountable value owners.
2	Separate engagement, intelligence, and record systems	Prevent AI interfaces from bypassing transactional integrity.
3	Make the governed path the easiest path	Provide reusable SDKs, prompts, policies, tests, and runbooks.
4	Treat security and ethics as non-functional requirements	Design for privacy, safety, fairness, accountability, and auditability.
5	Scale only after evidence is available	Use exit criteria, operational evidence, and value review before rollout.

TOGAF-style capability map for AI adoption

The framework is operationalized through capabilities that can be assessed, funded, governed, and matured.

Strategy & Governance	Business Architecture	Data Architecture	Application Architecture	Technology Architecture	Delivery & DevEx	Operations & Service
AI vision Portfolio management Risk and compliance	Capability alignment Process transformation Decision intelligence	Data readiness Knowledge management Provenance	AI-enabled apps Workflow orchestration Human-in-loop	Model gateways Vector stores Security operations	SDKs and templates Standards Evaluation and LLMOps	Runbooks KPI review Continuous assurance

Research interpretation: capability maturity becomes the bridge between AI strategy and implementation evidence.

GOVERNANCE PROCESS

Risk-based governance and evidence production

Governance gates should vary by AI risk and should produce evidence, not merely approval records.

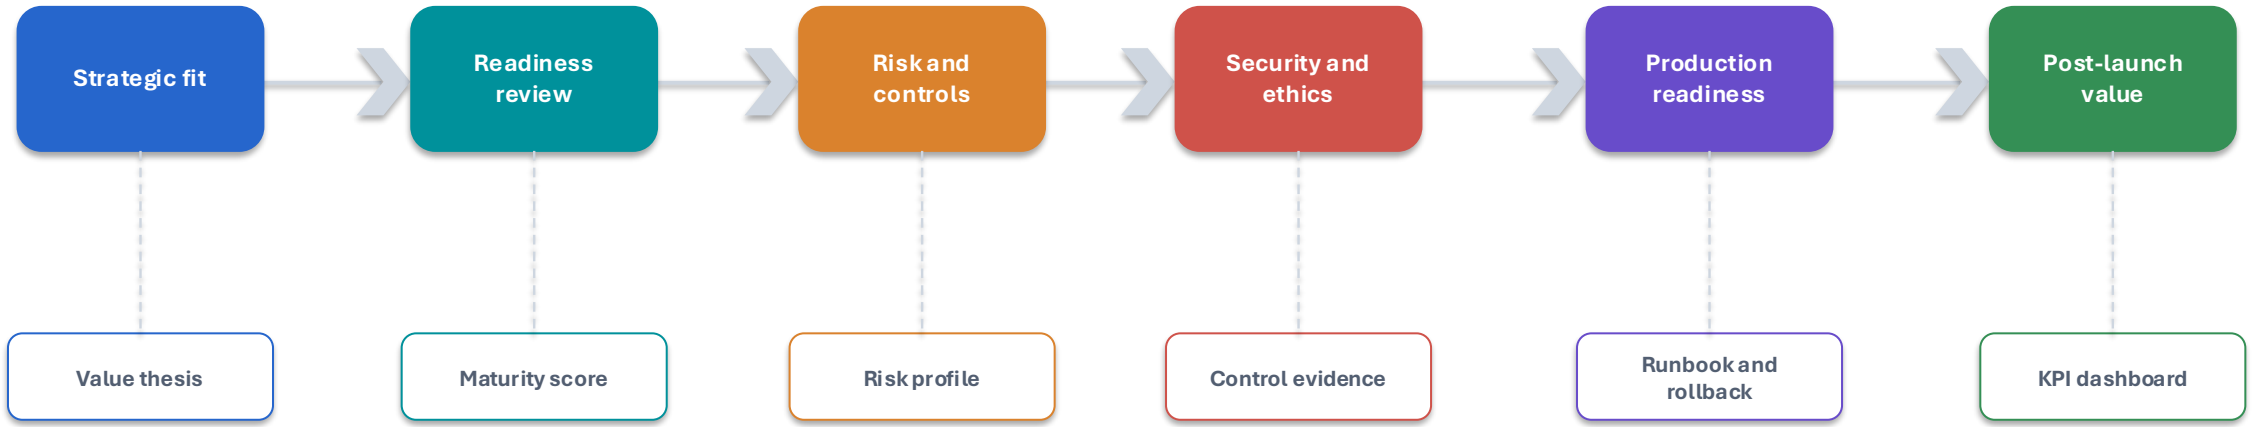

Risk basis: data sensitivity, autonomy of action, user impact, integration depth, regulatory exposure, and operational criticality.

Phased roadmap and scale conditions

The roadmap treats guardrails, operational support, trust evidence, and security assurance as prerequisites for broad scale-out.

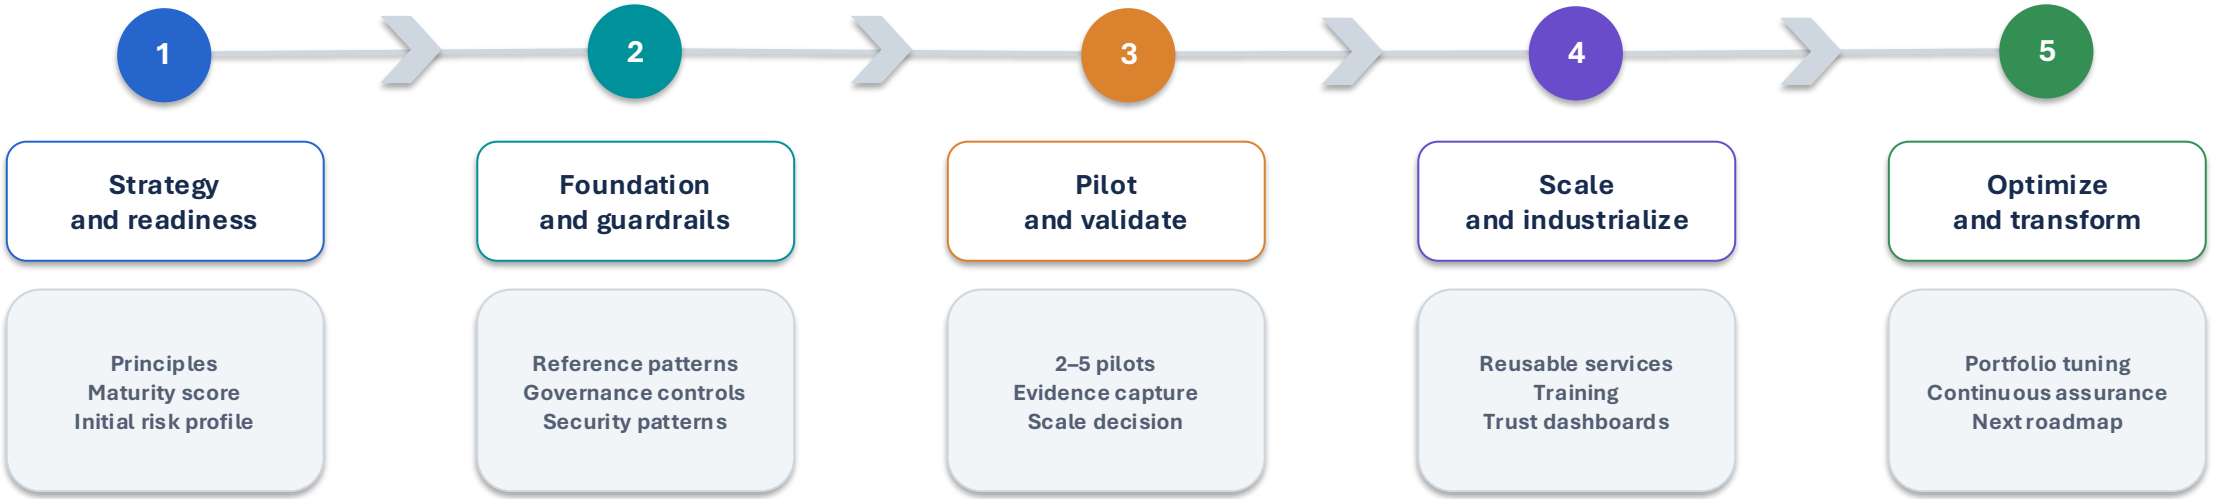

Exit criteria before broad scale-out: governance evidence • observability • user trust • support model • security assurance • ethical impact review

EVALUATION LENS

Suggested evaluation and evidence plan

Although conceptual, the framework implies measurable evaluation dimensions for future empirical validation.

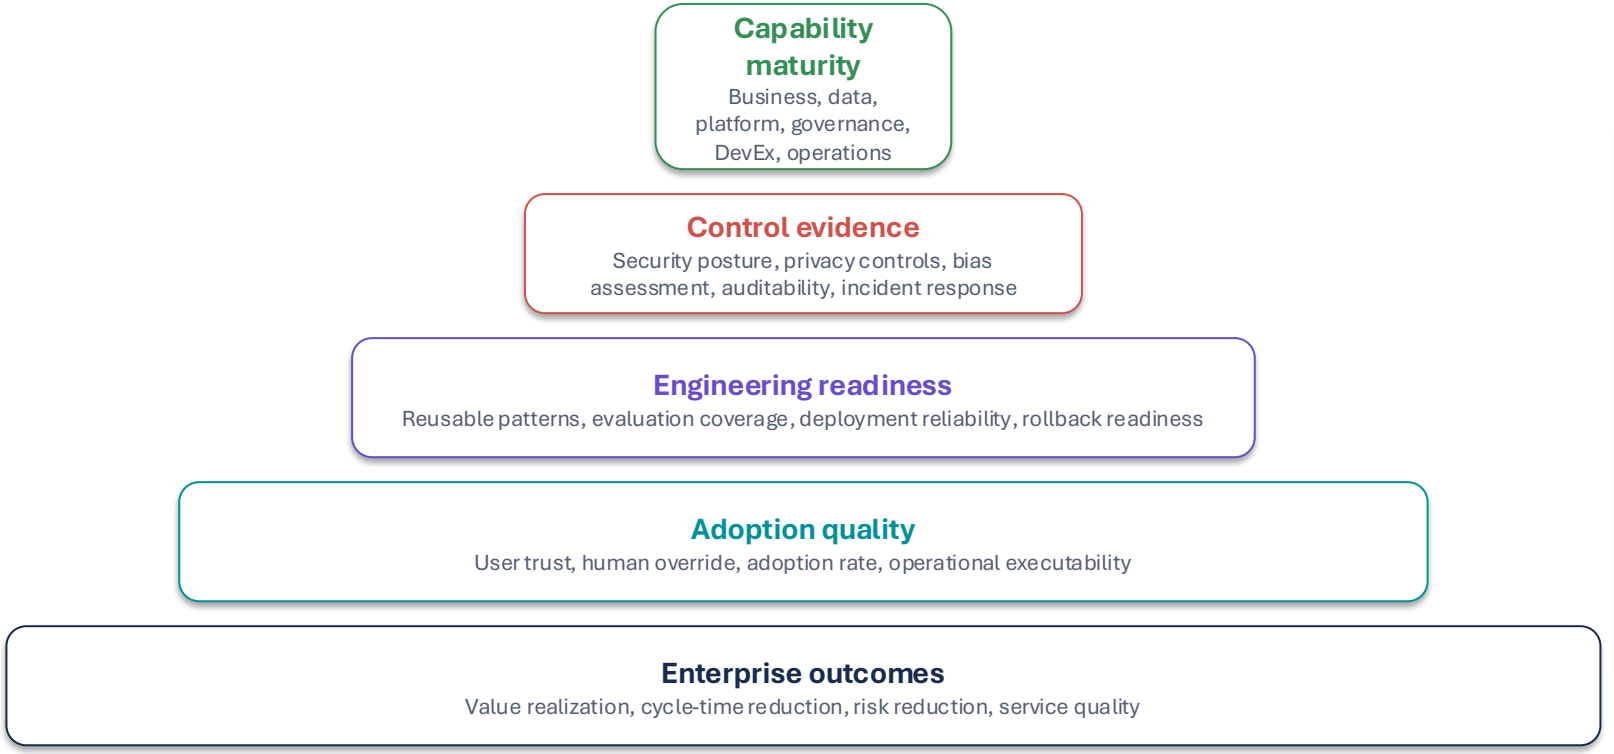

Future validation options

- Case study comparison
- Expert panel validation
- Maturity scoring
- Pilot outcome review
- Pre/post adoption survey

Research and practice contributions

The paper contributes a structured framework, operational toolkit, and governance logic for enterprise AI adoption.

Research contributions

- **C1: Frames enterprise AI adoption as a socio-technical architecture problem**
- **C2: Defines a 4+2 framework integrating value, readiness, enablement, trust, governance, and operating model**
- **C3: Embeds security and ethics as architecture constructs rather than downstream compliance checks**
- **C4: Connects stakeholder conflict resolution to AI architecture governance**

Practice implications

- **P1: Gives architects a capability map for portfolio planning**
- **P2: Provides governance gates that produce operational evidence**
- **P3: Supports reusable implementation assets and governed delivery paths**
- **P4: Defines scale conditions to reduce pilot-to-production failure**

Overall contribution: a traceable path from AI ambition to governable enterprise implementation.

Limitations and future research

The framework is intended as a conceptual architecture artifact and should be empirically validated across organizational contexts.

Limitations

- Conceptual framework rather than empirical causal model
- Requires contextual tailoring by industry and regulatory setting
- Does not prescribe a single AI platform or implementation stack
- Maturity levels require calibration before cross-enterprise comparison

Future research

- Validate through case studies and expert panels
- Develop a scored maturity instrument
- Compare adoption outcomes across governance models
- Study trade-offs between speed, control, trust, and cost

Conclusion: enterprise AI maturity is defined by alignment among architecture, controls, engineering systems, security, ethics, and human outcomes.

Thank You

Research Profile QR Code

Scan the QR code to open my research profile, publications, Google Scholar, ORCID, and professional links.