

ENTREPRENEURIAL AGENTIC AI: A CONCEPTUAL FRAMEWORK FOR AUTONOMOUS DECISION-MAKING, ETHICAL GOVERNANCE, AND HUMAN-AI COLLABORATION IN EARLY- STAGE VENTURES

Kavita Dahiya

Assistant Professor

Centre for Distance and Online Education, Manipal University Jaipur, Jaipur

Shubham Alria

Bachelor of technology in Computer Science and Engineering, CGC University, Mohali

ABSTRACT

Artificial intelligence (AI) is increasingly embedded in entrepreneurial activity, but the emergence of agentic AI—systems capable of pursuing goals through multi-step planning, tool use, and partially autonomous execution—creates a distinct set of opportunities and governance challenges for early-stage ventures. This conceptual paper examines the relationship between entrepreneurial agentic AI, autonomous decision-making, ethical governance, and human-AI collaboration. Building on research on AI and entrepreneurship, the automation-augmentation paradox, responsible AI, and entrepreneurial cognition, the paper develops an integrative framework in which agentic AI can extend entrepreneurial analytical capacity while human judgment remains essential for contextual interpretation, value selection, accountability, and exception handling. The framework identifies three interrelated pathways. First, agentic AI can improve the timeliness, consistency, and information-processing capacity of entrepreneurial decisions when task boundaries and escalation rules are clearly specified. Second, ethical governance—including accountability, transparency, privacy, fairness, monitoring, and risk controls—conditions whether increased autonomy produces legitimate and socially responsible outcomes. Third, human-AI collaboration can support innovation and operational effectiveness when trust is calibrated rather than assumed and when human and AI roles are explicitly negotiated. The paper contributes a conceptual model and three propositions for future empirical testing. It also provides managerial and policy implications for early-stage ventures seeking to adopt agentic AI without transferring responsibility from human decision-makers to machines.

Keywords: agentic AI; entrepreneurship; autonomous decision-making; responsible AI; human-AI collaboration; AI governance; early-stage ventures

1. INTRODUCTION

Artificial intelligence is changing how entrepreneurial opportunities are identified, evaluated, and exploited. Earlier work on AI and entrepreneurship emphasized automation and augmentation of tasks such as idea generation, selling, scaling, and organizational design. More recent developments in AI agents extend this discussion because systems can increasingly coordinate multiple steps, interact with tools, retrieve information, generate outputs, and execute bounded actions with limited human intervention. This creates a shift from AI as a passive decision-support tool toward AI as an active participant in entrepreneurial workflows.

The distinction matters particularly for early-stage ventures. Startups commonly operate with limited capital, small teams, incomplete information, and high environmental uncertainty.

These conditions make computational assistance attractive, but they also magnify the consequences of poorly specified objectives, biased data, automation errors, privacy failures, and unclear accountability. Research on AI and entrepreneurship already shows that AI can augment entrepreneurial activities while simultaneously creating new organizational and social liabilities (Chalmers et al., 2021). The automation–augmentation literature similarly warns that automation and augmentation are interdependent rather than mutually exclusive strategies (Raisch & Krakowski, 2021).

This paper therefore treats agentic AI as a socio-technical capability rather than as an autonomous substitute for the entrepreneur. The central question is not whether AI can make decisions independently, but under what conditions AI-enabled autonomy can improve entrepreneurial action while preserving human responsibility and legitimate oversight. This framing is consistent with responsible-AI approaches that emphasize validity, reliability, safety, security, accountability, transparency, explainability, privacy, and fairness (NIST, 2023).

The paper makes three contributions. First, it clarifies the conceptual distinction between conventional AI assistance and agentic AI in entrepreneurial settings. Second, it integrates autonomous decision-making, ethical governance, and human–AI collaboration into a single framework rather than treating technical capability and ethical responsibility as separate topics. Third, it develops three research propositions that can be tested in future empirical studies involving founders, startup teams, AI-enabled ventures, or controlled decision-making tasks.

2. CONCEPTUAL FOUNDATIONS

2.1 Artificial Intelligence, Agentic AI, and Entrepreneurship

AI is a broad field encompassing computational systems that perform functions associated with perception, prediction, learning, reasoning, language processing, and decision support. Machine learning is one important approach within AI, while deep learning is a family of machine-learning methods based on multilayer neural networks. Adversarial learning refers to methods in which models are trained or evaluated against deliberately challenging inputs, including adversarial examples; it should not be treated as a synonym for unsupervised learning (Jordan & Mitchell, 2015).

Agentic AI is treated here as an AI-enabled system that can pursue specified goals through a sequence of actions, such as planning, information retrieval, tool invocation, monitoring, and adaptation, with a degree of operational autonomy. The defining issue is therefore not a particular model architecture but the system's capacity to initiate and coordinate actions within an environment. Recent entrepreneurship research has begun to examine AI agents as potential entrepreneurial collaborators and co-agents, making the topic especially timely (Obschonka & Fisch, 2025).

In entrepreneurial contexts, agentic AI may support market scanning, customer interaction, workflow orchestration, scheduling, lead qualification, experimentation, resource allocation, and bounded operational decisions. Yet the delegation of action also changes the locus of control. The entrepreneur must specify goals, constraints, permissions, escalation conditions, and acceptable risk. Consequently, agentic AI should be conceptualized as a capability embedded in a governance arrangement rather than as an independent entrepreneurial actor.

2.2 AI as Augmentation Rather Than Simple Substitution

Chalmers et al. (2021) show that AI can alter new-venture processes by augmenting or replacing tasks and thereby changing how entrepreneurial work is organized. Raisch and Krakowski (2021) further show that automation and augmentation exist in a paradoxical relationship. For early-stage ventures, the useful unit of analysis is therefore the human–AI task configuration: which activities are delegated, which remain human-controlled, how decisions are reviewed, and how responsibility is assigned.

Agentic systems may increase the speed and scale of information processing, but speed is not equivalent to decision quality. Entrepreneurial decisions frequently involve ambiguity, stakeholder interests, ethical trade-offs, and consequences that are difficult to encode as optimization objectives. Human founders therefore retain an important role in defining what should be optimized, interpreting exceptions, evaluating legitimacy, and accepting responsibility for consequential decisions.

2.3 Entrepreneurial Cognition and Opportunity Processes

Entrepreneurial opportunity processes involve noticing signals, interpreting incomplete information, evaluating alternatives, and acting under uncertainty. AI can expand the volume and speed of information available to entrepreneurs, potentially improving opportunity recognition and experimentation. However, entrepreneurial judgment is not reducible to pattern recognition. It also involves contextual knowledge, values, experience, imagination, and interaction with stakeholders. (Shepherd, 2015).

Accordingly, the conceptual model treats AI as extending entrepreneurial cognitive capacity rather than replacing entrepreneurial cognition. This distinction is important because an AI system can generate recommendations or actions without possessing the normative authority to determine which outcomes ought to be pursued.

2.4 Responsible AI and Governance

Responsible AI requires more than a general statement that an organization will use AI ethically. Governance must translate principles into operational controls across the AI lifecycle. The NIST AI Risk Management Framework organizes risk management around Govern, Map, Measure, and Manage and identifies trustworthiness characteristics including validity and reliability, safety, security and resilience, accountability and transparency, explainability and interpretability, privacy enhancement, and fairness with harmful bias managed (NIST, 2023).

The European Commission's Ethics Guidelines for Trustworthy AI emphasize that trustworthy AI should be lawful, ethical, and technically and socially robust. Jobin et al. (2019) show that global AI ethics guidance converges around recurring principles but differs in interpretation and implementation. For early-stage ventures, governance should therefore be proportionate to risk while still establishing minimum requirements for data stewardship, human oversight, documentation, testing, monitoring, incident response, and accountability. Algorithmic ethics scholarship similarly emphasizes the need to examine transparency, responsibility, and potential harms rather than treating technical performance as sufficient (Mittelstadt et al., 2016).

3. LITERATURE SYNTHESIS

3.1 Agentic AI and Autonomous Entrepreneurial Decision-Making

Agentic AI can change entrepreneurial decision processes by moving from recommendation toward bounded execution. A conventional analytics system may identify a market trend,

whereas an agentic system could identify the trend, retrieve supporting information, generate alternative responses, initiate a permitted workflow, and report the outcome. Such capabilities may reduce coordination costs and shorten response times. However, autonomy introduces additional failure modes because errors can propagate across a sequence of actions.

The value of autonomy therefore depends on task characteristics. Highly repetitive, reversible, and measurable tasks are more suitable for delegation than decisions involving irreversible commitments, sensitive personal information, legal exposure, or significant stakeholder harm. A useful governance principle is graduated autonomy: the greater the potential impact and irreversibility of an action, the stronger the requirement for human authorization or review.

3.2 Ethical Boundaries as a Condition of AI Autonomy

Ethical boundaries define what an AI system is permitted to do, what information it may access, which decisions require escalation, and which outcomes are unacceptable regardless of potential business benefit. In a startup, these boundaries may be embedded in access controls, approval thresholds, data policies, audit trails, model evaluations, human review, and incident-response procedures.

Ethical governance is particularly important because early-stage ventures may lack specialized compliance teams. The absence of formal structures does not remove governance needs; it makes explicit minimum controls more important. Responsible governance can also support venture legitimacy by demonstrating that AI deployment is deliberate, reviewable, and aligned with applicable law and organizational values.

3.3 Human–AI Collaboration, Trust, and Role Negotiation

Human–AI collaboration should not be equated with unconditional trust in machine outputs. Human–AI interaction research likewise emphasizes calibrated interaction, appropriate reliance, and clear user responsibilities (Amershi et al., 2019). Appropriate collaboration requires calibrated trust: users should rely on AI when evidence of system reliability supports reliance and should intervene when uncertainty, novelty, or risk is high. Trust is shaped by transparency, experience, system performance, and clarity of responsibility.

Role negotiation is equally important. Founders and employees need to understand whether an AI system is acting as an information provider, recommender, workflow coordinator, executor of bounded actions, or escalation mechanism. Ambiguous roles can create automation bias, diffusion of responsibility, and accountability gaps. Clear role definitions can instead create complementary human–AI arrangements in which machines contribute scale and consistency while humans provide contextual judgment, ethical interpretation, and final responsibility.

3.4 Innovation and Operational Efficiency

Consistent with the broader literature on AI business value (Enholm et al., 2022), AI may contribute to innovation by increasing the number of alternatives that entrepreneurs can explore and by lowering the cost of experimentation. It may also improve operational efficiency through workflow automation, information synthesis, forecasting, customer support, and resource coordination. However, innovation benefits are not automatic. Excessive reliance on AI-generated alternatives may narrow exploration to patterns already present in training data, while poorly governed automation may create operational errors at scale.

The expected relationship is therefore conditional: AI autonomy is most likely to create value when it is combined with task-appropriate delegation, human oversight, trustworthy data, clear objectives, and mechanisms for detecting and correcting failure.

4. RESEARCH GAP

Recent scholarship confirms that AI-and-entrepreneurship research is still developing its theoretical foundations, definitions, methods, and evidence base (Obschonka et al., 2025). Existing research provides strong foundations but remains fragmented across entrepreneurship, information systems, AI management, and responsible-AI scholarship. Research on AI and entrepreneurship explains how AI can augment or replace entrepreneurial tasks and alter venture processes (Chalmers et al., 2021). Research on AI management emphasizes the tension between automation and augmentation (Raisch & Krakowski, 2021). Responsible-AI scholarship provides principles and governance mechanisms, while AI risk-management frameworks provide operational guidance (Jobin et al., 2019; NIST, 2023).

The unresolved problem is how these perspectives can be integrated specifically for early-stage ventures using agentic systems. The literature provides comparatively less conceptual guidance on the joint relationship among degrees of AI autonomy, ethical and governance boundaries, and human–AI role structures in entrepreneurial decision processes. This paper addresses the gap by positioning governance and human collaboration as mechanisms that condition the value of AI-enabled autonomy.

5. RESEARCH OBJECTIVES

- To develop a conceptual explanation of how agentic AI can influence autonomous decision-making in early-stage ventures.
- To identify the ethical and governance mechanisms required to bound AI autonomy and support responsible venture outcomes.
- To explain how human–AI collaboration, calibrated trust, and role clarity can shape innovation and operational effectiveness.

6. CONCEPTUAL REVIEW METHOD

This paper adopts a conceptual integrative-review approach rather than an empirical research design. The source manuscript was critically reorganized and its central arguments were synthesized around agentic AI, entrepreneurial decision-making, responsible AI governance, and human–AI collaboration. Key claims and foundational references were checked against authoritative scholarly and institutional sources during revision. The paper does not claim to be a systematic review because the original manuscript did not document a reproducible database search protocol, inclusion and exclusion criteria, screening process, or formal quality appraisal.

The purpose of the review is theory development: to integrate established ideas into a coherent framework and generate propositions suitable for subsequent empirical testing. Accordingly, no respondents, survey instruments, datasets, statistical tests, p-values, effect sizes, or empirical findings are reported.

7. PROPOSED CONCEPTUAL FRAMEWORK

The framework links three dimensions: agentic AI capability, ethical governance, and human–AI collaboration. Agentic AI capability refers to planning, information processing, tool use, monitoring, and bounded execution. Ethical governance refers to accountability, transparency, privacy, fairness, documentation, human oversight, and risk management.

Human–AI collaboration refers to task allocation, calibrated trust, role clarity, review, and escalation.

Construct	Core mechanism	Expected outcome	Boundary condition
Agentic AI capability	Planning, retrieval, tool use, bounded execution	Timelier and more consistent decisions	Task suitability, reversibility, data quality
Ethical governance	Oversight, accountability, transparency, privacy, fairness	Lower harmful or illegitimate outcomes	Risk and regulatory context
Human–AI collaboration	Role clarity, calibrated trust, review, escalation	Innovation and operational effectiveness	Human expertise and organizational capability




Figure 1. Conceptual framework linking agentic AI capability, autonomous entrepreneurial decision processes, ethical governance, human–AI collaboration, and responsible entrepreneurial outcomes.

8. RESEARCH PROPOSITIONS

Proposition 1: Agentic AI and decision-making

P1: Greater use of agentic AI for suitable, bounded entrepreneurial tasks is expected to improve the timeliness and information-processing capacity of venture decision-making, particularly when tasks are measurable, reversible, and supported by reliable data.

This proposition does not imply that greater autonomy always improves decision quality. Quality depends on the appropriateness of delegation, reliability of inputs, and escalation mechanisms.

Proposition 2: Ethical governance and responsible outcomes

P2: Stronger ethical governance is expected to weaken the relationship between AI autonomy and unintended negative outcomes and to strengthen the legitimacy of AI-enabled entrepreneurial activity.

Governance includes accountability, transparency, privacy controls, bias management, documentation, monitoring, human oversight, and incident response.

Proposition 3: Human–AI collaboration and venture performance

P3: Human–AI collaboration is expected to improve innovation and operational effectiveness when trust is appropriately calibrated and human and AI roles are clearly differentiated.

AI contributes scale, speed, consistency, and computational capacity; humans contribute contextual judgment, value interpretation, ethical reasoning, creativity, and accountability.

9. THEORETICAL CONTRIBUTION

The principal theoretical contribution is a governance-contingent model of entrepreneurial agentic AI. Rather than treating AI autonomy as an independent predictor of venture performance, the framework specifies autonomy as a capability whose consequences depend

on task suitability, ethical governance, and the quality of human–AI role arrangements. This extends the automation–augmentation perspective by locating the key unit of analysis in the delegation architecture: what is delegated, under what permissions, with what review, and with whom final responsibility remains (Raisch & Krakowski, 2021). It also responds to recent calls for stronger conceptual foundations and testable theory in AI-and-entrepreneurship research (Obschonka et al., 2025).

10. DISCUSSION

The framework shifts the discussion of entrepreneurial AI from a technology-centred question—what can AI do?—to a governance-centred question—what should AI do, under whose authority, with what controls, and with what consequences? This matters because agentic systems can connect information processing directly to action. A system that recommends a course of action and a system that executes it are not equivalent from an accountability perspective.

The first theoretical implication is that entrepreneurial AI should be examined at the level of task delegation and decision architecture. The relevant construct is not simply AI adoption, but the degree and type of autonomy granted to AI within a venture. This provides a bridge between entrepreneurship research on action and management research on automation and augmentation.

The second implication is that ethical governance should be treated as a boundary condition of entrepreneurial AI rather than a separate compliance topic. When AI autonomy increases, governance mechanisms determine whether the venture can detect errors, explain decisions, assign responsibility, and intervene before harm becomes consequential.

The third implication concerns the changing role of the entrepreneur. As AI systems assume more analytical and operational tasks, founders may spend relatively more time on goal setting, exception management, stakeholder relationships, ethical judgment, and system oversight. This redistributes entrepreneurial agency across a human–AI system rather than eliminating it.

11. MANAGERIAL IMPLICATIONS

- Begin with low-risk, reversible workflows before delegating consequential decisions to agentic systems.
- Define explicit permissions, action limits, escalation thresholds, and human-approval requirements.
- Maintain auditable records of important AI-generated recommendations and actions.
- Use calibrated trust rather than assuming that high-performing AI systems are universally reliable.
- Maintain organizational accountability; do not transfer responsibility to the AI system.
- Assess data quality, privacy, bias, security, and model performance throughout the AI lifecycle.
- Train founders and employees in AI literacy, including limitations, failure modes, and appropriate intervention.

12. POLICY AND GOVERNANCE IMPLICATIONS

Policymakers and startup-support institutions should encourage proportionate AI governance rather than identical controls for every use case. Risk-sensitive approaches can combine baseline requirements—privacy protection, documentation, human accountability, and incident reporting—with stronger controls for high-impact applications. Frameworks such as the NIST AI RMF provide a lifecycle-oriented structure, while the European approach to trustworthy AI emphasizes legal, ethical, and technical robustness.

For early-stage ventures, policy support could include standardized governance templates, model-risk checklists, AI-literacy programs, responsible procurement guidance, and accessible auditing resources. Such measures can lower the cost of responsible adoption without requiring startups to reproduce the governance structures of large corporations.

13. LIMITATIONS AND FUTURE RESEARCH

This paper is conceptual and does not provide empirical estimates of the proposed relationships. Its integrative review is not systematic and should not be interpreted as exhaustive. Agentic AI is also evolving rapidly, so capabilities, regulatory expectations, and organizational practices may change.

Future research should test the propositions using mixed methods. Surveys could examine AI autonomy, governance maturity, calibrated trust, role clarity, innovation, and operational performance across startups. Case studies could examine how founders establish delegation boundaries and respond to AI failures. Experiments could compare human-only, AI-assisted, and agentic-AI decision conditions. Longitudinal research could examine how governance evolves as ventures move from experimentation to scale.

CONCLUSION

Agentic AI has the potential to reshape entrepreneurial work by extending the speed, scale, and coordination capacity of early-stage ventures. Its value, however, does not arise from autonomy alone. The central argument of this paper is that entrepreneurial value from agentic AI depends on the interaction among bounded autonomy, ethical governance, and effective human–AI collaboration.

Rather than positioning AI as an independent replacement for the entrepreneur, the framework conceptualizes agentic AI as a socio-technical capability embedded in human goals, organizational processes, and governance structures. Agentic systems may execute increasingly complex workflows, but humans remain responsible for defining legitimate objectives, setting boundaries, interpreting uncertainty, managing exceptions, and accepting accountability.

The framework and propositions provide a basis for future empirical research. The central research challenge is not simply to determine whether entrepreneurs can delegate more tasks to AI, but to identify when delegation creates value, when it creates unacceptable risk, and how ventures can design human–AI systems that remain innovative, accountable, and socially legitimate.

DECLARATIONS

Funding No external funding was received for the development of this conceptual manuscript.

Conflict of Interest The author declares no conflict of interest.

Data Availability No primary or secondary dataset was generated or analyzed in this conceptual study.

Ethics Statement Ethical approval was not required because the paper does not involve human participants, identifiable personal data, or primary empirical data.

REFERENCES

1. Amershi, S., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human–AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. <https://doi.org/10.1145/3290605.3300233>
2. Chalmers, D., MacKenzie, N. G., & Carter, S. (2021). Artificial intelligence and entrepreneurship: Implications for venture creation in the Fourth Industrial Revolution. *Entrepreneurship Theory and Practice*, 45(5), 1028–1053. <https://doi.org/10.1177/1042258720934581>
3. Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2022). Artificial intelligence and business value: A literature review. *Information Systems Frontiers*, 24, 1709–1734. <https://doi.org/10.1007/s10796-021-10186-w>
4. European Commission High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. European Commission. <https://doi.org/10.2759/346720>
5. Ferreira, R., Araújo, M., Tereso, A., & Novais, P. (2026). Agentic AI architecture for multi-criteria decision-making: A collaborative human-AI framework. *Group Decision and Negotiation*, 35, 45. <https://doi.org/10.1007/s10726-026-10004-1>
6. Haenlein, M., & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California Management Review*, 61(4), 5–14. <https://doi.org/10.1177/0008125619864925>
7. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
8. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
9. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2). <https://doi.org/10.1177/2053951716679679>
10. NIST. (2023). Artificial intelligence risk management framework (AI RMF 1.0). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
11. Obschonka, M., & Fisch, C. (2025). From humans to machines: Researching entrepreneurial AI agents built on large language models. *Journal of Business Venturing Insights*, 24, e00581. <https://doi.org/10.1016/j.jbvi.2025.e00581>
12. Obschonka, M., Grégoire, D. A., Nikolaev, B., Ooms, F., Lévesque, M., Pollack, J. M., & Behrend, T. S. (2025). Artificial intelligence and entrepreneurship: A call for research to prospect and establish the scholarly AI frontiers. *Entrepreneurship Theory and Practice*, 49(3), 620–641. <https://doi.org/10.1177/10422587241304676>
13. Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
14. Shepherd, D. A. (2015). Party on! A call for entrepreneurship research that is more interactive, activity based, cognitively hot, and compassionate. *Journal of Business Venturing*, 30(4), 489–507. <https://doi.org/10.1016/j.jbusvent.2015.02.001>