

MSc DISSERTATION

How EdTech Startups Implement AI Agents to Reduce Operational Costs and Accelerate Growth

Jeffrey Nana Boakye

MSc IT Management

Berlin School of Business and Innovation, in partnership with the University for the Creative Arts

2026

Awarded 71% (Distinction)

This edition corrects grammatical, typographical and numerical errors in the version submitted for assessment and revises the wording for clarity. The data, findings and recommendations are unchanged.

Abstract

The education technology (EdTech) sector has grown rapidly, placing increasing pressure on startups to scale with limited resources. Many of these startups now deploy AI agents: systems capable of managing administrative workflows, responding to learner queries and personalising content delivery with limited human involvement.

This study examines how EdTech startups implement AI agents to reduce operational costs and accelerate growth. A quantitative research design was adopted, using a structured online questionnaire completed by 45 startup professionals across 10 countries. Because the target population is niche, the sample was drawn from across the startup sector, with EdTech respondents forming the second-largest industry group (24.4 per cent); findings are therefore interpreted for EdTech in light of this cross-sector evidence.

The findings show that AI agent adoption was near-universal (91.1 per cent actively using or piloting), with customer support, administrative automation, content generation and software development the most common deployment areas. Respondents reported substantial reductions in operational costs (mean 3.82 out of 5) and meaningful contributions to growth (mean 3.82 out of 5), with scaling without proportional increases in headcount the most frequently reported outcome. Barriers were also evident, chiefly data privacy concerns, high implementation costs, data quality issues and skills shortages. Read alongside the literature, these barriers appear to reinforce one another rather than act in isolation.

The study concludes with practical recommendations for EdTech startup managers considering AI agent adoption as part of their growth strategy.

Keywords: EdTech, AI agents, operational costs, startup growth, artificial intelligence, business automation, education technology.

Acknowledgements

I would like to express my sincere gratitude to my supervisor for their guidance and support throughout this research. I also thank the Berlin School of Business and Innovation for providing the resources needed to complete this dissertation.

Special thanks go to my family and friends for their encouragement and support throughout this journey.

Statement of Compliance with Academic Ethics and the Avoidance of Plagiarism

I honestly declare that this dissertation is entirely my own work and none of its part has been copied from printed or electronic sources, translated from foreign sources and reproduced from essays of other researchers or students. Wherever I have been based on ideas or other people texts I clearly declare it through the good use of references following academic ethics.

(In the case that is proved that part of the essay does not constitute an original work, but a copy of an already published essay or from another source, the student will be expelled permanently from the postgraduate program).

Name and Surname (Capital letters): **JEFFREY NANA BOAKYE**

Date: **10 June 2026**

Contents

Abstract.....	2
Acknowledgements.....	3
Statement of Compliance.....	4
Introduction.....	7
Trigger and Rationale.....	7
Aims and Objectives.....	8
Research Questions.....	8
Methodology.....	9
Synopsis of Chapters.....	9
Chapter One – Literature Review I.....	10
1.1 The Growth and Transformation of the EdTech Sector.....	10
1.2 The Rise of Artificial Intelligence in Business Operations.....	11
1.3 Artificial Intelligence and Entrepreneurship.....	12
1.4 Theoretical Frameworks for Understanding AI Adoption.....	14
1.5 AI-Driven Business Models and Value Creation.....	16
1.6 Chapter Summary.....	17
Chapter Two – Literature Review II.....	18
2.1 AI Agent Deployment in EdTech Operational Functions.....	18
2.2 AI and Operational Cost Reduction in Startups.....	19
2.3 AI-Driven Content Personalisation and Adaptive Learning.....	21
2.4 Barriers to AI Implementation in EdTech Startups.....	21
2.5 Ethical Considerations in AI-Enabled Education.....	23
2.6 Gaps in the Existing Literature.....	24
2.7 Chapter Summary.....	25
Chapter Three – Methodology.....	26
3.1 Research Philosophy and Design.....	26
3.2 Justification for the Chosen Methodology.....	26
3.3 Rejection of Alternative Methodologies.....	27
3.4 Data Collection Method.....	27
3.5 Sample Selection.....	28
3.6 Data Analysis Approach.....	29
3.7 Ethical Considerations.....	30
3.8 Limitations of the Methodology.....	30
Chapter Four – Findings, Analysis and Discussion.....	32
4.1 Findings.....	32
4.2 Analysis.....	41
4.3 Discussion.....	43

Concluding Remarks.....	50
5.1 Summary of Key Findings.....	50
5.2 Answering the Central Research Question.....	51
5.3 Recommendations.....	51
5.4 Limitations and Future Research.....	52
Bibliography.....	53
Appendices.....	55
Appendix A: Survey Questionnaire.....	55
Appendix B: Survey Data Summary.....	59

Introduction

The education technology sector, commonly referred to as EdTech, has undergone a major transformation over recent years, shifting from a supplementary tool to a core pillar of modern instruction. Advances in artificial intelligence, machine learning and digital infrastructure have changed the way educational services are delivered, accessed and managed. HolonIQ (2019) projected that the global EdTech market would reach USD 404 billion by 2025, a figure that reflects the scale of demand for digital learning solutions worldwide. At the same time, the World Economic Forum (2023) estimated that 83 million jobs would be eliminated globally by 2027, against 69 million created, a churn that has intensified the urgency of reskilling and upskilling. EdTech companies are well positioned to meet this need.

Within this growing market, EdTech startups occupy a distinctive position. They move quickly, innovate readily and are often well attuned to the needs of learners. However, they also face significant operational challenges. Most operate with small teams and tight budgets, yet they are expected to grow quickly, serve large numbers of users and maintain high-quality educational experiences. To reconcile these pressures, many have begun to adopt AI agents: tools that can complete tasks with limited step-by-step human supervision. Such agents can respond to learner queries, write code, process administrative requests, personalise marketing and educational content, and analyse user behaviour in real time. For a startup seeking to scale without hiring large numbers of additional staff, this kind of automation is not merely convenient; it can be critical. Chopra (2025) notes that agentic AI systems have been identified as one of the most important emerging technology trends for organisations seeking to improve efficiency and reduce overhead costs.

Despite widespread interest in AI, there is still little academic research on how EdTech startups specifically are using it. Much of the existing research focuses either on AI in education from a teaching and learning perspective (OECD, 2020; UNESCO, 2023) or on AI adoption within large enterprises (Farayola et al., 2023; Chinnaraju, 2025). Comparatively little examines how small, early-stage EdTech companies use AI agents to run their businesses. This study was undertaken to help address that gap.

Trigger and Rationale

This research topic was chosen for two main reasons. First, agentic AI, meaning software that can complete a sequence of tasks without continuous human direction, has rapidly gained prominence and was named a top strategic technology trend for 2025 (Gartner, 2024, cited in Chopra, 2025). Yet practitioner-level understanding of how such systems translate into measurable business outcomes in EdTech startups remains sparse. Second, professional

engagement with the EdTech sector revealed a persistent tension between the growth expectations placed on startups by investors and the operational capacity constraints those organisations face. AI agents appeared to offer a viable means of resolving this tension, yet no comprehensive academic investigation had examined the relationship in a focused, empirical manner.

This study was therefore undertaken to generate practical, evidence-based insight into how EdTech startups are implementing AI agents, what outcomes they are experiencing and what challenges they encounter along the way.

Aims and Objectives

The aim of this dissertation is to investigate how EdTech startups use AI agents as a tool for reducing operational costs and accelerating business growth. To achieve this aim, the following objectives were set:

- to examine the types of AI agent technologies being adopted by EdTech startups;
- to identify the operational areas in which AI agents generate the most significant cost savings;
- to evaluate the relationship between AI agent adoption and growth within EdTech startups;
- to explore the barriers that EdTech startups face when implementing AI agents; and
- to provide practical recommendations for EdTech startup managers who are considering, or are in the process of, adopting AI agents.

Research Questions

The central research question guiding this study is: *How do EdTech startups implement AI agents to reduce operational costs and accelerate growth, and what are the principal challenges and outcomes associated with this process?*

Three supporting questions were also addressed:

- What types of AI agents are most commonly deployed by EdTech startups, and for what purposes?
- To what extent does AI agent implementation contribute to measurable reductions in operational costs?
- What organisational and technical barriers do EdTech startups encounter during implementation, and what factors do startups identify as most helpful in overcoming them?

Methodology

This study adopted a quantitative research design, using a structured online questionnaire to collect primary data from professionals working within startups. A quantitative approach was selected because the topic, while still emerging, required the identification of patterns and trends across a broader sample than qualitative methods alone could provide (Saunders, Lewis and Thornhill, 2019). The questionnaire was distributed to startup professionals via email, LinkedIn and an online startup community. A total of 45 completed responses were received from professionals across 10 countries, representing sectors including education technology, software as a service, financial technology and health technology. The primary data were complemented by an extensive review of secondary literature drawn from peer-reviewed academic sources, industry reports and scholarly frameworks relevant to AI adoption, EdTech business models and operational efficiency. The limitations of the methodology, including the sample size and the constraints of non-probability sampling, are discussed in detail in Chapter Three.

Synopsis of Chapters

Chapter One provides a broad review of the literature on the general elements of the research topic, including the growth of the EdTech sector, the rise of AI in business and the theoretical frameworks used to understand AI adoption in startup organisations.

Chapter Two focuses more specifically on the concepts central to the research question. It examines how AI agents have been deployed within EdTech operational functions, what the literature says about the relationship between AI and cost efficiency, and what barriers organisations typically encounter during implementation.

Chapter Three sets out the research design in detail. It explains why a quantitative methodology was chosen, describes the structured questionnaire and the sampling strategy, and discusses the limitations of the approach.

Chapter Four presents the primary data gathered through the survey, analyses the findings using descriptive statistics and synthesises the primary and secondary research into a discussion of the study's broader implications for EdTech startup management and AI strategy.

The **Concluding Remarks** summarise the key findings, provide a direct answer to the central research question and offer recommendations for EdTech startup managers and practitioners.

Chapter One – Literature Review I

This chapter reviews the academic and industry literature relevant to the general dimensions of the research topic. Its purpose is to establish the theoretical and contextual foundations on which the more specific analysis of AI agent implementation in EdTech startups, presented in Chapter Two, is built. The discussion covers the growth and transformation of the education technology sector, the emergence of artificial intelligence as a force in business operations, the evolving relationship between AI and entrepreneurship, and the principal theoretical frameworks used to understand technology adoption. A critical approach is adopted throughout: rather than merely summarising existing studies, the review interrogates the assumptions, limitations and contextual boundaries of the literature to identify both its contributions and its shortcomings.

1.1 The Growth and Transformation of the EdTech Sector

The education technology sector has undergone an unprecedented period of expansion over the last decade, moving from a collection of supplementary digital tools to a central pillar of modern educational infrastructure. HolonIQ (2019) projected that the global EdTech market would reach USD 404 billion by 2025, a figure that reflects both the scale of demand for digital learning solutions and the accelerating pace of investment. This projection sits within a wider context of workforce disruption. The World Economic Forum (2023) estimated that 83 million jobs would be eliminated globally by 2027, which has intensified the urgency of the reskilling and upskilling programmes that EdTech companies are well placed to deliver.

The COVID-19 pandemic accelerated this trajectory by compelling educational institutions worldwide to adopt digital platforms at scale. UNESCO (2023) observed that this shift increased the uptake of AI-enabled learning technologies while highlighting the potential of artificial intelligence to personalise educational delivery. Andersson and Rosenqvist (n.d.) examined how AI integration is transforming EdTech business models, noting that companies such as Duolingo, Coursera and Byju's have used AI to build scalable learning solutions. Their analysis showed that incorporating AI requires new approaches to value creation and capture, fundamentally altering the economic models on which the sector was built.

This trend towards AI-embedded business models has not been confined to large, well-funded platforms. Alam and Mohanty (2022) observed that EdTech companies are integrating AI and learning analytics into their business models and strategies in order to create personalised educational experiences, reflecting a wider shift towards data-driven, customised approaches across the sector.

These growth narratives nevertheless deserve critical scrutiny. The HolonIQ projection, although widely cited, was produced by an industry-aligned market intelligence firm and may rest on optimistic assumptions about adoption rates. Similarly, the Duolingo and Coursera case studies, while useful examples of good practice, describe well-funded, mature platforms with access to world-class engineering talent and substantial venture capital. Whether such models can realistically be replicated by early-stage startups with limited resources remains an open question that the literature has not adequately addressed. Byju's is instructive in this respect: it subsequently experienced significant financial difficulties and workforce reductions, suggesting that AI-enhanced growth does not automatically translate into sustainable business performance (Andersson and Rosenqvist, n.d.). The literature's tendency to foreground success stories while underreporting failures introduces a form of survivorship bias that risks overstating the benefits of AI adoption for startups.

The EdTech landscape also remains highly fragmented, comprising a large number of startups competing alongside established technology firms. Most of these startups operate with tight budgets and small teams, yet face the same expectations of rapid scaling (Andersson and Rosenqvist, n.d.). This tension between growth expectations and operational capacity forms the practical impetus for the present study. Moreover, the literature generally treats EdTech startups as a homogeneous category, overlooking significant variations in maturity, geographic context and regulatory environment, all of which shape both the feasibility and the outcomes of AI adoption.

1.2 The Rise of Artificial Intelligence in Business Operations

Artificial intelligence, broadly defined as the simulation of human cognitive processes by computer systems, has become one of the most transformative technologies shaping today's business environment. Brynjolfsson and McAfee (2017) argue that AI is fundamentally a general-purpose technology, capable of restructuring entire industries by automating tasks, generating data-driven insights and enabling new forms of decision-making. In their view, AI is not merely a source of efficiency gains but, more fundamentally, a driver of competitive advantage.

Dwivedi et al. (2021) offer a broader picture in what is arguably one of the most comprehensive multidisciplinary assessments of AI to date, drawing on the perspectives of 43 experts from computer science, marketing, information systems and management. The picture that emerges is mixed. AI opens up genuine opportunities for innovation, but it also raises substantial challenges, with ethics, transparency and workforce displacement recurring throughout. Of particular note, Dwivedi et al. observed that AI adoption is uneven: its extent depends on sector, organisational size and geographic location, and smaller firms tend to face disproportionate barriers compared with larger competitors. This finding is significant for the present study because it directly

challenges an assumption that runs through much of the literature, namely that the benefits of AI are equally available to organisations of all sizes.

On the technological side, Ojika et al. (2024) identify three principal AI technologies reshaping business process automation. Machine learning enables systems to improve their performance as they are exposed to more data. Natural language processing allows machines to interpret and generate human language. Robotic process automation uses software robots to handle routine, rule-based work. Together, these technologies have extended automation beyond repetitive tasks into work that requires judgement and a degree of contextual decision-making.

The more recent concept of agentic AI represents a qualitative shift in capability. Gartner (2024) identified it as one of the top strategic technology trends for 2025, predicting that by 2028 at least 15 per cent of day-to-day work decisions will be made autonomously by AI agents. The significance of this lies in how agentic systems differ from traditional AI. Traditional AI systems are reactive, responding to predefined prompts, whereas agentic AI systems are characterised by their ability to plan, execute and adapt workflows independently (Chopra, 2025). This distinction matters for the present study because it positions AI agents as active participants in organisational decision-making rather than passive tools awaiting instruction. It should be acknowledged, however, that the Gartner prediction is a forward-looking estimate produced by a commercial technology advisory firm; it reflects projected rather than empirically validated trends. This gap between industry projection and empirical evidence is a limitation that recurs across this body of literature.

1.3 Artificial Intelligence and Entrepreneurship

The intersection of artificial intelligence and entrepreneurship has attracted growing academic attention, driven largely by recognition that AI technologies are altering how new ventures are created, scaled and managed. Giuggioli and Pellegrini (2023) conducted a systematic literature review of 60 studies and proposed an "AI-enabled entrepreneurial process" model, which maps the impact of AI across four dimensions: opportunity recognition, decision-making, performance improvement and education. Their framework portrays AI not as a factor confined to a single stage of the entrepreneurial journey but as one that runs through all of it: opportunity identification is enhanced by pattern recognition, and strategic decision-making is sharpened by data analytics. The framework is valuable chiefly because of its breadth.

The framework does, however, have a notable limitation: it relies predominantly on literature drawn from well-resourced startups in developed economies. Whether its conclusions hold for early-stage ventures in emerging markets, where digital infrastructure, regulatory frameworks and access to AI talent differ markedly, has not been empirically tested. This contextual limitation is relevant to the present study, whose participants are drawn from both European and African

startup ecosystems, where the conditions enabling AI-enhanced entrepreneurship differ considerably.

Farayola et al. (2023) extended this analysis by examining how AI technologies are driving the emergence of innovative business models. Their review showed that AI-driven models are typically characterised by enhanced operational efficiency and the capacity to create value through data-driven personalisation. For startups in particular, the ability to use AI as a substitute for human labour in routine functions represents a genuine competitive advantage, allowing lean teams to achieve output levels that would traditionally have required a much larger workforce. The authors are careful, however, not to present an overly optimistic picture. They note that AI adoption introduces new dependencies and vulnerabilities, including over-reliance on third-party AI platforms, exposure to algorithmic bias that is difficult to audit and the risk of over-automating functions in which human judgement should be retained.

Building on this dual perspective, Lu, Zhao and Zheng (forthcoming) propose the Empowerment–Entrapment Framework, which holds that generative AI cuts both ways for entrepreneurs. Their model identifies that while AI can improve idea quality, boost productivity and broaden functional capacity, it also introduces risks: hallucinations in AI-generated outputs, entrepreneurial overconfidence arising from the perception of AI as infallible, reduced relational embeddedness as automated processes replace human interaction, and the gradual erosion of critical thinking skills. This balanced view is essential to the present study, which examines not only the benefits of AI agent implementation but also its costs and unintended consequences.

The challenges of AI adoption in smaller firms have been further documented by Mohd Rasdi and Umar Baki (2025). Their qualitative study of six small and medium-sized enterprises (SMEs) identified several significant internal barriers: reluctance to change, fear of job displacement among staff, restricted financial resources and a shortage of in-house technical expertise. Externally, a rapidly changing technological landscape and tightening regulatory requirements demand a degree of agility that many SMEs are not in a position to deliver. These findings are directly relevant to the EdTech startup context, although the study's single-country sample limits how far its conclusions can be extended to other regulatory and economic settings.

The OECD (2020) brings an institutional perspective to the discussion. Its report observes that AI-powered chatbots and virtual assistants are increasingly used to provide continuous learner support, while AI-driven translation tools are reducing the linguistic barriers that once made cross-border learning more difficult. The OECD also acknowledges that these developments have real implications for how EdTech companies structure their operations and business models. Nevertheless, like much of the policy literature, the OECD analysis tends to describe potential applications rather than evaluate their effectiveness or examine the conditions under which they succeed or fail.

A further dimension that the entrepreneurship literature has not explored in sufficient depth is the role of regional context in shaping AI adoption outcomes. Empirical studies of AI and entrepreneurship draw predominantly on data from North America, Western Europe and parts of East Asia, settings in which digital infrastructure is established, venture capital is accessible and the technical talent pool is deep. Obiye (2024) departs from this pattern by examining AI adoption among Nigerian startups. The study found a more mixed reality: AI-powered systems contributed to revenue growth and operational efficiency in the firms studied, but infrastructural and financial constraints meant that adoption clustered within a narrow band of industries rather than spreading widely. The implication is clear: the relationship between AI adoption and startup performance is not linear but is mediated by contextual factors that the dominant literature, with its focus on well-resourced ecosystems, tends to underplay. For the present study, which includes participants from both European and African startup ecosystems, this contextual sensitivity is methodologically important.

1.4 Theoretical Frameworks for Understanding AI Adoption

The academic literature on technology adoption offers several well-established theoretical frameworks that help explain how and why EdTech startups adopt AI agents. Four are particularly applicable: the Technology Acceptance Model, the Resource-Based View, Diffusion of Innovation Theory and the Dynamic Capabilities Framework. Each examines adoption through a different lens, and each has limitations that should not be overlooked.

Technology Acceptance Model (TAM)

Originally proposed by Davis (1989), the Technology Acceptance Model holds that perceived usefulness and perceived ease of use are the two primary determinants of technology adoption. In the startup context, perceived usefulness refers to the extent to which a manager believes that AI will enhance operational performance, while perceived ease of use relates to the technical complexity of integration. Chinnaraju (2025) applied TAM specifically to AI adoption in startups and found it useful in explaining why some startups adopted AI rapidly while others remained hesitant.

TAM nevertheless has limitations in this context. It was designed to explain individual-level technology acceptance within organisations, not firm-level strategic adoption decisions. Its applicability to startups, where decisions are often made by a single founder under extreme resource constraints, has been questioned (Dwivedi et al., 2021). Furthermore, TAM does not account for external factors such as competitive pressure, regulatory requirements or access to funding, which frequently determine outcomes in the startup environment.

Resource-Based View (RBV)

The Resource-Based View, popularised by Barney (1991), holds that competitive advantage derives not from simply possessing resources but from strategically deploying resources that are valuable, rare, inimitable and non-substitutable. Applied to AI, the framework suggests that startups able to leverage AI effectively gain a resource advantage that competitors will struggle to replicate quickly. Chinnaraju (2025) argued that AI constitutes a strategic resource enabling differentiation in crowded markets. For EdTech startups, deploying AI agents as a force multiplier is a compelling application of RBV principles. The difficulty with this perspective is that it assumes AI capabilities are rare and inimitable. As AI tools become increasingly commoditised through platforms such as OpenAI, Anthropic and Google Cloud, whether AI genuinely constitutes an inimitable resource, rather than a widely available commodity, warrants critical examination.

Diffusion of Innovation Theory

Rogers (2003) proposed Diffusion of Innovation Theory to explain how new technologies spread through social systems over time. The model identifies five adopter categories and five characteristics that influence the rate of adoption: relative advantage, compatibility, complexity, trialability and observability. Applied to EdTech, the framework helps explain why AI adoption has been so uneven across the sector. Technically proficient founders are more likely to fall within the innovator category, while those without technical expertise tend to adopt later, if at all (Chinnaraju, 2025). The theory's limitation in this context is its implicit assumption that the technology being diffused is relatively stable and moves through a defined population. AI agent technology is evolving so rapidly that the "innovation" being diffused is a moving target rather than a fixed product, which makes Rogers' sequential adoption model harder to apply cleanly.

Dynamic Capabilities Framework

Developed by Teece, Pisano and Shuen (1997), the Dynamic Capabilities Framework emphasises a firm's ability to integrate, build and reconfigure competencies in response to changing environments. AI agents, through their capacity to learn from data and adapt their behaviour autonomously, align closely with this concept. Startups that integrate AI agents into their operations are doing more than adding a tool: they are building adaptive capacity that positions them to respond to market shifts faster than competitors reliant on labour-intensive processes (Giuggioli and Pellegrini, 2023; Chinnaraju, 2025). The framework's principal limitation is that it was developed to explain the behaviour of established firms with existing resource bases to reconfigure. Whether early-stage startups, which are still building rather than reconfiguring capabilities, can be meaningfully analysed through this lens remains a matter of theoretical debate.

Synthesis

Taken together, these four frameworks offer complementary rather than competing perspectives on AI adoption. TAM addresses individual-level adoption perceptions; RBV explains why AI can function as a source of competitive advantage; Diffusion of Innovation accounts for the uneven pace of adoption across firms and contexts; and Dynamic Capabilities captures the organisational agility that AI can enable. No single framework is sufficient on its own, which is why the present study draws on all four while acknowledging their respective limitations. This multi-theoretical approach is consistent with Chinnaraju (2025), who argued that a comprehensive understanding of AI adoption in startups requires the integration of multiple perspectives.

1.5 AI-Driven Business Models and Value Creation

The integration of AI into business operations has given rise to new models of value creation that depart considerably from traditional approaches. Farayola et al. (2023) identified several categories of AI-driven business model, including platform-based models that use AI for personalisation and service models in which AI agents perform tasks previously undertaken by human employees. Their review found that the most successful AI-driven models are those that align AI capabilities with clearly defined customer needs, rather than those that adopt AI for its own sake.

Ojika et al. (2024) proposed a model explaining how AI reduces operational costs through business process automation. They identified three mechanisms: the replacement of manual tasks with automated workflows, the improvement of decision-making accuracy through predictive analytics, and the integration of continuous learning capabilities that allow processes to be optimised on an ongoing basis. These findings are directly relevant to the present study. However, the model was developed from a general business perspective and does not distinguish between large enterprises with dedicated IT departments and early-stage startups in which technical implementation is often managed by the founders themselves. How well the model scales down to resource-constrained environments requires further empirical validation.

Chinnaraju (2025) introduced a scalability perspective through the Lean Scalability Model, demonstrating that startups using AI-powered automation can scale their operations without a proportional increase in costs. Customer support chatbots, automated scheduling and AI-driven analytics were identified as areas of significant efficiency gain. Andersson and Rosenqvist (n.d.) complemented this analysis with EdTech-specific case studies showing that AI-driven personalised learning paths have enhanced user engagement and opened up opportunities for scalable educational delivery. The combined implication is that AI adoption is no longer merely a technological choice but is becoming a strategic imperative for startup survival. This conclusion should, however, be tempered by the recognition that the evidence base remains narrow: it is

dominated by a small number of high-profile case studies and lacks the rigorous causal analysis that would strengthen its claims.

A further observation concerns the methodological orientation of the existing literature. Most of the studies reviewed in this chapter rely on secondary data analysis, conceptual modelling or systematic literature review. Primary empirical research that captures the actual adoption practices and decision-making of startup practitioners is notably scarce. This gap is consequential, because the adoption of AI agents in startups is shaped by contextual, organisational and individual factors that are difficult to capture through secondary analysis alone. The present study addresses this gap by collecting primary quantitative data through a structured online questionnaire, as described in Chapter Three.

1.6 Chapter Summary

This chapter has reviewed the general literature relevant to the research topic. The EdTech sector was established as a rapidly growing industry with significant demand for AI-enabled solutions. AI was positioned as a transformative technology with implications spanning business operations and entrepreneurship more broadly. Four theoretical frameworks were identified as lenses for understanding AI adoption, each contributing a distinct perspective while carrying its own contextual limitations. The literature on AI-driven business models demonstrated real potential for cost reduction and scalable growth, although, on closer inspection, the evidence base is dominated by case studies of well-resourced platforms and lacks rigorous causal analysis.

The review also revealed a notable gap. While the general relationship between AI and business performance is reasonably well documented, empirical research at the specific intersection of AI agents and EdTech startups remains scarce. Existing studies tend to focus either on large enterprises or on AI in education from a pedagogical rather than an operational perspective, and they tend to treat startups as a homogeneous category, overlooking variation in maturity, geography and regulatory context. This gap reinforces the rationale for the present study and provides the foundation for the more specific literature reviewed in the following chapter.

Chapter Two – Literature Review II

The preceding chapter established the general context of the research topic, covering the growth of the EdTech sector, the role of AI in business and the theoretical frameworks applicable to technology adoption. This chapter narrows the focus to the literature that speaks directly to the central research question: how EdTech startups implement AI agents to reduce operational costs and accelerate growth. It addresses four specific dimensions: the operational functions in which AI agents have been deployed within EdTech environments; the mechanisms through which AI contributes to cost reduction; the barriers and challenges that arise during implementation; and the ethical considerations raised by the deployment of AI in educational contexts. As in the preceding chapter, a critical approach is adopted throughout. The review interrogates the strength of the available evidence, identifies methodological limitations in existing studies and highlights areas in which claims outpace empirical support.

2.1 AI Agent Deployment in EdTech Operational Functions

The literature identifies several operational domains in which AI agents have been deployed within EdTech environments. The most commonly cited are customer support, content generation, administrative automation and learner analytics.

Chinnaraju (2025) observed that chatbots equipped with natural language processing capabilities have become one of the most prevalent forms of AI deployment in startups. These systems respond to customer and learner queries in real time without direct human involvement. They handle a high volume of routine enquiries, escalate more complex cases to human agents and learn from each interaction, gradually improving their accuracy over time.

The effectiveness of AI-powered customer support in the EdTech context nonetheless deserves closer scrutiny than the literature typically offers. Most studies report deployment without rigorous evaluation of user satisfaction, resolution rates or the quality of learner interactions. A chatbot that deflects queries away from a human support team reduces costs, but whether it delivers educationally meaningful responses to learners facing genuine comprehension difficulties is a separate question, and one the literature has largely left unanswered. The conflation of cost efficiency with service quality is a recurring analytical weakness in this body of work.

Content generation and personalisation constitute another critical operational function. Through their case studies of Duolingo and Coursera, Andersson and Rosenqvist (n.d.) showed that AI-driven content personalisation is no longer a peripheral feature of successful EdTech models but sits close to their core. The underlying algorithms draw on learner behaviour, performance data and engagement patterns to construct learning paths tailored to the individual rather than the

cohort. The OECD (2020), approaching the topic from a different angle, highlights AI-enabled translation and localisation tools as a further contribution, enabling EdTech companies to expand into linguistic markets that until recently would have posed a significant barrier to scaling.

Administrative automation represents a further area of deployment. According to Ojika et al. (2024), AI-driven systems are increasingly used in the administrative layer of EdTech operations, with scheduling, enrolment management, progress tracking and reporting all automated to varying degrees. In practical terms, this frees time previously spent on routine administrative work and allows teams to devote greater attention to strategic activities.

Learner analytics is the fourth major domain. UNESCO (2023) draws attention to the use of AI-powered analytics systems to monitor learner progress, identify students at risk and generate actionable insights. These systems produce not only static reports but also real-time dashboards and predictive models, which inform both educational interventions and business decisions. For EdTech startups, the value of analytics is therefore twofold: it improves educational outcomes and provides data-driven insights for product development and marketing strategy. The limitation lies in the data itself. The accuracy of predictions depends heavily on the volume and quality of the training data, and early-stage startups with small user bases often lack such data. This creates a paradox: the startups with the most to gain from AI-driven analytics are frequently those least able to generate the data needed to train the underlying models effectively.

2.2 AI and Operational Cost Reduction in Startups

The link between AI adoption and operational cost reduction is one of the most prominent themes in recent literature. Chopra (2025) reports that organisations deploying AI-driven automation are realising substantial savings, which arise through several channels. AI reduces costs by automating repetitive work across customer service, data entry, compliance monitoring and financial reporting, freeing human capital for the strategic and innovation-focused activities the business requires.

Ojika et al. (2024) proposed a model identifying three mechanisms through which AI contributes to cost reduction: the replacement of manual tasks with automated workflows; improved decision-making accuracy through predictive analytics, which reduces costly errors; and the integration of continuous learning capabilities, which enable ongoing rather than one-off optimisation. Organisations adopting AI-powered business process automation have achieved measurable reductions in overheads, most visibly in customer support, finance and human resources. As noted in Chapter One, however, the model was developed from a general business perspective and does not distinguish between large enterprises with dedicated IT infrastructure and early-stage startups in which implementation is often managed by founders with limited

technical support. Whether the same cost-reduction mechanisms operate equally in both contexts remains an unresolved empirical question.

Chinnaraju (2025) brought the analysis more directly into the startup context, arguing through the Lean Scalability Model that AI agents play a crucial role in enabling scalable growth. The functions with the most pronounced efficiency gains, namely customer support chatbots, automated scheduling and AI-driven analytics, are consistent with the lean operations literature, which has long emphasised resource efficiency in early-stage ventures (Ries, 2011). Obiye (2024) adds evidence from a different geography, reporting that AI-powered customer support and data analytics contributed to genuine reductions in operational expenditure among the Nigerian startups studied.

Farayola et al. (2023) contributed a structural perspective, showing that AI allows firms to reduce the marginal cost of serving additional customers to close to zero, particularly in functions such as automated support and algorithmic content recommendation. For EdTech startups pursuing rapid user acquisition, this characteristic is especially advantageous, as it effectively decouples the cost of service delivery from user growth.

The cost-reduction literature has, however, frequently overlooked an important distinction concerning the maturity of the startup in question. A pre-revenue startup with ten users and a growth-stage startup with ten thousand users face neither the same cost-reduction opportunities nor the same challenges. For the former, the overhead of implementing AI agents may exceed the costs being replaced, making automation economically irrational in the short term. For the latter, the economies of scale enabled by AI-driven automation are substantial and well documented. Failing to distinguish between these stages weakens the practical value of many cost-reduction claims and creates a concrete risk that early-stage founders will be encouraged to invest in AI infrastructure prematurely, before their operational scale can justify it.

Examined critically, these claims are subject to further caveats. Most studies report cost savings in percentage terms without disclosing baseline figures, which makes it difficult to assess absolute impact. The literature also consistently underreports the upfront costs of AI implementation, including infrastructure development, talent acquisition, integration testing and ongoing maintenance. A startup might reduce its customer support costs by 40 per cent through chatbot deployment, but if implementation consumes a disproportionate share of its available capital, the net financial effect may be negative in the short term. The temporal dimension of cost-benefit analysis has not been adequately addressed: the literature tends to present cost reduction as an immediate outcome rather than a deferred one.

2.3 AI-Driven Content Personalisation and Adaptive Learning

Content personalisation through AI is one of the most extensively researched applications in the education sector. UNESCO (2023) noted that AI technologies create personalised learning experiences by analysing individual learner profiles, performance histories and engagement patterns, and by tailoring content delivery in real time. For learners, personalised content tends to increase engagement and improve outcomes. For EdTech providers, it strengthens the value proposition and contributes to higher retention rates, both of which are critical revenue drivers in subscription-based models.

Andersson and Rosenqvist (n.d.) illustrated this through Duolingo's adaptive algorithms, which continuously assess performance and adjust difficulty, pacing and content as the learner progresses. Chinnaraju (2025) noted that AI-powered recommendation systems enable EdTech platforms to maximise cross-engagement by suggesting supplementary courses tailored to individual profiles. The operational implication for startups is significant: by automating content curation and adaptation, AI agents reduce the need for large content development teams, allowing personalised experiences to be delivered at a fraction of the former manual cost (Chopra, 2025).

The evidence for the efficacy of AI-driven personalisation is not, however, as unequivocal as some of the literature implies. Personalisation algorithms are only as effective as the data on which they are trained, and for early-stage EdTech startups with limited user data, the cold-start problem presents a significant obstacle. Without sufficient interaction data, algorithms cannot generate meaningful personalisation and instead produce generic recommendations, undermining the very advantage the system was introduced to provide. Furthermore, there is little evidence on whether AI-generated learning paths deliver pedagogical outcomes equivalent to those designed by qualified educators. The assumption that algorithmic efficiency and educational effectiveness are aligned is widespread, but it has not yet been adequately tested.

2.4 Barriers to AI Implementation in EdTech Startups

The literature on AI adoption makes clear that its benefits, while real, are not achieved without difficulty. A range of technical, financial, organisational and regulatory barriers has slowed the pace of AI adoption among smaller ventures and limited how far it spreads in practice.

Technical and infrastructure challenges

Ojika et al. (2024) emphasise that integrating AI into existing business operations is not a plug-and-play exercise; it requires specialised expertise and careful planning. Many startups lack high-performance computing resources, robust data management frameworks or scalable cloud architectures. Mohd Rasdi and Umar Baki (2025) further note that the rapid pace of technological

change creates additional complexity, as organisations must maintain the stability of existing systems while introducing capabilities that are still maturing. For EdTech startups this presents a particular difficulty: they must keep their platforms reliable for the learners who depend on them while undertaking AI integrations that may, by their nature, introduce instability.

Skills gap and talent shortage

The shortage of skilled AI talent has consistently been identified as one of the most significant barriers in this field. Chinnaraju (2025) noted that AI development requires specialised knowledge of machine learning, data science and software engineering, skills that are scarce globally and particularly difficult to secure in the startup ecosystem, where compensation cannot compete with that offered by large technology firms. Mohd Rasdi and Umar Baki (2025) observed that the absence of in-house expertise has forced many SMEs to rely on external vendors, introducing additional costs and dependency risks. This reliance also raises questions about knowledge transfer: startups that outsource AI implementation risk remaining dependent on vendors without developing internal capabilities, a structural vulnerability that the literature has identified but not examined in depth.

Financial constraints

Implementation costs represent a substantial barrier for early-stage startups. The costs of acquiring AI tools, hiring personnel, building infrastructure and maintaining systems are often prohibitive (Chinnaraju, 2025). Dwivedi et al. (2021) noted that while the long-term cost savings of AI are increasingly well established, the upfront investment required creates a significant barrier for startups without venture capital backing. This tension between short-term cost and long-term benefit is a fundamental challenge, particularly because the literature rarely specifies with any precision the time horizon over which AI investments yield returns.

Resistance to change

Organisational resistance constitutes a further barrier. Mohd Rasdi and Umar Baki (2025) reported that employees and managers have resisted AI adoption owing to fear of displacement, limited understanding of AI's capabilities and concerns about reliability. Ojika et al. (2024) similarly note that AI-driven automation tends to stall when staff and management resist it, concluding that change management must be planned from the outset rather than added as an afterthought.

Maintaining educational quality

In the EdTech context specifically, a distinct barrier concerns the preservation of pedagogical integrity when AI is used for content generation and learner interaction. UNESCO (2023) cautions that while AI can automate content production, the quality and accuracy of its outputs remain a

concern, particularly in subject areas that demand nuance or cultural sensitivity. For startups whose differentiation rests on content quality, the risks of factual errors, biased outputs or material lacking pedagogical depth are considerable (Lu, Zhao and Zheng, forthcoming). Addressing these risks requires robust human oversight mechanisms, which, paradoxically, reintroduce the labour costs that AI adoption was intended to reduce.

Regulatory uncertainty

Huang (2023) notes that there are as yet no harmonised international standards governing AI in education, creating considerable uncertainty for startups operating across multiple jurisdictions. The European Union sits at the stricter end of the spectrum, with the General Data Protection Regulation (GDPR) imposing more demanding data protection obligations than those found in many other markets. Navigating this patchwork of rules requires legal expertise that early-stage ventures can rarely afford (Chinnaraju, 2025). The recent introduction of the EU AI Act has added a further layer of regulatory complexity that the existing literature has yet to address fully.

Interconnected barriers

These barriers should not be read as isolated obstacles, because they reinforce one another. A startup facing financial constraints is, for that reason alone, less able to hire technical talent; this increases its dependence on external vendors, which raises costs and reduces control, placing further strain on already limited financial resources. This compounding dynamic is underexplored in the existing literature, which has tended to present barriers as discrete categories rather than as elements of an interconnected, self-reinforcing system.

2.5 Ethical Considerations in AI-Enabled Education

The deployment of AI in educational settings raises ethical concerns that do not map neatly onto the general adoption challenges discussed above. On privacy and data protection, Huang (2023) argues that the collection of learner data by AI systems introduces risks that are not always immediately apparent, including surveillance, inadequate consent and the potential for future misuse of data. Young learners are particularly exposed: they often lack the legal capacity to give informed consent and rarely understand how their personal data will be used once collected.

UNESCO (2023) observed that fewer than 30 per cent of countries outside Europe have comprehensive data protection laws, creating a fragmented regulatory environment for EdTech companies operating across international markets. Chinnaraju (2025) noted that compliance with regulations such as the GDPR and the California Consumer Privacy Act (CCPA) is both a legal obligation and a practical challenge, particularly for startups with limited compliance resources.

Lu, Zhao and Zheng (forthcoming) add a further ethical dimension by highlighting the risk of AI-induced erosion of critical thinking and human judgement. Their Empowerment–Entrapment Framework warns that over-reliance on AI for decision-making can lead to entrepreneurial overconfidence, reduced relational embeddedness and the gradual degradation of interpersonal skills. For EdTech startups, this concern is particularly salient in content development, where the uncritical adoption of AI-generated material risks compromising the pedagogical mission of the platform. The ethical tension between efficiency and educational integrity is acknowledged in the literature but rarely resolved, leaving practitioners to navigate these trade-offs without established guidance.

A related concern that has received insufficient attention is accountability: who is responsible when AI systems produce errors in educational contexts. An AI agent might give a learner incorrect feedback, flag a student as at risk when they are not or, more seriously, fail to identify a student who genuinely needs intervention. AI-generated content containing factual inaccuracies might reach learners without human review. In such cases, it is unclear whether responsibility lies with the startup that deployed the system, the vendor that built it or the educator who chose to use it; legally and ethically, the boundaries remain blurred (Huang, 2023). For EdTech startups, this ambiguity is not merely an academic question but a source of reputational risk and potential liability, and one for which the literature has yet to propose concrete governance frameworks.

2.6 Gaps in the Existing Literature

Taken together, the literature reviewed in this and the preceding chapter reveals several significant gaps.

First, although the general impact of AI on business operations and on education has been written about extensively, empirical work focusing on the intersection of AI agents and EdTech startups remains surprisingly thin. Existing studies tend to focus either on large enterprises (Dwivedi et al., 2021; Farayola et al., 2023) or on AI in education from a pedagogical perspective, with operations receiving far less attention (OECD, 2020; UNESCO, 2023).

Second, much of the literature on barriers to AI adoption is situated within the broader SME or enterprise context rather than reflecting the realities of early-stage EdTech ventures specifically.

Third, the existing research is predominantly conceptual or based on secondary analysis. Primary empirical research, whether quantitative surveys of practitioners or qualitative studies of their lived experience, is comparatively scarce.

Fourth, and perhaps most critically, the literature tends to treat AI adoption as a binary phenomenon. Organisations are presented as either adopting or not adopting, with little attention paid to the degree, scope or sophistication of implementation. A startup using a single chatbot to

handle customer enquiries and one that has deployed multiple autonomous agents across its entire operational workflow are fundamentally different cases, yet the literature frequently groups them under the same broad category of "AI adoption". This lack of granularity limits the analytical value of existing findings.

Fifth, the literature exhibits a pronounced publication bias towards positive outcomes. Studies of successful AI implementations clearly outnumber those reporting failures, abandoned initiatives or outcomes that fell short of expectations. The resulting survivorship bias presents a picture of AI adoption in startups that does not reflect the full reality and risks leaving practitioners with expectations higher than the evidence supports. The present study seeks to counterbalance this bias by asking respondents not only about the benefits they have experienced but also about the barriers they encountered and whether implementation proved more complex or costly than anticipated.

These gaps shape the rationale for the research design adopted in this study: a structured online questionnaire distributed to professionals working in startups, with respondents drawn predominantly from Europe and Africa. The aim is to generate empirical, experience-based insight that complements rather than repeats the theoretical literature. The methodology is set out in detail in the following chapter.

2.7 Chapter Summary

This chapter has reviewed the literature most directly relevant to the central research question. AI agents are deployed across customer support, content personalisation, administrative automation and learner analytics, although the evidence for their effectiveness is stronger in some domains than in others. The literature establishes that AI contributes to cost reduction through three main routes: task automation, improved decision-making and the enabling of scalable growth. Examined critically, however, cost-reduction claims are often presented without adequate baseline data, the upfront costs of implementation are routinely underreported and the temporal dimension of cost-benefit analysis is insufficiently addressed.

The review also identified significant implementation barriers, namely technical limitations, talent shortages, financial constraints, organisational resistance, quality concerns and regulatory uncertainty, and highlighted that these barriers interact rather than operate in isolation. Ethical considerations, particularly those concerning data privacy and the preservation of human judgement, form a further important layer. The chapter concluded by identifying the critical gaps in the existing literature that the present study is designed to address through primary quantitative research.

Chapter Three – Methodology

This chapter sets out the research methodology adopted to investigate the central research question: how EdTech startups implement AI agents to reduce operational costs and accelerate growth. It begins with the research philosophy and design, then describes the data collection method, sampling strategy and approach to data analysis. It also explains why the chosen methodology was preferred over alternatives and concludes by acknowledging the limitations of the research design. Throughout, methodological choices were guided by the need to generate primary empirical data addressing the gaps identified in the literature review, while remaining feasible within the constraints of an MSc dissertation.

3.1 Research Philosophy and Design

The study is situated within a positivist research philosophy, which holds that knowledge is best generated through the systematic collection and analysis of observable, measurable data (Saunders, Lewis and Thornhill, 2019). A positivist orientation was appropriate because the study sought to identify patterns, frequencies and relationships in AI agent adoption across a sample of startup professionals, rather than to interpret subjective experiences or construct deeply contextualised narratives. This orientation aligned well with the use of a structured questionnaire as the primary data collection instrument, enabling the collection of standardised responses that could be analysed and compared systematically across the sample.

The research adopted a quantitative, cross-sectional survey design. Cross-sectional studies collect data at a single point in time and are well suited to exploratory and descriptive research that aims to capture the current state of a phenomenon (Saunders, Lewis and Thornhill, 2019). Given that AI agent adoption in startups is evolving rapidly and that empirical evidence remains limited, a cross-sectional survey offered an efficient means of gathering data from a relatively broad and geographically dispersed sample within the time available.

3.2 Justification for the Chosen Methodology

A quantitative approach was chosen for three reasons. First, the literature review demonstrated a scarcity of primary empirical data on AI agent adoption in startups; existing work is dominated by conceptual models, secondary analyses and a small number of case studies of well-resourced platforms. Collecting structured primary data from 45 startup professionals addresses this gap directly.

Second, a quantitative approach enabled the identification of patterns across a broader sample than qualitative interviews alone would have allowed. Although interviews would have offered

richer narrative depth, the structured questionnaire captured responses from professionals across 10 countries and multiple industries, producing a dataset with greater geographic and sectoral diversity than a smaller interview-based study could have achieved.

Third, the Likert-scale and multiple-choice items produced data suitable for descriptive statistical analysis, including frequency distributions, percentages and mean scores. These forms of analysis are appropriate for an MSc dissertation and produce findings that are accessible, reproducible and directly comparable with claims made in the existing literature. The structured format also ensured consistency across all 45 responses, removing the variability that can arise from less standardised data collection methods.

3.3 Rejection of Alternative Methodologies

A purely qualitative approach based on semi-structured interviews was considered during the design phase. Qualitative methods would have offered greater insight into individual experiences and decision-making processes than a survey can provide. However, practical constraints made this approach unfeasible: participants scheduled for interviews were unable to attend at the agreed times, and repeated attempts to reschedule were unsuccessful. Rather than proceed with a qualitative sample too small to achieve the depth and saturation that rigorous thematic analysis requires, the decision was taken to focus exclusively on the survey, which by that stage had already generated a usable dataset of 45 responses.

A mixed-methods approach, combining the questionnaire with a smaller set of interviews to triangulate the findings, was also considered. Although methodologically desirable, it faced the same obstacle: interview participants could not be secured within the available timeframe. The decision to proceed with a single-method quantitative design was therefore pragmatic. It was taken in consultation with the dissertation supervisor and is consistent with the positivist philosophy the study had already adopted.

Experimental and longitudinal designs were rejected for different reasons. The study aimed to capture a snapshot of current AI adoption practices across a sample of startups; it did not seek to measure change over time or to test causal hypotheses under controlled conditions. Neither design was appropriate to the exploratory and descriptive objectives of the research.

3.4 Data Collection Method

Primary data were collected through a structured online questionnaire designed by the researcher and administered through Microsoft Forms. The questionnaire comprised 18 questions organised into five sections:

- demographic and organisational profile (Questions 1–5);

- AI agent adoption (Questions 6–9);
- impact on operations and growth (Questions 10–13);
- barriers and challenges (Questions 14–15); and
- future outlook (Questions 16–18).

The question types included multiple-choice items, select-all-that-apply checkboxes, five-point Likert-scale items, ranking and top-choice items, and a single optional open-ended question at the end of the survey. The full instrument is reproduced in Appendix A.

The questionnaire was designed to align directly with the research questions and with the themes emerging from the literature review. Questions 6 and 7 addressed the types of AI agents deployed and their operational areas, corresponding to the first supporting question. Questions 10, 11 and 13 addressed the extent and location of operational cost reduction, corresponding to the second. Questions 14 and 15 addressed the barriers and challenges encountered during implementation, corresponding to the third. This alignment ensured that every question served the research objectives and that no data were collected without a clear analytical purpose.

The questionnaire was distributed between late May and early June 2026 through three channels: direct email invitations to professionals within the researcher's professional network; a LinkedIn post targeted at startup founders and technology professionals; and an online community of startup founders. These channels were chosen to maximise the reach and diversity of the sample within the available timeframe. The survey remained open for approximately two weeks, by the end of which 45 completed responses had been received.

For the purposes of this study, AI agents were defined broadly as AI-powered systems that perform operational tasks with varying degrees of autonomy. This range extends from AI-assisted tools that augment human work, such as ChatGPT used for content drafting, to fully autonomous agents that execute multi-step workflows, such as AI chatbots that manage customer support end to end without human involvement. This definition was presented to respondents in the survey introduction to ensure a shared understanding of the term.

3.5 Sample Selection

The study employed a non-probability sampling strategy combining purposive and snowball techniques. Purposive sampling was used to target professionals meeting three inclusion criteria: they were currently working in a startup environment; they had direct experience with, or sound knowledge of, the use of AI agents or AI-powered tools within their organisation; and they were willing to complete the survey voluntarily. Snowball sampling was then layered on top, with recipients invited to forward the questionnaire to other professionals meeting the same criteria. This combination was appropriate for a study whose target population is specific and relatively

niche, and therefore difficult to reach through random sampling (Saunders, Lewis and Thornhill, 2019).

The final sample comprised 45 respondents from 10 countries, with the largest concentrations in Ghana (37.8 per cent) and Germany (28.9 per cent). Respondents represented a range of industries, of which software as a service (31.1 per cent) and education technology (24.4 per cent) were the most represented. The sample also included a spread of roles: founders and co-founders (26.7 per cent), software developers and engineers (17.8 per cent), operations managers (13.3 per cent), product managers and CTOs or technical leads (11.1 per cent each), and marketing or growth leads (6.7 per cent), among others. Just under half of respondents (48.9 per cent) worked in early-stage startups with one to ten employees, while 22.2 per cent worked in growth-stage ventures. This diversity of role, geography, industry and startup maturity strengthened the descriptive value of the dataset, although the non-probability sampling approach means that the findings cannot be generalised to the wider population of startup professionals.

3.6 Data Analysis Approach

The survey data were exported from Microsoft Forms as a CSV file and analysed using descriptive statistical techniques, focusing on frequency distributions, percentages and mean scores for the Likert-scale items. Responses to multiple-choice questions were tallied and expressed as raw frequencies and as proportions of the total sample. For select-all-that-apply and top-choice questions, the number of respondents selecting each option was counted and expressed as a percentage of all 45 respondents; because respondents could select more than one option, these percentages sum to more than 100 per cent.

The Likert-scale items (Questions 10, 12 and 15) were analysed by calculating the frequency distribution across the five-point scale and computing the mean score for each item. The use of means for ordinal Likert data has been the subject of methodological debate; nevertheless, calculating mean scores for descriptive purposes is widely accepted in survey-based research and was adopted here to facilitate comparison across items and with the existing studies reviewed in the literature (Saunders, Lewis and Thornhill, 2019).

The optional open-ended question (Question 18) was analysed separately. Responses were read carefully and recurring themes were identified using a simple content analysis approach. Given the small number of open-ended responses and the absence of a formal qualitative methodology, these responses were treated as supplementary data that enrich the quantitative findings rather than as standalone qualitative evidence.

The findings are presented in Chapter Four through a combination of tables, bar charts and narrative interpretation. The analysis follows the sequence of the research questions to ensure

that each is addressed systematically and that the relationship between the primary data and the existing literature is made explicit in the discussion.

3.7 Ethical Considerations

The research was conducted in accordance with the ethical guidelines of the Berlin School of Business and Innovation. Participation was entirely voluntary. Respondents were informed of the purpose of the study, that their responses would remain anonymous and that they could withdraw at any point before submitting the form. No personally identifiable information was collected: the Microsoft Forms settings were configured so that names and email addresses were not recorded. The introductory text of the survey served as the informed consent mechanism; by proceeding to complete the questionnaire, respondents indicated their consent to take part.

3.8 Limitations of the Methodology

Several limitations of the research methodology should be acknowledged.

The first concerns sample size. The 45 respondents were sufficient for descriptive analysis at MSc level but not for inferential statistical testing or claims of statistical significance. The findings should therefore be read as indicative of patterns and trends rather than as generalisable conclusions about the wider population of startup professionals. Relatedly, the unit of analysis is the individual respondent rather than the firm; because the survey was anonymous, it cannot be ruled out that more than one respondent worked for the same startup.

The second concerns sampling. Non-probability sampling introduces the possibility of selection bias. Respondents were drawn primarily from the researcher's professional network and from an online startup community whose members were already engaged with the topic, which may have skewed the sample towards individuals favourably disposed to AI adoption. The geographic concentration in Ghana and Germany further limits representativeness. In addition, although the study focuses on EdTech, respondents working in EdTech made up 24.4 per cent of the sample ($n = 11$); findings for the sample as a whole therefore describe startups more broadly, and conclusions about EdTech specifically should be read with this in mind.

The third concerns the reliance on self-reported data, which carries a risk of response bias, including social desirability bias. Respondents may have overstated the benefits of AI adoption or understated the challenges encountered. The structured format, while ensuring consistency, also constrained depth and left little room to explore unexpected themes or nuances that qualitative methods would have been better placed to capture. The broad definition of AI agents adopted for the survey, which encompassed AI-assisted tools as well as fully autonomous agents, also means that the results describe AI-enabled automation generally rather than autonomous agents alone.

The fourth concerns time. Because the study captured data at a single point in time, it could not assess how AI agent adoption and its outcomes evolve. A longitudinal design would have provided richer insight into adoption trajectories but was not feasible within the constraints of an MSc dissertation.

Despite these limitations, the methodology was appropriate for the exploratory and descriptive objectives of the study. The data collected make a meaningful empirical contribution to a topic on which primary evidence remains scarce, and while the findings are not generalisable in a statistical sense, they offer valuable insight into the patterns, outcomes and challenges associated with AI agent adoption in startups.

Chapter Four – Findings, Analysis and Discussion

This chapter presents the results of the primary data collection, analyses the patterns that emerged from the survey responses and discusses the findings in relation to the literature reviewed in Chapters One and Two. It is organised into three sections. Section 4.1 presents the findings from the structured questionnaire completed by 45 startup professionals. Section 4.2 analyses these findings by interpreting the patterns and their implications. Section 4.3 synthesises the primary data with the secondary literature to discuss the study's broader significance for EdTech startup management and AI strategy.

4.1 Findings

The findings in this section are drawn from the responses of 45 professionals working in startups across 10 countries. They are presented in the sequence of the research questions: respondent profile, AI agent adoption patterns, impact on cost reduction and growth, barriers to implementation and future outlook.

4.1.1 Respondent Profile

The survey attracted responses from a diverse range of startup professionals. Founders and co-founders formed the largest group, comprising 26.7 per cent of the sample ($n = 12$), followed by software developers and engineers (17.8 per cent, $n = 8$), operations managers (13.3 per cent, $n = 6$), product managers and CTOs or technical leads (11.1 per cent each, $n = 5$) and marketing or growth leads (6.7 per cent, $n = 3$). The remaining respondents comprised interns (6.7 per cent), HR professionals (4.4 per cent) and one finance professional (2.2 per cent). The distribution of roles is presented in Table 1.

Table 1: Respondent roles (n = 45)

Role	Number (n)	Percentage (%)
Founder / Co-founder	12	26.7
Software Developer / Engineer	8	17.8
Operations Manager	6	13.3
Product Manager	5	11.1
CTO / Technical Lead	5	11.1
Marketing / Growth Lead	3	6.7
Intern	3	6.7
HR Professional	2	4.4
Finance	1	2.2

Geographically, the sample was concentrated in two countries. Ghana accounted for the largest share at 37.8 per cent (n = 17), followed by Germany at 28.9 per cent (n = 13). The remaining respondents were based in Kenya (8.9 per cent); Spain, Nigeria, China and the United States (4.4 per cent each); and Greece, France and Sweden (one respondent each). The distribution is shown in Table 2.

Table 2: Geographic distribution (n = 45)

Country	Number (n)	Percentage (%)
Ghana	17	37.8
Germany	13	28.9
Kenya	4	8.9
Spain	2	4.4
Nigeria	2	4.4
China	2	4.4
United States	2	4.4
Greece	1	2.2
France	1	2.2
Sweden	1	2.2

In terms of industry, software as a service (SaaS) was the most represented sector at 31.1 per cent ($n = 14$), followed by education technology (EdTech) at 24.4 per cent ($n = 11$), financial technology at 15.6 per cent ($n = 7$), health technology at 11.1 per cent ($n = 5$) and e-commerce at 8.9 per cent ($n = 4$). The distribution is shown in Table 3. Just under half of respondents (48.9 per cent) worked in early-stage startups with one to ten employees and 22.2 per cent in growth-stage ventures with 11 to 50 employees; the remaining 28.9 per cent were at either the pre-revenue stage or the scaling stage (51 or more employees). Most startups had been operating for between one and five years (62.2 per cent), with 20.0 per cent operating for less than one year and 17.8 per cent for more than five years.

Table 3: Industry distribution ($n = 45$)

Industry	Number (n)	Percentage (%)
Software as a Service (SaaS)	14	31.1
Education Technology (EdTech)	11	24.4
Financial Technology (FinTech)	7	15.6
Health Technology (HealthTech)	5	11.1
E-commerce	4	8.9
Other (logistics, interior design, telecoms)	4	8.9

4.1.2 AI Agent Adoption Patterns

AI agent adoption was widespread across the sample. A substantial majority of respondents (75.6 per cent, $n = 34$) reported that their startups were actively using AI agents in their operations, and a further 15.6 per cent ($n = 7$) were implementing or piloting AI solutions. Three respondents (6.7 per cent) indicated that adoption was planned but not yet started, and one respondent reported no current use of AI agents. Taking active users and those piloting together, 91.1 per cent of the sample were engaged in some form of AI agent deployment. Adoption status is presented in Table 4.

Table 4: AI agent adoption status ($n = 45$)

Status	Number (n)	Percentage (%)
Yes, actively using	34	75.6
Currently implementing / piloting	7	15.6
Planned but not yet started	3	6.7
No	1	2.2

Customer support and chatbots were the most common area of deployment, reported by 62.2 per cent of respondents (n = 28). They were followed by administrative automation, covering scheduling, reporting and enrolment (55.6 per cent, n = 25); content generation and personalisation (53.3 per cent, n = 24); software development and code generation (46.7 per cent, n = 21); marketing and lead generation (40.0 per cent, n = 18); data analytics and insights (33.3 per cent, n = 15); human resources and recruitment (24.4 per cent, n = 11); and onboarding and user engagement (22.2 per cent, n = 10). The distribution is shown in Table 5 and Figure 1.

Table 5: Operational areas of AI agent deployment (n = 45, multi-select)

Operational area	Number (n)	Percentage (%)
Customer support / chatbots	28	62.2
Administrative automation	25	55.6
Content generation / personalisation	24	53.3
Software development / code generation	21	46.7
Marketing and lead generation	18	40.0
Data analytics and insights	15	33.3
Human resources / recruitment	11	24.4
Onboarding and user engagement	10	22.2




Figure 1: Operational areas of AI agent deployment (share of respondents, multi-select)

On the platform side, OpenAI products, including ChatGPT and the GPT API, were the most widely used, reported by 71.1 per cent of respondents (n = 32). Anthropic's Claude was second at 64.4 per cent (n = 29), followed by Google products, including Gemini and Vertex AI, at 46.7 per cent (n = 21) and Microsoft Copilot at 22.2 per cent (n = 10). Custom-built AI models and n8n, an open-source workflow automation platform, were each reported by 15.6 per cent (n = 7), and Zapier or Make by 11.1 per cent (n = 5). The distribution of platforms is presented in Table 6 and Figure 2.

Table 6: AI tools and platforms used (n = 45, multi-select)

Platform	Number (n)	Percentage (%)
OpenAI (ChatGPT, GPT API)	32	71.1
Anthropic (Claude)	29	64.4
Google (Gemini, Vertex AI)	21	46.7
Microsoft Copilot	10	22.2
Custom-built AI models	7	15.6
n8n	7	15.6
Zapier / Make	5	11.1




Figure 2: AI tools and platforms used (share of respondents, multi-select)

Respondents were asked to identify their top three motivations for adopting AI agents. Reducing operational costs was the most frequently selected motivation, chosen by 73.3 per cent of respondents (n = 33), followed by improving efficiency and speed (60.0 per cent, n = 27). Scaling without increasing headcount and staying competitive were selected by 33.3 per cent (n = 15) and

31.1 per cent (n = 14) respectively, and founder or leadership interest in AI by 28.9 per cent (n = 13). Investor expectations were selected by only 6.7 per cent (n = 3).

4.1.3 Impact on Cost Reduction and Growth

Respondents rated the extent to which AI agent adoption had reduced their operational costs on a five-point Likert scale. The mean score was 3.82 out of 5, indicating a clearly positive perception of cost reduction across the sample. Almost a third of respondents (31.1 per cent, n = 14) reported a transformative reduction, a further 31.1 per cent (n = 14) a significant reduction and 28.9 per cent (n = 13) a moderate reduction. Three respondents (6.7 per cent) reported a slight reduction, and one (2.2 per cent) reported no reduction at all. The distribution is presented in Table 7.

Table 7: Impact of AI agents on operational cost reduction (n = 45)


Response	Number (n)	Percentage (%)
5 = Transformative reduction	14	31.1
4 = Significant reduction	14	31.1
3 = Moderate reduction	13	28.9
2 = Slight reduction	3	6.7
1 = No reduction at all	1	2.2
Mean score		3.82

Content creation was the area in which the most respondents reported significant cost savings (62.2 per cent, n = 28), followed by customer support (51.1 per cent, n = 23), marketing (48.9 per cent, n = 22) and administrative tasks (46.7 per cent, n = 21). Software development was cited by 37.8 per cent (n = 17), data analysis by 31.1 per cent (n = 14) and human resources by 15.6 per cent (n = 7).

The mean score for AI agents' contribution to growth was also 3.82 out of 5. The largest group of respondents (48.9 per cent, n = 22) rated the contribution as significant, followed by moderate (24.4 per cent, n = 11), transformative (20.0 per cent, n = 9) and slight (6.7 per cent, n = 3). Unlike the cost reduction item, on which one respondent reported no reduction, no respondent indicated that AI agents had made no contribution to growth. The results are presented in Table 8, and Figure 3 compares the two distributions.

Table 8: Contribution of AI agents to startup growth (n = 45)

Response	Number (n)	Percentage (%)
5 = Transformative contribution	9	20.0
4 = Significant contribution	22	48.9
3 = Moderate contribution	11	24.4
2 = Slight contribution	3	6.7
1 = No contribution	0	0.0
Mean score		3.82

**Figure 3:** Ratings of AI agents' impact on operational costs and on growth

When asked to select the single statement that best reflected their experience, the largest group of respondents (26.7 per cent, n = 12) chose "AI agents have enabled us to scale without proportionally increasing staff". This was followed by "AI agents have improved the quality of our product or service" (24.4 per cent, n = 11) and "AI agent implementation was more complex than we initially expected" (20.0 per cent, n = 9). The remaining 13 respondents (28.9 per cent) were distributed across the other three statements: a greater focus on higher-value strategic work, improved response times and higher-than-anticipated implementation costs. That implementation complexity was selected almost as often as improved product quality is a notable finding, explored further in the analysis below.

4.1.4 Barriers to Implementation

The most frequently reported barrier to AI agent implementation was data privacy and regulatory compliance, cited by 46.7 per cent of respondents (n = 21). It was followed closely by high implementation costs (42.2 per cent, n = 19) and data quality or availability issues (37.8 per cent, n = 17). A lack of technical expertise, commonly described as the skills gap, was reported by 31.1 per cent (n = 14), integration with existing systems by 26.7 per cent (n = 12) and vendor dependency by 24.4 per cent (n = 11). Concerns about AI output quality and accuracy were raised by 15.6 per cent (n = 7), resistance from team members by 8.9 per cent (n = 4) and difficulty measuring return on investment by 4.4 per cent (n = 2). Only 8.9 per cent of respondents (n = 4) reported facing no significant challenges. The barriers are presented in Table 9 and Figure 4.

Table 9: Barriers to AI agent implementation (n = 45, multi-select)

Barrier	Number (n)	Percentage (%)
Data privacy / regulatory compliance	21	46.7
High implementation costs	19	42.2
Data quality / availability issues	17	37.8
Skills gap / lack of expertise	14	31.1
Integration with existing systems	12	26.7
Vendor dependency	11	24.4
AI output quality / accuracy	7	15.6
Resistance from team members	4	8.9
No significant challenges	4	8.9
Difficulty measuring ROI	2	4.4




Figure 4: Barriers to AI agent implementation (share of respondents, multi-select)

Asked specifically how significant data privacy and regulatory concerns had been, respondents gave a mean score of 3.22 out of 5. The largest group (40.0 per cent, $n = 18$) rated these concerns as very significant, followed by moderately significant (37.8 per cent, $n = 17$); only 8.9 per cent ($n = 4$) rated them as not significant at all.

4.1.5 Future Outlook

The data revealed a strongly positive outlook on future AI expansion. Over half of respondents (55.6 per cent, $n = 25$) planned to expand their use of AI agents significantly over the next 12 months, and a further 40.0 per cent ($n = 18$) planned a moderate expansion. One respondent planned no change and one planned to reduce AI usage.

Asked what would most help their startup adopt or expand AI agents (selecting up to two options), respondents most often chose more affordable AI tools (64.4 per cent, $n = 29$), followed by better availability of technical talent (40.0 per cent, $n = 18$), internal training and upskilling (35.6 per cent, $n = 16$), clearer regulatory guidance (33.3 per cent, $n = 15$) and proven case studies and best practices (28.9 per cent, $n = 13$).

4.2 Analysis

The findings reveal several patterns that warrant closer interpretation. This section analyses the key themes emerging from the data and considers what they indicate about the current state of AI agent adoption in startups.

4.2.1 The Ubiquity of AI Adoption

The near-universal adoption rate of 91.1 per cent (combining active users and those piloting) was the most striking finding of the survey. It suggests that, for the startups sampled, the question is no longer whether to adopt AI agents but how to deploy and scale them. The finding indicates that AI agents have moved beyond the early adopter phase and into the early majority phase of the adoption curve described by Rogers (2003), at least within the population sampled. That this level of adoption was observed in a sample spanning multiple industries and countries reinforces the interpretation that AI adoption in startups has reached a tipping point, although the broad definition of AI agents used in the survey, and the likely favourable disposition of respondents towards AI, mean that this figure should be read with some caution.

4.2.2 Cost Reduction as a Primary Driver

The prominence of cost reduction, both as the most frequently selected motivation for adoption (73.3 per cent) and as an area of strong perceived impact (mean 3.82 out of 5), indicates that startups approach AI primarily as an operational efficiency tool rather than as a source of product innovation or competitive differentiation. This pragmatic orientation is consistent with the resource constraints that characterise early-stage ventures. The concentration of reported savings in content creation (62.2 per cent), customer support (51.1 per cent) and marketing (48.9 per cent) indicates where AI agents are being deployed most effectively. These functions involve repetitive, language-based tasks, precisely the kind of work to which the current capabilities of large language models are well suited.

The lower impact reported in human resources (15.6 per cent) and data analysis (31.1 per cent) indicates that AI adoption has not yet penetrated all operational domains equally. Functions requiring more complex judgement, contextual sensitivity or specialised data infrastructure appear slower to benefit from AI agent deployment. This is particularly likely in early-stage startups, where the data volumes and workflow complexity needed to justify AI-driven analytics are often not yet present.

4.2.3 Growth Without Proportional Headcount

One of the more revealing findings was that 26.7 per cent of respondents selected scaling without proportionally increasing staff as the statement that best reflected their experience, making it the most common response. This indicates that AI agents are functioning as a workforce multiplier,

allowing startups to grow their user base, output and revenue without the corresponding increases in team size that traditional scaling models would require. The pattern has direct implications for the financial sustainability of early-stage startups, as AI-enabled operational leverage may be reducing the burn rate that typically constrains early growth.

The finding that 20.0 per cent of respondents found implementation more complex than expected tempers this positive picture considerably. It suggests that the operational benefits of AI agents carry hidden costs, such as implementation time, technical debugging and workflow redesign, that are not always anticipated at the point of adoption. The gap between expected simplicity and experienced complexity has real implications for how startups should plan and resource their AI adoption.

4.2.4 Interconnected Barriers

The barrier data reveal an important pattern. The three most frequently reported barriers, namely data privacy (46.7 per cent), high costs (42.2 per cent) and data quality issues (37.8 per cent), were reported at similar frequencies, suggesting that startups face not a single dominant barrier but a cluster of related challenges. These barriers are plausibly linked. A startup struggling with data quality is likely also to face higher implementation costs because of the additional data cleaning and preparation required. Similarly, the skills gap (31.1 per cent) is likely to compound the cost barrier, as startups without in-house expertise must rely on external vendors at greater expense. Although the descriptive data cannot establish these links directly, the pattern is consistent with the interconnected dynamic identified in the literature review.

The low frequency of resistance from team members (8.9 per cent) is notable and suggests that cultural barriers to AI adoption within startups are considerably lower than those reported in the broader SME literature. This probably reflects the younger, more technology-oriented workforce typical of startups, as well as a founder-driven decision-making culture in which AI adoption is championed from the top.

4.2.5 The Demand for Affordability

That 64.4 per cent of respondents identified more affordable AI tools as one of the most helpful factors for expanding AI adoption reinforces the cost sensitivity of early-stage startups. Despite strongly positive perceptions of AI's impact on costs, the cost of the tools themselves remains a barrier to deeper adoption. This apparent paradox, that AI reduces costs yet is itself costly to implement, is a defining tension in the data, with practical implications for AI tool providers, policymakers and startup managers alike.

4.2.6 Platform Dominance and Vendor Concentration Risk

The platform data reveal a high degree of concentration around a small number of commercial AI providers. OpenAI (71.1 per cent) and Anthropic (64.4 per cent) were by far the most widely used, with Google third at 46.7 per cent. Custom-built AI models, by comparison, were reported by only 15.6 per cent of respondents. Most startups are therefore building their AI capabilities on third-party infrastructure rather than developing proprietary systems. While this approach lowers the technical barrier to entry and enables rapid deployment, it also creates a dependency structure in which core operational processes rely on platforms controlled by external entities. A change in pricing, terms of service or API availability by any of these dominant providers could have immediate and significant operational consequences for the startups that depend on them.

4.2.7 The EdTech Subset

Although the survey captured responses across multiple industries, the EdTech subset (24.4 per cent, $n = 11$) is of particular relevance to the research question. The patterns within this subset were broadly consistent with those observed in the overall sample: content generation, customer support and administrative automation were the most commonly reported deployment areas, and perceptions of cost reduction were positive. Given the small size of the subset, these observations are indicative only. EdTech respondents also operate within a context of additional complexity not present to the same degree in other sectors: the need to maintain pedagogical quality in AI-generated content, the ethical sensitivity of learner data and regulatory requirements specific to education. These sector-specific dimensions were not fully captured by the quantitative instrument and represent an area in which future qualitative research could provide valuable insight.

4.3 Discussion

This section synthesises the primary findings with the secondary literature reviewed in Chapters One and Two. Its aim is to situate the survey data within the broader academic discourse on AI adoption in startups and to evaluate the extent to which the findings support, extend or challenge existing theoretical and empirical claims.

4.3.1 AI Agents and Operational Cost Reduction

The survey data support the argument that AI agents contribute to operational cost reduction in startups. The mean score of 3.82 out of 5, together with the finding that 62.2 per cent of respondents rated the impact as significant or transformative, is consistent with the claims of Ojika et al. (2024) and Chopra (2025), both of whom argue that AI-driven automation generates measurable cost savings through the replacement of manual tasks, improved decision-making accuracy and continuous process optimisation.

The concentration of savings in content creation, customer support and marketing aligns with the operational domains identified by Chinnaraju (2025) and supports the Lean Scalability Model's prediction that AI agents are most impactful in functions involving high-volume, language-based tasks. The finding extends the existing literature by providing primary, cross-industry empirical data in support of claims previously argued mainly on the basis of case studies and conceptual models.

The data also lend support to the critical observations raised in the literature review. That 20.0 per cent of respondents found implementation more complex than expected, and that 42.2 per cent cited high implementation costs as a barrier, echoes the concern identified in Chapter Two that the cost-reduction literature systematically underreports the upfront costs of AI implementation. The temporal dimension of cost-benefit analysis, flagged as a gap in the literature, is also relevant here: the survey captured perceptions at a single point in time and cannot distinguish between startups that have already recouped their implementation investment and those still operating at a net cost. Longitudinal research would be required to resolve this ambiguity.

4.3.2 AI Agents and Startup Growth

The growth data (mean 3.82 out of 5) support the broader argument of Giuggioli and Pellegrini (2023) that AI enhances entrepreneurial performance across multiple dimensions. That scaling without proportional increases in headcount was the most frequently selected experience is consistent with the Resource-Based View's framing of AI as a strategic resource that enables competitive advantage through operational leverage (Barney, 1991; Chinnaraju, 2025). On this interpretation, AI agents function not merely as cost-cutting tools but as enablers of a fundamentally different growth model, one in which headcount and output are decoupled.

That investor expectations were selected as a motivation by only 6.7 per cent of respondents is a noteworthy departure from the assumption, present in parts of the literature, that AI adoption in startups is driven by investors seeking technology-forward portfolio companies. The data suggest instead that adoption is predominantly internally motivated, driven by operational need and founder initiative rather than by external stakeholders. This finding adds a new empirical dimension to the literature on AI adoption motivations in startups.

Farayola et al. (2023) argued that the most successful AI-driven business models align AI capabilities with clearly defined customer needs rather than adopting AI for its own sake. The present data provide empirical support for this claim. The concentration of deployment in functions linked to customer-facing operations, namely content creation, customer support and marketing, indicates that the startups sampled are deploying AI where it addresses concrete operational pain points rather than following a technology trend. The pragmatic orientation of the sample, in which cost reduction and efficiency dominate as motivations, further reinforces the

argument that alignment between AI capability and business need is a prerequisite for successful adoption.

Andersson and Rosenqvist (n.d.) observed that AI integration is transforming EdTech business models and requires new approaches to value creation and capture. The survey data provide partial support for this claim within the EdTech subset. Positive perceptions of AI's contribution to growth, combined with deployment across content, support and analytics functions, suggest that AI is becoming embedded in the operational core of EdTech startups rather than remaining a peripheral enhancement. Whether this constitutes a fundamental transformation of the business model, as Andersson and Rosenqvist argue, or merely an incremental improvement in operational efficiency cannot be determined from the quantitative data alone and remains a question for future qualitative investigation.

4.3.3 Barriers: Privacy, Cost and Their Interaction

The barrier findings are broadly consistent with the existing literature. The prominence of data privacy concerns (46.7 per cent), the most frequently reported barrier, aligns with the ethical and regulatory complexities of deploying AI in contexts involving personal data highlighted by Huang (2023) and UNESCO (2023). That 40.0 per cent of respondents rated privacy concerns as very significant also reinforces Chinnaraju's (2025) observation that compliance with regulations such as the GDPR and CCPA is a real practical challenge for startups with limited legal and compliance resources.

The skills gap (31.1 per cent) and high implementation costs (42.2 per cent) are consistent with the findings of Dwivedi et al. (2021) and Mohd Rasdi and Umar Baki (2025), both of which identified financial and human capital constraints as significant barriers to AI adoption in smaller firms. The present data are also consistent with the view, developed in Chapter Two, that these barriers interact rather than operate in isolation: the skills gap forces reliance on external vendors, that reliance increases costs, and those costs consume capital that could otherwise be invested in product development or market expansion. The similar frequencies of the leading barriers fit this interpretation, although establishing the interaction formally would require analysis of how barriers co-occur within individual startups, which the descriptive design of this study did not permit.

The low level of resistance from team members (8.9 per cent) represents a notable divergence from the existing literature. Mohd Rasdi and Umar Baki (2025) and Ojika et al. (2024) both identified organisational resistance as a significant barrier in SMEs. The substantially lower figure in the present sample probably reflects the distinctive organisational culture of startups, which are generally smaller, younger, more technology-oriented and more willing to experiment with new tools than the SMEs studied by Mohd Rasdi and Umar Baki. This suggests that barrier

frameworks developed in broader SME contexts should not be applied uncritically to startups; cultural and structural differences must first be taken into account.

4.3.4 Theoretical Implications

The findings carry implications for each of the four theoretical frameworks discussed in Chapter One.

The **Technology Acceptance Model** (Davis, 1989) is partially supported. The high adoption rate and positive perceptions of cost reduction suggest that perceived usefulness is high, while the complexity reported by 20.0 per cent of respondents and the prominence of the skills gap indicate that perceived ease of use remains a constraint. The data suggest that, in the startup context, perceived usefulness may override perceived ease of use as the dominant driver of adoption, a nuance the original model does not fully capture.

The **Resource-Based View** (Barney, 1991) is supported by the finding that AI agents enable startups to achieve growth without proportional increases in headcount, with AI functioning as a strategic resource that creates competitive advantage. However, the widespread use of commercially available platforms, with OpenAI at 71.1 per cent and Anthropic at 64.4 per cent, raises an important question: does AI constitute a rare and inimitable resource, as the RBV requires, or has it become a tool available to anyone willing to pay for a subscription? The data suggest that competitive advantage may lie not in access to the technology itself, which is increasingly commoditised, but in the capability to integrate and deploy it effectively within the specific operational context of each startup.

Diffusion of Innovation Theory (Rogers, 2003) is supported by the near-universal adoption rate observed, which suggests that AI has moved well beyond the innovator and early adopter phases. With 91.1 per cent of respondents either actively using AI agents or piloting them, the technology appears to be well into the early majority phase, at least within the startup population captured by this survey. The strong intention to expand AI usage, with 95.6 per cent of respondents planning moderate or significant expansion, suggests that diffusion is ongoing and has not yet reached saturation.

The **Dynamic Capabilities Framework** (Teece, Pisano and Shuen, 1997) is supported by evidence that startups are using AI agents to build adaptive operational capacity. Deployment across multiple operational functions, combined with positive perceptions of growth, suggests that startups are using AI not as a static tool but as a dynamic capability enabling continuous adaptation to their environment. The limitations identified in the literature review regarding the framework's applicability to early-stage ventures nonetheless remain relevant: the data cannot distinguish between startups genuinely building dynamic capabilities and those simply adopting available tools without strategic integration.

4.3.5 The Empowerment–Entrapment Tension

The Empowerment–Entrapment Framework proposed by Lu, Zhao and Zheng (forthcoming) is perhaps the theoretical model most directly supported by the data. The simultaneous presence of strongly positive outcomes (cost reduction, growth and the ability to scale) and significant barriers (privacy, cost, complexity and vendor dependency) reflects precisely the double-edged dynamic the framework describes. Startups are empowered by AI to achieve outcomes that would not have been possible without automation, yet they are also becoming increasingly dependent on third-party AI platforms, vulnerable to regulatory change and exposed to implementation complexities that strain their limited resources.

The finding that vendor dependency was reported as a barrier by 24.4 per cent of respondents, while 71.1 per cent relied on OpenAI and 64.4 per cent on Anthropic, illustrates this tension concretely. Startups are building core operational processes on third-party AI infrastructure over which they have limited control, creating precisely the structural vulnerability the framework anticipates. Changes to pricing, API policies or platform capabilities could have disproportionate effects on startups that have integrated these tools deeply into their workflows.

The entrapment dimension was also evident in the open-ended responses. One respondent noted that the cost of AI tokens continues to rise, suggesting that the financial benefits of AI adoption could erode over time as providers adjust their pricing models. Although drawn from a single response, this observation points to a broader structural risk. Startups that have optimised their operations around AI-driven workflows at current price points could become locked into dependencies that grow progressively more expensive, at which point the cost of switching platforms or reverting to manual processes may be prohibitive. In this sense, the Empowerment–Entrapment Framework provides the most nuanced lens through which to interpret the data, capturing the coexistence of genuine benefit and emerging vulnerability that characterises AI adoption in the startup context.

4.3.6 Regional Considerations

The geographic distribution of the sample, with 37.8 per cent of respondents in Ghana and 28.9 per cent in Germany, offered an opportunity to consider whether regional context shapes AI adoption patterns. Although the sample size did not permit rigorous cross-regional comparison, the data are broadly consistent with Obiya's (2024) observation that AI adoption in African startups is shaped by infrastructural and financial constraints that differ from those in European contexts. That more affordable AI tools was the most frequently selected enabler across the full sample (64.4 per cent) suggests that cost sensitivity is a universal concern, although its relative weight is likely to be greater where venture capital and digital infrastructure are less accessible.

The OECD's (2020) observation that AI-enabled translation tools help EdTech companies scale across linguistic boundaries is also relevant, as the cross-continental composition of the sample demonstrates that AI adoption in startups is not confined to a single geographic or economic context. The findings suggest that the broad patterns of AI adoption, including the types of tools used, the operational areas targeted and the barriers encountered, are similar across regions, even if their severity and configuration differ according to local conditions.

4.3.7 Practical Implications for EdTech Startup Managers

Beyond their theoretical significance, the findings carry several practical implications for EdTech startup managers considering or engaged in AI agent implementation.

First, the most immediate and accessible returns on AI investment are found in content creation, customer support and marketing, the functions in which large language model capabilities are most mature and in which automating language-based tasks yields rapid savings. EdTech startups seeking to maximise their return on AI investment would be well advised to prioritise these functions before deploying AI in more complex domains such as analytics.

Second, the barrier data underline the importance of planning for implementation complexity and cost. With one in five respondents finding implementation more complex than anticipated and over two in five identifying high costs as a barrier, managers should budget not only for tool subscriptions but also for the time, expertise and workflow redesign required for effective integration.

Third, the vendor dependency finding highlights the strategic risk of building core operations on a single AI platform. EdTech startup managers would benefit from maintaining flexibility across multiple providers or investing in modular architectures that allow platforms to be switched without disrupting operations.

Finally, the data privacy finding has particular implications for the EdTech sector, where the sensitivity of learner data and the regulatory requirements of jurisdictions such as the European Union demand careful attention to compliance. That 46.7 per cent of respondents cited data privacy as a barrier reinforces the message that AI adoption in educational contexts cannot be pursued on the basis of operational efficiency alone. It requires concurrent investment in governance, transparency and compliance frameworks that protect learners' interests and meet the legal obligations of each jurisdiction in which the startup operates.

4.3.8 Summary of Discussion

This section has shown that the primary survey data are broadly consistent with the theoretical and empirical claims of the existing literature, while extending the evidence base in several directions. The findings provide primary empirical support for the argument that AI agents reduce

operational costs and contribute to growth in startups, confirm the relevance of TAM, the RBV, Diffusion of Innovation and Dynamic Capabilities to the startup context, and corroborate the barriers identified in the literature. At the same time, the data reveal nuances that the literature has not fully captured: the low level of internal resistance in startup environments, the predominantly internal rather than investor-driven motivation for adoption, the interconnected character of implementation barriers, and the emerging tension between the benefits of AI-enabled operational leverage and the risks of vendor dependency. Together, these contributions position the present study as a meaningful addition to the nascent empirical literature on AI agent adoption in startups.

Concluding Remarks

This dissertation set out to investigate how EdTech startups implement AI agents in practice, focusing on whether these implementations reduce operational costs and accelerate growth in the way the broader literature suggests, and examining the principal challenges and outcomes associated with the process. The research was motivated by the observation that, although AI adoption in business has been discussed extensively in both academic and industry literature, primary empirical research focusing on the intersection of AI agents and EdTech startups remains scarce. To address this, the study conducted a quantitative survey of 45 startup professionals across 10 countries, with the aim of generating evidence-based insights that contribute to closing the gap identified in the literature review.

5.1 Summary of Key Findings

AI agent adoption among the startups sampled was close to universal: 91.1 per cent of respondents reported either actively using AI-powered tools or piloting them with a view to wider deployment. The most common areas of deployment were customer support, administrative automation, content generation and software development. OpenAI and Anthropic were the dominant platforms, used by 71.1 per cent and 64.4 per cent of respondents respectively. Reducing operational costs was the most frequently selected motivation for adoption (73.3 per cent), followed closely by improving efficiency and speed (60.0 per cent).

The perceived impact of AI agent adoption on costs was strongly positive, with a mean score of 3.82 out of 5 and 62.2 per cent of respondents rating the impact as significant or transformative. The greatest savings were reported in content creation, customer support and marketing. The contribution to growth was similarly positive, also with a mean score of 3.82 out of 5, and the most commonly reported experience was the ability to scale operations without proportionally increasing staff.

The research also identified significant barriers. Data privacy and regulatory compliance was the most frequently cited challenge (46.7 per cent), followed by high implementation costs (42.2 per cent), data quality issues (37.8 per cent) and the skills gap (31.1 per cent). Read together with the literature, these barriers appear to operate as an interconnected system rather than as isolated obstacles: financial constraints limit the ability to hire technical talent, which forces reliance on external vendors, which in turn increases costs. Resistance from team members, by contrast, was notably low at 8.9 per cent, suggesting that cultural barriers to AI adoption in startups are considerably lower than those reported in the broader SME literature.

5.2 Answering the Central Research Question

The central research question of this study was: *How do EdTech startups implement AI agents to reduce operational costs and accelerate growth, and what are the principal challenges and outcomes associated with this process?*

On the basis of the evidence gathered, which draws on EdTech respondents alongside a wider cross-sector sample of startups, the answer is as follows. Startups, including those in EdTech, implement AI agents primarily through commercially available platforms, with OpenAI and Anthropic the most common choices. These platforms are deployed across customer-facing and administrative functions, where language-based automation yields the most immediate returns. Implementation is driven predominantly by internal operational need rather than investor pressure, and follows a pragmatic pattern in which AI is applied first to the highest-volume, most repetitive tasks before any broader roll-out.

The outcomes of this approach are broadly positive. AI agents are perceived to reduce costs, enable scaling without proportional increases in headcount and contribute meaningfully to business growth, as the literature anticipated. The principal challenges, however, are real: data privacy compliance, implementation costs, data quality issues and a persistent skills gap. These challenges appear to reinforce one another in practice, creating a barrier environment more complex than any single obstacle considered in isolation.

5.3 Recommendations

On the basis of these findings, the following recommendations are offered to EdTech startup managers who are considering AI agents or are already implementing them.

1. Prioritise deployment in content creation, customer support and marketing. These are the areas in which respondents reported the most consistent and significant cost savings and in which the technology is most mature. Starting with these functions maximises the likelihood of a rapid and visible return on investment, which matters particularly for early-stage ventures that need to demonstrate traction quickly.

2. Budget realistically for implementation complexity. One in five respondents found the implementation process more complex than expected, and over two in five cited high costs as a barrier. Managers should plan for the full cost of adoption, not merely the subscription fee, including integration time, workflow redesign and ongoing maintenance. AI is not a plug-and-play solution, and treating it as one tends to create problems later.

3. Take vendor dependency seriously and mitigate it where possible. With most startups in the sample relying on two or three commercial AI platforms, the risk of dependency is real rather than

theoretical. Managers should maintain flexibility across multiple providers and adopt modular architectures that allow providers to be switched without major operational disruption.

4. Invest in data privacy and compliance from the outset. Data privacy was the most frequently reported barrier, and the regulatory landscape continues to evolve. For EdTech startups handling learner data in particular, compliance is not optional; it must be built into the AI adoption strategy from the start rather than retrofitted after deployment.

5. Develop internal AI capabilities over time. The skills gap and vendor dependency are interconnected. Startups that invest in upskilling their existing teams, even modestly, will reduce their reliance on external vendors and build the internal knowledge needed to adapt as AI technology continues to advance.

5.4 Limitations and Future Research

The limitations of this study should be acknowledged. The sample of 45 respondents, while adequate for descriptive analysis, did not permit the inferential statistical testing that a larger dataset would support. The non-probability sampling approach, the geographic concentration of the sample in Ghana and Germany, and the fact that EdTech respondents formed a quarter of the sample all limit the extent to which the findings can be generalised. The cross-sectional design captured a snapshot rather than a trajectory over time, and the reliance on self-reported data introduces the possibility of response bias. Finally, the inability to conduct the interviews originally planned reduced the depth of insight that a mixed-methods approach could have achieved.

Future research could address these limitations in several ways. Larger-scale surveys, ideally with a larger EdTech-specific sample, would enable inferential analysis. Longitudinal studies could track AI adoption outcomes over time rather than at a single point. Qualitative research could capture in greater depth the decision-making processes and lived experiences of EdTech startup professionals. Comparative studies examining AI adoption across different startup maturity levels and regulatory environments would also make a valuable contribution to this developing field.

Notwithstanding these limitations, this study makes a meaningful contribution to an area in which primary empirical evidence remains scarce. Its findings offer both academic insight and practical guidance for the growing number of EdTech startup managers navigating the opportunities and challenges of AI agent adoption.

Bibliography

- Alam, A. and Mohanty, A. (2022) 'Business models, business strategies, and innovations in EdTech companies: integration of learning analytics and artificial intelligence in higher education', in *2022 IEEE 6th Conference on Information and Communication Technology (CICT)*. IEEE, pp. 1–6. doi: 10.1109/CICT56698.2022.9997887.
- Andersson, P. and Rosenqvist, C. (no date) *The evolution and impact of AI on EdTech business models* [Working paper]. Stockholm School of Economics.
- Barney, J. (1991) 'Firm resources and sustained competitive advantage', *Journal of Management*, 17(1), pp. 99–120. doi: 10.1177/014920639101700108.
- Brynjolfsson, E. and McAfee, A. (2017) 'The business of artificial intelligence', *Harvard Business Review*, July. Available at: <https://hbr.org/2017/07/the-business-of-artificial-intelligence>.
- Chinnaraju, A. (2025) 'AI-driven strategic decision-making on innovation: scalable, ethical approaches and AI agents for startups', *World Journal of Advanced Research and Reviews*, 25(2), pp. 2219–2248. doi: 10.30574/wjarr.2025.25.2.0575.
- Chopra, A. (2025) 'Adoption of AI and agentic systems: value, challenges, and pathways', *California Management Review Insight*, 15 August. Available at: <https://cmr.berkeley.edu>.
- Davis, F.D. (1989) 'Perceived usefulness, perceived ease of use, and user acceptance of information technology', *MIS Quarterly*, 13(3), pp. 319–340. doi: 10.2307/249008.
- Dwivedi, Y.K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P.V., Janssen, M., Jones, P., Kar, A.K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., Medaglia, R., Le Meunier-FitzHugh, K., Le Meunier-FitzHugh, L.C., Misra, S., Mogaji, E., Sharma, S.K., Singh, J.B., Raghavan, V., Raman, R., Rana, N.P., Samothrakis, S., Spencer, J., Tamilmani, K., Tubadji, A., Walton, P. and Williams, M.D. (2021) 'Artificial Intelligence (AI): multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy', *International Journal of Information Management*, 57, 101994. doi: 10.1016/j.ijinfomgt.2019.08.002.
- Farayola, O.A., Abdul, A.A., Irabor, B.O. and Okeleke, E.C. (2023) 'Innovative business models driven by AI technologies: a review', *Computer Science & IT Research Journal*, 4(2), pp. 85–110. doi: 10.51594/csitj.v4i2.608.
- Gartner (2024) *Gartner identifies the top 10 strategic technology trends for 2025*. Available at: <https://www.gartner.com/en/newsroom/press-releases/2024-10-21-gartner-identifies-the-top-10-strategic-technology-trends-for-2025>.
- Giuggioli, G. and Pellegrini, M.M. (2023) 'Artificial intelligence as an enabler for entrepreneurs: a systematic literature review and an agenda for future research', *International Journal of*

- Entrepreneurial Behavior & Research*, 29(4), pp. 816–837. doi: 10.1108/IJEBr-05-2021-0426.
- HoloniQ (2019) *Global EdTech market to reach \$404 billion by 2025*. Available at: <https://www.holoniq.com/notes/global-edtech-market-to-reach-404b-by-2025/>.
- Huang, L. (2023) 'Ethics of artificial intelligence in education: student privacy and data protection', *Science Insights Education Frontiers*, 16(2), pp. 2577–2587. doi: 10.15354/sief.23.re095.
- Lu, J.G., Zhao, G.G. and Zheng, A.M. (forthcoming) 'Generative AI use in entrepreneurship: an integrative review and an Empowerment–Entrapment Framework', *Strategic Entrepreneurship Journal* (Special Collection on AI and Entrepreneurship).
- Mohd Rasdi, R. and Umar Baki, N. (2025) 'Navigating the AI landscape in SMEs: overcoming internal challenges and external obstacles for effective integration', *PLOS ONE*. doi: 10.1371/journal.pone.0323249.
- Obiya, S.O. (2024) 'Impact of artificial intelligence (AI) on the growth of startups in Nigeria', *Journal of Economics and Allied Research*, 9(4), pp. 1–16.
- OECD (2020) 'The role of AI in education', in *OECD Digital Education Outlook 2020*. Paris: OECD Publishing.
- Ojika, F.U., Owobu, W.O., Abieba, O.A., Esan, O.J., Ubamadu, B.C. and Daraojimba, A.I. (2024) 'The role of artificial intelligence in business process automation: a model for reducing operational costs and enhancing efficiency', *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), pp. 1449–1462. doi: 10.62225/2583049X.2024.4.6.4046.
- Ries, E. (2011) *The lean startup: how today's entrepreneurs use continuous innovation to create radically successful businesses*. New York: Crown Business.
- Rogers, E.M. (2003) *Diffusion of innovations*. 5th edn. New York: Free Press.
- Saunders, M., Lewis, P. and Thornhill, A. (2019) *Research methods for business students*. 8th edn. Harlow: Pearson Education.
- Teece, D.J., Pisano, G. and Shuen, A. (1997) 'Dynamic capabilities and strategic management', *Strategic Management Journal*, 18(7), pp. 509–533. doi: 10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z.
- UNESCO (2023) *Artificial intelligence in education: challenges and opportunities for sustainable development*. Paris: UNESCO.
- World Economic Forum (2023) *The future of jobs report 2023*. Geneva: World Economic Forum. Available at: <https://www.weforum.org/publications/the-future-of-jobs-report-2023/>.

Appendices

Appendix A: Survey Questionnaire

Study title: How EdTech Startups Implement AI Agents to Reduce Operational Costs and Accelerate Growth

Instrument: Structured online questionnaire (Microsoft Forms)

Respondents: 45 startup professionals across 10 countries

Period: May–June 2026

Introduction displayed to respondents:

Thank you for taking the time to complete this questionnaire. This survey is part of an MSc dissertation at the Berlin School of Business and Innovation (BSBI), investigating how startups use AI agents to reduce operational costs and accelerate growth. Your responses are completely anonymous and will be used for academic purposes only. The survey takes approximately 8–10 minutes to complete. By proceeding, you consent to participate in this study.

Section 1: Demographic and Organisational Profile

Q1. What is your current role? *[Multiple choice]*

- ☐ Founder / Co-founder
- ☐ CTO / Technical Lead
- ☐ Operations Manager
- ☐ Product Manager
- ☐ Software Developer / Engineer
- ☐ Marketing / Growth Lead
- ☐ Other (please specify)

Q2. In which country is your startup primarily based? *[Short answer]*

Q3. What industry does your startup primarily operate in? *[Multiple choice]*

- ☐ Education Technology (EdTech)
- ☐ Financial Technology (FinTech)
- ☐ Health Technology (HealthTech)
- ☐ Software as a Service (SaaS)
- ☐ E-commerce
- ☐ Other (please specify)

Q4. What is the current stage of your startup? *[Multiple choice]*

- ☐ Pre-revenue / Idea stage
- ☐ Early-stage (1–10 employees)
- ☐ Growth-stage (11–50 employees)
- ☐ Scaling / Established (51+ employees)

Q5. How long has your startup been in operation? *[Multiple choice]*

- ☐ Less than 1 year
- ☐ 1–2 years
- ☐ 3–5 years
- ☐ More than 5 years

Section 2: AI Agent Adoption

Q6. Does your startup currently use AI agents or AI-powered automation tools in its operations?
[Multiple choice]

- ☐ Yes, actively using
- ☐ Currently implementing / piloting
- ☐ Planned but not yet started
- ☐ No

Q7. In which operational areas does your startup use AI agents? *[Select all that apply]*

- ☐ Customer support / chatbots
- ☐ Content generation / personalisation
- ☐ Marketing and lead generation
- ☐ Administrative automation (scheduling, reporting, enrolment)
- ☐ Data analytics and insights
- ☐ Software development / code generation
- ☐ Onboarding and user engagement
- ☐ Human resources / recruitment
- ☐ Other (please specify)

Q8. Which AI tools or platforms does your startup use? *[Select all that apply]*

- ☐ OpenAI (ChatGPT, GPT API)
- ☐ Anthropic (Claude)
- ☐ Google (Gemini, Vertex AI)
- ☐ Custom-built AI models
- ☐ Microsoft Copilot
- ☐ Zapier / Make (AI automation)
- ☐ Other (please specify)

Q9. What was the primary motivation for adopting AI agents? *[Rank your top 3]*

- ☐ Reducing operational costs
- ☐ Improving efficiency and speed
- ☐ Scaling without increasing headcount
- ☐ Staying competitive
- ☐ Improving customer / user experience
- ☐ Founder / leadership interest in AI
- ☐ Investor expectations

Section 3: Impact on Operations and Growth

Q10. To what extent has AI agent adoption reduced your startup's operational costs? *[Likert scale]*

- ☐ 1 = No reduction at all
- ☐ 2 = Slight reduction
- ☐ 3 = Moderate reduction
- ☐ 4 = Significant reduction
- ☐ 5 = Transformative reduction

Q11. In which areas have you experienced the most significant cost savings from AI agents? *[Select your top 3]*

- ☐ Customer support
- ☐ Content creation
- ☐ Marketing
- ☐ Administrative tasks
- ☐ Software development
- ☐ Data analysis
- ☐ Human resources
- ☐ Other (please specify)

Q12. To what extent has AI agent adoption contributed to the growth of your startup? *[Likert scale]*

- ☐ 1 = No contribution
- ☐ 2 = Slight contribution
- ☐ 3 = Moderate contribution
- ☐ 4 = Significant contribution
- ☐ 5 = Transformative contribution

Q13. Which of the following statements best reflects your experience? *[Select one]*

- ☐ AI agents have allowed our team to focus on higher-value strategic work.
- ☐ AI agents have improved the quality of our product or service.
- ☐ AI agents have enabled us to scale without proportionally increasing staff.
- ☐ AI agents have improved our response time to customers or users.
- ☐ AI agent implementation was more complex than we initially expected.
- ☐ The cost of implementing AI agents was higher than we initially anticipated.

Section 4: Barriers and Challenges

Q14. What challenges has your startup faced when implementing AI agents? *[Select all that apply]*

- ☐ Lack of technical expertise / skills gap
- ☐ High implementation costs
- ☐ Data quality or data availability issues
- ☐ Integration with existing systems
- ☐ Resistance from team members
- ☐ Data privacy and regulatory compliance (e.g. GDPR)
- ☐ Difficulty measuring ROI
- ☐ AI output quality / accuracy concerns
- ☐ Vendor dependency
- ☐ None – no significant challenges
- ☐ Other (please specify)

Q15. How significant were data privacy and regulatory concerns in your AI adoption process?
[Likert scale]

- ☐ 1 = Not significant at all
- ☐ 2 = Slightly significant
- ☐ 3 = Moderately significant
- ☐ 4 = Very significant
- ☐ 5 = Extremely significant / a major blocker

Section 5: Future Outlook

Q16. Do you plan to expand your use of AI agents in the next 12 months? *[Multiple choice]*

- ☐ Yes, significantly
- ☐ Yes, moderately
- ☐ No change planned
- ☐ Planning to reduce AI usage
- ☐ Unsure

Q17. What would most help your startup adopt or expand AI agents? *[Select your top 2]*

- ☐ More affordable AI tools
- ☐ Better technical talent availability
- ☐ Clearer regulatory guidance
- ☐ Proven case studies and best practices
- ☐ Internal training and upskilling
- ☐ Other (please specify)

Q18. Is there anything else you would like to share about your experience with AI agents in your startup? *[Open-ended, optional]*

Appendix B: Survey Data Summary

Total respondents: 45 · **Countries represented:** 10 · **Data collection period:** May–June 2026

Table B1: Respondent roles (n = 45)

Role	Number (n)	Percentage (%)
Founder / Co-founder	12	26.7
Software Developer / Engineer	8	17.8
Operations Manager	6	13.3
Product Manager	5	11.1
CTO / Technical Lead	5	11.1
Marketing / Growth Lead	3	6.7
Intern	3	6.7
HR Professional	2	4.4
Finance	1	2.2

Table B2: Geographic distribution (n = 45)

Country	Number (n)	Percentage (%)
Ghana	17	37.8
Germany	13	28.9
Kenya	4	8.9
Spain	2	4.4
Nigeria	2	4.4
China	2	4.4
United States	2	4.4
Greece	1	2.2
France	1	2.2
Sweden	1	2.2

Table B3: Industry distribution (n = 45)

Industry	Number (n)	Percentage (%)
Software as a Service (SaaS)	14	31.1
Education Technology (EdTech)	11	24.4
Financial Technology (FinTech)	7	15.6
Health Technology (HealthTech)	5	11.1
E-commerce	4	8.9
Other (logistics, interior design, telecoms)	4	8.9

Table B4: AI agent adoption status (n = 45)

Status	Number (n)	Percentage (%)
Yes, actively using	34	75.6
Currently implementing / piloting	7	15.6
Planned but not yet started	3	6.7
No	1	2.2

Table B5: Operational areas of AI agent deployment (n = 45, multi-select)

Operational area	Number (n)	Percentage (%)
Customer support / chatbots	28	62.2
Administrative automation	25	55.6
Content generation / personalisation	24	53.3
Software development / code generation	21	46.7
Marketing and lead generation	18	40.0
Data analytics and insights	15	33.3
Human resources / recruitment	11	24.4
Onboarding and user engagement	10	22.2

Table B6: AI tools and platforms used (n = 45, multi-select)

Platform	Number (n)	Percentage (%)
OpenAI (ChatGPT, GPT API)	32	71.1
Anthropic (Claude)	29	64.4
Google (Gemini, Vertex AI)	21	46.7
Microsoft Copilot	10	22.2
Custom-built AI models	7	15.6
n8n	7	15.6
Zapier / Make	5	11.1

Table B7: Impact of AI agents on operational cost reduction (n = 45)

Response	Number (n)	Percentage (%)
5 = Transformative reduction	14	31.1
4 = Significant reduction	14	31.1
3 = Moderate reduction	13	28.9
2 = Slight reduction	3	6.7
1 = No reduction at all	1	2.2
Mean score		3.82

Table B8: Contribution of AI agents to startup growth (n = 45)

Response	Number (n)	Percentage (%)
5 = Transformative contribution	9	20.0
4 = Significant contribution	22	48.9
3 = Moderate contribution	11	24.4
2 = Slight contribution	3	6.7
1 = No contribution	0	0.0
Mean score		3.82

Table B9: Barriers to AI agent implementation (n = 45, multi-select)

Barrier	Number (n)	Percentage (%)
Data privacy / regulatory compliance	21	46.7
High implementation costs	19	42.2
Data quality / availability issues	17	37.8
Skills gap / lack of expertise	14	31.1
Integration with existing systems	12	26.7
Vendor dependency	11	24.4
AI output quality / accuracy	7	15.6
Resistance from team members	4	8.9
No significant challenges	4	8.9
Difficulty measuring ROI	2	4.4

Table B10: Future AI expansion plans (n = 45)

Response	Number (n)	Percentage (%)
Yes, significantly	25	55.6
Yes, moderately	18	40.0
No change planned	1	2.2
Planning to reduce AI usage	1	2.2
Unsure	0	0.0