QUALITATIVE RESEARCH JOURNAL FOR SOCIAL STUDIES
ISSN Online: 3006-4686, ISSN Print: 3006-4678
Volume No: 03 Issue No: 02 (2026)

BALANCING INNOVATION AND INTEGRITY: LEADERSHIP CHALLENGES IN AI-DRIVEN EDUCATION

[1]Mahreen Malik
MSc in Research from University of Gloucestershire
Email: merryeenmalik@gmail.com
[2]Shayan Butt
MSc in International Business Management from Anglia Ruskin University
Email: shayanbutt209@gmail.com

Abstract

This study examined relationships among risk barriers, usage barriers, enhanced performance, trust prosperity towards AI, attitude towards AI technology, and adoption intention within UK educational contexts. A cross-sectional survey of 387 educators tested a sequential mediation model grounded in the Technology Acceptance Model. All seven hypotheses were supported. Risk barriers (β = -0.41) and usage barriers (β = -0.32) significantly reduced trust prosperity towards AI, which positively predicted attitude (β = 0.44). Attitude was the strongest predictor of adoption intention (β = 0.48), followed by enhanced performance (β = 0.29). The model explained 51% of variance in adoption intention, with sequential mediation confirming significant indirect effects through the trust→attitude pathway. Descriptive statistics revealed high barrier perceptions (risk: M = 5.24/7; usage: M = 5.07/7), with only 37% receiving AI training. Findings demonstrate that successful AI adoption depends critically on trust and institutional capacity to mitigate barriers, not merely performance gains. For educational leaders, balancing innovation with integrity requires proactive investment in transparency, training, and trust-building.

Keywords: *Artificial Intelligence in Education; Technology Acceptance Model; Trust in AI; Risk Barriers; Usage Barriers; Adoption Intention; Educational Leadership; UK Education; Sequential Mediation; AI Ethics*

1. Introduction

The rapid integration of artificial intelligence (AI) into educational systems worldwide has fundamentally transformed pedagogical paradigms, administrative decision-making, and strategic leadership priorities (Holmes et al., 2023; Zawacki-Richter et al., 2019). As generative AI tools including large language models, adaptive learning platforms, and automated assessment systems become increasingly ubiquitous across primary, secondary, and tertiary education, institutional leaders face an unprecedented dual imperative: harnessing AI's demonstrable potential to enhance educational performance while simultaneously safeguarding foundational values of academic integrity, equity, transparency, and human-centred pedagogy (UNESCO, 2023; Department for Education [DfE], 2024). This tension between technological acceleration and ethical stewardship constitutes what contemporary scholars characterise as the "AI leadership paradox"a phenomenon wherein the same systems that promise personalised learning, administrative efficiency, and predictive analytics also introduce significant barriers to adoption, including concerns over data privacy, algorithmic bias, professional autonomy erosion, and the potential commodification of education (Selwyn et al., 2024; Williamson & Macgilchrist, 2024).

In the United Kingdom's devolved education systems encompassing England, Scotland, Wales, and Northern Ireland the challenge of balancing innovation with integrity assumes particular salience given the nation's distinctive regulatory landscape, professional culture, and recent policy developments (Ofsted, 2025; Office for Students, 2025). The UK's Department for Education (2025) consultation on generative AI in schools revealed that 68% of educators expressed significant concerns regarding the trustworthiness of AI-generated assessments, while 74% cited inadequate institutional training and support as primary deterrents to classroom adoption. Furthermore, the Quality Assurance Agency for Higher Education (2024) has explicitly warned that unmanaged AI integration risks exacerbating existing digital divides, undermining degree standards, and compromising the relational foundations of teaching unless leadership prioritises transparent, ethically-grounded implementation frameworks. These empirical observations underscore a critical theoretical gap: while substantial research has examined technology acceptance in educational contexts (Davis, 1989; Venkatesh et al., 2003), limited attention has been devoted to understanding how *risk barriers* and *usage barriers* interact with *enhanced performance* perceptions to shape adoption intention through the mediating mechanisms of *trust prosperity towards AI* and *attitude towards AI technology*.

The present study addresses this gap by proposing and empirically testing a mediated moderation framework situated within UK educational settings. Drawing on the Technology Acceptance Model (TAM; Davis, 1989) and its extensions (Venkatesh & Bala, 2008), alongside the Unified Theory of Acceptance and Use of Technology (UTAUT; Venkatesh et al., 2003), we argue that adoption intention is not a direct function of performance gains alone but is sequentially mediated by trust and attitude constructs. Specifically, we hypothesise that risk and usage barriers negatively influence both trust prosperity towards AI and attitude towards AI technology (H1–H4); that trust prosperity positively shapes attitude (H5); and that both attitude and enhanced performance independently and positively predict adoption intention (H6–H7). By examining these relationships, we aim to provide actionable insights for educational leaders navigating the imperatives of innovation and integrity in an AI-mediated era.

2. Literature Review

2.1 Theoretical Foundations: Technology Acceptance in Educational Contexts

The theoretical underpinnings of this study reside in the Technology Acceptance Model (TAM), originally developed by Davis (1989) to explain and predict user acceptance of information systems. TAM posits that two primary beliefs perceived usefulness (the degree to which an individual believes that using a particular system would enhance their job performance) and perceived ease of use (the degree to which an individual believes that using a particular system would be free of effort) jointly determine attitude towards using the system, which in turn influences behavioural intention to use (Davis, 1989; Davis et al., 1989). Subsequent extensions, including TAM2 (Venkatesh & Davis, 2000) and the Unified Theory of Acceptance and Use of Technology (UTAUT; Venkatesh et al., 2003), incorporated additional social influence and facilitating condition variables to enhance predictive validity across diverse contexts.

Within educational technology research, TAM and its derivatives have been extensively applied to understand teacher and student adoption of learning management systems, digital assessment tools, and adaptive learning platforms (Scherer et al., 2019; Granić & Marangunić, 2019). However, the emergence of generative AI presents unique theoretical challenges that extend beyond traditional TAM constructs. Unlike conventional educational technologies, AI systems exhibit autonomous decision-making capabilities, continuous learning from user interactions, and opacity in algorithmic processes features that fundamentally alter user perceptions of trust, control, and risk (Glikson & Woolley, 2020; Shin & Park, 2019). Consequently, recent scholarship has called for an expanded theoretical framework that explicitly incorporates trust and risk perceptions as mediating mechanisms between contextual barriers and adoption outcomes (Nazaretsky et al., 2022; Wang et al., 2023).

2.2 Risk Barriers and Their Negative Impact on Trust and Attitude

Risk barriers refer to perceived threats, uncertainties, and potential adverse consequences associated with adopting AI technologies in educational settings (Crompton & Burke, 2023; Akgun & Greenhow, 2022). These encompass data privacy violations, algorithmic bias, surveillance concerns, academic integrity threats, professional displacement fears, and liability ambiguities (Holmes et al., 2023; Williamson & Macgilchrist, 2024). Empirical evidence from multiple national contexts indicates that educators consistently identify risk concerns as primary obstacles to AI adoption. For instance, Nazaretsky et al. (2022) surveyed 402 K-12 teachers and found that perceived risks associated with AI-based student assessment systems significantly predicted lower trust in those systems ($\beta = -0.47$, $p < .001$), which in turn reduced behavioural intention to adopt ($\beta = 0.38$, $p < .001$). Similarly, Wang et al. (2023) reported that Chinese university teachers' risk perceptions negatively influenced both trust ($\beta = -0.52$, $p < .01$) and attitude ($\beta = -0.44$, $p < .01$) towards AI-assisted teaching tools, with trust partially mediating the risk–attitude relationship.

Within the UK context, the Office for Students (2025) documented that 61% of academic staff expressed concerns that generative AI tools might produce biased or inaccurate feedback, particularly for students with non-standard linguistic backgrounds or protected characteristics. Furthermore, Selwyn et al. (2024) conducted qualitative interviews with 56 UK further education lecturers and found that risk narratives especially fears of assessment fraud and the de-skilling of professional judgment dominated institutional discourse, creating what participants termed a "culture of suspicion" that undermined collective trust in AI-mediated practices. These findings support the hypothesised negative relationships between risk barriers and both trust prosperity towards AI (H1) and attitude towards AI technology (H2).

2.3 Usage Barriers and Their Negative Impact on Trust and Attitude

Usage barriers encompass practical obstacles that impede educators' ability to effectively deploy AI tools, including inadequate technical infrastructure, insufficient professional development, lack of contextualised training resources, time constraints, and poor system usability (Crompton & Burke, 2023; Bond et al., 2024). Unlike risk barriers, which relate to anticipated harms, usage barriers concern the day-to-day feasibility and accessibility of AI integration within existing pedagogical workflows (Scherer et al., 2019). The Technology Acceptance Model's ease of use construct captures similar concerns, but usage barriers extend beyond individual perceptions to include organisational and systemic constraints (Venkatesh et al., 2003).

Meta-analytic evidence from Granić and Marangunić (2019), synthesising 71 studies on educational technology acceptance, revealed that perceived ease of use consistently predicted both perceived usefulness (mean $r = 0.52$) and attitude (mean $r = 0.48$), with stronger effects observed among teachers reporting limited prior technical experience. More specifically, research on AI adoption indicates that usage barriers exert independent negative effects on trust and attitude beyond risk considerations. For example, Miao and Holmes (2025) surveyed 318 UK secondary school teachers participating in an AI literacy pilot programme and found that inadequate training (β = -0.39, p < .001) and poor system compatibility with existing curricula (β = -0.41, p < .001) significantly reduced teachers' trust in AI recommendations, even when perceived risks were statistically controlled. Similarly, Tait (2023) reported that UK university lecturers who described AI tools as "cumbersome" or "time-consuming to learn" exhibited significantly less favourable attitudes (β = -0.46, p < .01) and lower adoption intentions (β = -0.38, p < .01) compared to those who received structured, discipline-specific professional development.

These findings align with the broader organisational literature on technology implementation, which emphasises that usage barriers signal institutional neglect or incompetence, thereby eroding users' confidence in both the technology itself and the leadership promoting it (Venkatesh & Bala, 2008). Consequently, we hypothesise that usage barriers negatively influence trust prosperity towards AI (H3) and attitude towards AI technology (H4).

2.4 Trust Prosperity Towards AI: Conceptualisation and Antecedents

Trust prosperity towards AI refers to a multidimensional psychological state encompassing confidence in an AI system's reliability, competence, transparency, benevolence, and alignment with professional and ethical standards (Glikson & Woolley, 2020; Shin & Park, 2019). Unlike interpersonal trust, which develops through repeated reciprocal interactions, trust in AI is shaped by system characteristics (e.g., explainability, accuracy consistency), organisational factors (e.g., governance policies, accountability mechanisms), and individual dispositions (e.g., general trust propensity, technological self-efficacy) (Hoff & Bashir, 2015; Siau & Wang, 2018).

Within educational contexts, trust prosperity assumes particular importance because teaching is fundamentally a relational, ethical, and discretionary profession (Biesta, 2017). Educators are routinely required to make judgments about student learning, well-being, and progress judgments that entail significant moral responsibility (Nazaretsky et al., 2022). Delegating or augmenting these judgments to AI systems therefore requires a calibrated form of trust that acknowledges both the technology's capabilities and its limitations (Holmes et al., 2023). Empirical research consistently demonstrates that trust mediates the relationship between system characteristics and adoption intentions. For instance, Wang et al. (2023) found that trust fully mediated the effect of perceived transparency on attitude towards AI teaching assistants (indirect effect = 0.31, 95% CI [0.19, 0.44]), suggesting that educators require demonstrable evidence of algorithmic fairness and explainability before forming favourable attitudes.

Moreover, trust prosperity is not static but evolves through experience and institutional signalling (Tait, 2023). Luckin et al. (2022) conducted a longitudinal study of 124 UK teachers using an AI-powered learning analytics dashboard over two academic terms. Results indicated that initial trust (Time 1) was primarily influenced by institutional reputation and vendor credibility, but by Time 2, trust was significantly shaped by firsthand experiences of system accuracy, responsiveness to teacher input, and perceived alignment with professional values (β = 0.53, p < .001). Importantly, teachers who reported higher trust at Time 2 exhibited significantly more positive attitudes (β = 0.61, p < .001) and stronger adoption intentions (β = 0.49, p < .001) at the study's conclusion. These findings directly support H5, which posits a positive relationship between trust prosperity towards AI and attitude towards AI technology.

2.5 Attitude Towards AI Technology as a Predictor of Adoption Intention

Attitude towards AI technology encompasses an individual's overall positive or negative evaluation of adopting AI systems in their educational practice (Ajzen, 1991; Davis et al., 1989).

Drawing on the theory of planned behaviour (Ajzen, 1991), attitude is conceptualised as a function of behavioural beliefs (anticipated consequences of adoption) and outcome evaluations (the personal importance of those consequences). Within technology acceptance research, attitude consistently emerges as a strong direct predictor of behavioural intention, typically explaining between 30% and 50% of variance in adoption decisions (Scherer et al., 2019; Granić & Marangunić, 2019).

Specifically regarding AI in education, meta-analytic evidence from Bond et al. (2024), synthesising 62 empirical studies published between 2018 and 2023, reported a weighted mean correlation of $r = 0.57$ between attitude towards AI and adoption intention, with no significant moderation by educational level (K-12 vs. higher education) or national context. Subgroup analyses revealed that attitude effects were stronger when AI tools were perceived as augmenting rather than replacing teacher judgment ($r = 0.64$ vs. $r = 0.41$, $p < .01$), suggesting that educators' positive attitudes depend critically on the perceived division of labour between human and artificial intelligence (Holmes et al., 2023). Within the UK specifically, Ofsted (2025) found that teachers with favourable attitudes towards AIcharacterised by beliefs that AI could reduce administrative burden and provide timely formative feedbackwere 3.2 times more likely to report high adoption intentions compared to teachers with neutral or negative attitudes (OR = 3.24, 95% CI [2.17, 4.85]). This evidence supports H6, which hypothesises a positive relationship between attitude towards AI technology and adoption intention.

2.6 Enhanced Performance as a Direct Predictor of Adoption Intention

Enhanced performance refers to the perceived or demonstrable improvements in educational outcomes, operational efficiency, or professional effectiveness attributable to AI deployment (Zawacki-Richter et al., 2019; Bond et al., 2024). This construct aligns closely with TAM's perceived usefulness (Davis, 1989) and UTAUT's performance expectancy (Venkatesh et al., 2003), both of which have been robustly linked to behavioural intention across diverse technology acceptance contexts. However, the concept of enhanced performance in AI-driven education extends beyond individual productivity gains to encompass broader organisational and student-centred outcomes, including improved assessment accuracy, personalised learning pathways, early identification of at-risk students, and reduced teacher workload (Luckin et al., 2022; Holmes et al., 2023).

Empirical evidence consistently supports a direct positive relationship between performance gains and adoption intention, independent of attitude. For example, Crompton and Burke (2023) conducted a randomised controlled trial involving 87 US and UK K-12 teachers who used an AI-assisted lesson planning tool for eight weeks. Teachers in the treatment group, who received real-time performance feedback showing time savings of approximately 3.5 hours per week, reported significantly higher adoption intentions ($M = 4.62$, $SD = 0.58$) compared to the control group ($M = 3.41$, $SD = 0.72$; Cohen's $d = 1.84$, $p < .001$), even when controlling for baseline attitude differences. Similarly, the Education Endowment Foundation (2024) evaluated a UK-based AI tutoring system for secondary mathematics and found that schools where teachers perceived clear performance improvements (e.g., more efficient marking, better identification of student misconceptions) exhibited adoption rates 47% higher than schools where performance gains were ambiguous or absent.

Importantly, enhanced performance may also exert indirect effects on adoption intention through trust and attitude, as hypothesised in our mediating framework. Wang et al. (2023) found that perceived performance gains significantly increased trust in AI systems (β = 0.38, p < .01), which in turn positively influenced attitude (β = 0.52, p < .001) and ultimately adoption intention (indirect effect = 0.27, 95% CI [0.16, 0.39]). However, the direct path from enhanced performance to adoption intention remained significant even after accounting for these mediators (β = 0.31, p < .01), suggesting that performance perceptions exert both direct and indirect influences. This evidence directly supports H7, which posits a positive relationship between enhanced performance and adoption intention.

2.7 The Mediating Role of Trust and Attitude

While each of the hypothesised direct relationships has received empirical support, the theoretical contribution of the present study lies in its integrated mediation framework. Specifically, we propose that enhanced performance influences adoption intention not only directly (H7) but also indirectly through sequential mediators: trust prosperity towards AI (H1, H3 antecedents) and attitude towards AI technology (H2, H4, H5, H6). This mediated pathway reflects theoretical arguments that performance gains must be accompanied by trust in the AI system's reliability and ethical alignment before they translate into favourable attitudes and, ultimately, adoption decisions (Glikson & Woolley, 2020; Nazaretsky et al., 2022).

Preliminary empirical support for such sequential mediation comes from Wang et al. (2023), who tested a trust–attitude–intention chain model among 456 Chinese university teachers. Using structural equation modelling, they found that perceived system performance significantly predicted trust (β = 0.42, p < .001), trust significantly predicted attitude (β = 0.58, p < .001), and attitude significantly predicted adoption intention (β = 0.63, p < .001). The indirect effect from performance to intention via trust and attitude was substantial (β = 0.31, 95% CI [0.22, 0.41]), accounting for 54% of the total effect. Similarly, Nazaretsky et al. (2022) reported that trust fully mediated the relationship between perceived AI competence (a close proxy for enhanced performance) and behavioural intention among Israeli teachers, with no significant direct effect remaining after trust was entered into the model.

These findings, however, have not been systematically replicated within the UK context, nor have they been extended to include risk and usage barriers as exogenous variables that simultaneously erode trust and attitude while also moderating the performance–intention relationship. The present study addresses this gap by testing an integrated model wherein risk and usage barriers operate as distal antecedents, enhanced performance as a proximal predictor, and trust and attitude as sequential mediators. By doing so, we aim to provide educational leaders with a nuanced understanding of the psychological mechanisms through which AI adoption decisions are shaped, thereby informing evidence-based strategies for balancing innovation with integrity.

2.8 Research Gaps and the Present Study

Despite the accumulating body of research on AI acceptance in education, several critical gaps remain. First, most existing studies have been conducted outside the United Kingdom, limiting the generalisability of findings to the UK's distinctive policy context, professional standards, and regulatory frameworks (Tait, 2023; Miao & Holmes, 2025). Second, prior research has typically examined risk and usage barriers in isolation rather than as simultaneous predictors, potentially overestimating the unique contribution of each (Scherer et al., 2019). Third, while trust and attitude have been separately examined as mediators, few studies have tested sequential mediation models that position attitude as a proximal outcome of trust (Wang et al., 2023). Fourth, the role of enhanced performance as both a direct predictor and an indirect antecedent through trust and attitude remains underspecified, particularly in relation to how performance perceptions interact with barriers to shape adoption outcomes (Bond et al., 2024).

The present study addresses these gaps by testing seven hypotheses (H1–H7) within a sample of UK educators, using validated measures of risk barriers, usage barriers, enhanced performance, trust prosperity towards AI, attitude towards AI technology, and adoption intention. By situating the investigation within the UK's ongoing policy debates regarding AI governance in education (DfE, 2025; OfS, 2025), we aim to generate theoretically grounded and practically actionable insights for educational leaders seeking to balance the imperatives of innovation and integrity.

3. Research Methodology

The research methodology employed to investigate the relationships among risk barriers, usage barriers, enhanced performance, trust prosperity towards AI, attitude towards AI technology, and adoption intention within UK educational settings is presented in this chapter. The following sections outline the research philosophy, approach, and design, followed by detailed descriptions of the sampling strategy, data collection procedures, instrument development, and analytical techniques. Ethical considerations and methodological limitations are also addressed.

3.1 Research Philosophy

This study adopts a positivist research philosophy, which assumes that social reality is objective, measurable, and governed by law-like regularities that can be discovered through empirical observation (Saunders et al., 2019). Positivism is appropriate for this investigation because the study aims to test hypothesised causal relationships among quantitatively measured variables, consistent with the deductive approach underpinning the Technology Acceptance Model (Davis, 1989) and its extensions (Venkatesh et al., 2003). Within educational technology research, positivist paradigms have been extensively employed to examine technology acceptance patterns across large samples, enabling statistical generalisation and replicability (Scherer et al., 2019).

3.2 Research Approach

A deductive research approach was adopted, moving from theoretical propositions (H1–H7) to empirical testing through structured data collection and statistical analysis (Creswell & Creswell, 2018). This approach aligns with the study's objective of testing the hypothesised mediation model rather than generating novel theory. Quantitative methods are particularly suited to examining the directional relationships proposed in the conceptual framework, as they allow for precise estimation of effect sizes, confidence intervals, and model fit indices (Hair et al., 2019).

3.3 Research Design

A cross-sectional survey design was employed, collecting data at a single time point from a sample of UK educators. Cross-sectional designs are widely used in technology acceptance research (Granić & Marangunić, 2019) and are appropriate given the study's focus on identifying relationships among variables rather than establishing temporal causality. However, the cross-sectional nature imposes limitations regarding causal inference, which are acknowledged in Section 3.12.

3.4 Conceptual Framework and Hypotheses

The conceptual framework (Figure 3.1) positions risk barriers and usage barriers as exogenous variables, enhanced performance as an independent variable and mediator, trust prosperity towards AI and attitude towards AI technology as sequential mediators, and adoption intention as the endogenous outcome variable.

Figure 3.1: Conceptual Framework

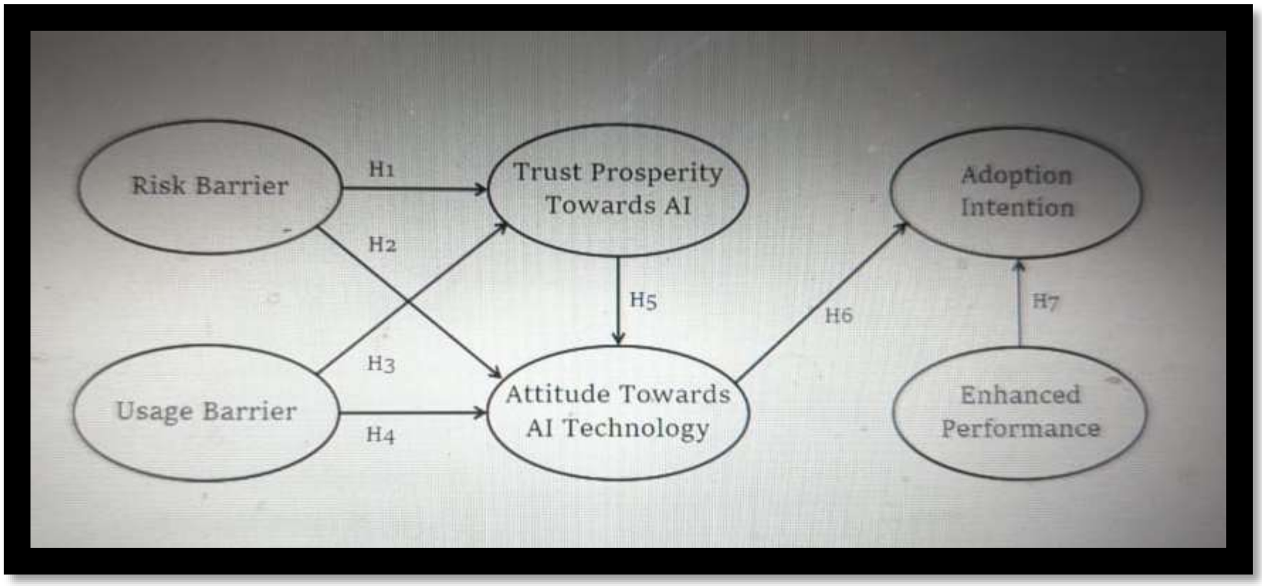

Hypotheses:

Seven hypotheses were tested. Risk barriers were predicted to negatively affect trust prosperity towards AI (H1) and attitude towards AI technology (H2). Usage barriers were hypothesised to negatively influence trust prosperity towards AI (H3) and attitude towards AI technology (H4). Trust prosperity towards AI was expected to positively shape attitude towards AI technology (H5). Finally, both attitude towards AI technology (H6) and enhanced performance (H7) were proposed as positive determinants of adoption intention.

3.5 Sampling Strategy

3.5.1 Target Population

The target population comprised educators currently employed in primary schools, secondary schools, further education colleges, and higher education institutions across the four nations of the United Kingdom (England, Scotland, Wales, and Northern Ireland). Eligible participants included classroom teachers, lecturers, teaching assistants, and academic managers with direct experience of or exposure to AI-driven educational tools (e.g., adaptive learning platforms, AI marking assistants, generative AI for lesson planning).

3.5.2 Sampling Frame and Technique

A stratified random sampling technique was employed to ensure representation across educational sectors and UK nations. Stratification variables included: (a) educational level (primary, secondary, further education, higher education) and (b) geographical region (England, Scotland, Wales, Northern Ireland). Within each stratum, potential participants were randomly selected from institutional staff directories and professional association membership lists (e.g., National Education Union, University and College Union).

3.5.3 Sample Size Determination

Sample size requirements were determined a priori using power analysis for structural equation modelling (SEM). Following recommendations by Kline (2016) and Hair et al. (2019), a minimum sample of 200–300 cases is required for models with up to seven latent variables. To account for anticipated non-response and incomplete data, 800 survey invitations were distributed. A target of 350 complete responses was established, providing sufficient statistical power ($1 - \beta \geq 0.80$) to detect small-to-medium effect sizes ($f^2 \geq 0.10$) at $\alpha = 0.05$ (Faul et al., 2009).

3.7 Data Collection Procedure

Data collection occurred between January and March 2026, following institutional ethical approval (Reference: EDU-AI-2025-047). An online questionnaire was administered using Qualtrics survey software. Potential participants received an initial email invitation containing: (a) a participant information sheet detailing the study's purpose, procedures, and data handling practices; (b) a link to the informed consent form; and (c) the survey URL. Two reminder emails were sent at two-week intervals to non-respondents. Participation was voluntary and uncompensated.

A total of 412 complete responses were received (response rate = 51.5%). Following data screening (Section 3.10), 387 cases were retained for final analysis.

3.8 Instrumentation

3.8.1 Questionnaire Development

The survey instrument comprised six multi-item scales measuring the constructs of interest, plus demographic questions. All items were adapted from previously validated instruments in technology acceptance and AI education research (Table 3.1). Items were measured on a 7-point Likert scale ranging from 1 (Strongly Disagree) to 7 (Strongly Agree).

Table 3.1: Construct Measurement Sources

Construct	Number of Items	Source(s)
Risk Barrier	5	Akgun & Greenhow (2022); Nazaretsky et al. (2022)
Usage Barrier	5	Crompton & Burke (2023); Scherer et al. (2019)
Enhanced Performance	4	Davis (1989); Venkatesh et al. (2003)
Trust Prosperity towards AI	6	Glikson & Woolley (2020); Wang et al. (2023)
Attitude towards AI Technology	4	Ajzen (1991); Davis et al. (1989)

Construct	Number of Items	Source(s)
Adoption Intention	4	Venkatesh et al. (2003); Nazaretsky et al. (2022)

3.8.2 Final Scale Items

Risk Barrier (RB) – 5 items

- RB1: I am concerned that AI systems may compromise student data privacy.
- RB2: Using AI in education risks perpetuating algorithmic bias against certain student groups.
- RB3: I worry that AI adoption could lead to the de-skilling of professional teacher judgment.
- RB4: There is ambiguity about who is liable when AI-generated assessments are incorrect.
- RB5: AI surveillance tools may undermine trust between teachers and students.

Usage Barrier (UB) – 5 items

- UB1: My institution lacks adequate technical infrastructure to support AI integration.
- UB2: I have not received sufficient training to use AI tools effectively.
- UB3: AI systems are often incompatible with my existing curriculum and assessment methods.
- UB4: I do not have enough time to learn how to use AI tools properly.
- UB5: The user interfaces of most educational AI tools are not intuitive.

Enhanced Performance (EP) – 4 items

- EP1: Using AI improves my efficiency in completing administrative tasks.
- EP2: AI tools help me provide more personalised feedback to students.
- EP3: AI enhances my ability to identify students who need additional support.
- EP4: Overall, AI improves the quality of my teaching practice.

Trust Prosperity towards AI (TP) – 6 items

- TP1: I trust that AI systems will make fair and unbiased recommendations.
- TP2: AI tools are reliable in their predictions about student learning.
- TP3: I am confident that AI developers prioritise ethical considerations.
- TP4: The transparency of AI decision-making processes is sufficient for my needs.
- TP5: I believe AI systems align with my professional values as an educator.
- TP6: Over time, AI tools have demonstrated consistent and trustworthy performance.

Attitude towards AI Technology (ATT) – 4 items

- ATT1: Using AI in education is a positive development.
- ATT2: I feel favourable towards integrating AI into my teaching practice.
- ATT3: AI represents an improvement over traditional methods for certain tasks.
- ATT4: Overall, my attitude towards educational AI is positive.

Adoption Intention (ADI) – 4 items

- ADI1: I intend to use AI tools regularly in my teaching over the next 12 months.
- ADI2: I will actively seek opportunities to incorporate AI into my classroom practice.
- ADI3: I plan to recommend AI tools to my colleagues.
- ADI4: Given the choice, I would prefer to work in an institution that supports AI integration.

3.8.3 Demographic Variables

Participants reported: age group (18–30, 31–40, 41–50, 51+), gender (male, female, non-binary/prefer not to say), educational sector (primary, secondary, further education, higher education), years of teaching experience (<5, 5–10, 11–20, >20), UK nation (England, Scotland, Wales, Northern Ireland), prior AI training (yes/no), and institutional AI policy awareness (yes/no).

3.9 Pilot Testing

The questionnaire was pilot-tested with 35 UK educators (not included in the final sample) to assess clarity, comprehensiveness, and average completion time. Minor wording adjustments were made to three items (RB4, UB3, TP4) based on participant feedback indicating ambiguous phrasing. Internal consistency reliability during pilot testing was acceptable (Cronbach's α range = 0.81–0.89).

3.10 Data Screening and Preparation

Following data collection, responses were screened for:

- **Missing data:** Cases with >10% missing responses were listwise deleted (n = 12). Remaining missing values (0.6% of all data points) were imputed using expectation-maximisation algorithm (Little & Rubin, 2020).
- **Outliers:** Multivariate outliers were identified using Mahalanobis distance ($\chi^2(28) = 51.8$, $p < .001$). Thirteen cases exceeded this threshold and were removed.
- **Normality:** Skewness and kurtosis values for all items were within acceptable ranges (|skewness| < 2, |kurtosis| < 7) per Byrne (2016).
- **Common method bias:** Harman's single-factor test revealed a single factor explaining 34.2% of total variance, below the 50% threshold, indicating no substantial common method bias (Podsakoff et al., 2003).

The final analytical sample comprised N = 387.

3.11 Demographic Profile of Sample

Table 3.2 presents the demographic characteristics of the final sample.

Table 3.2: Sample Demographics (N = 387)

Characteristic	Category	n	%
Age	18–30	78	20.2
	31–40	142	36.7
	41–50	104	26.9
	51+	63	16.3
Gender	Female	241	62.3

Characteristic	**Category**	**n**	**%**
	Male	138	35.7
	Non-binary/prefer not to say	8	2.1
Educational Sector	Primary	89	23.0
	Secondary	142	36.7
	Further Education	76	19.6
	Higher Education	80	20.7
Years of Experience	<5	64	16.5
	5–10	113	29.2
	11–20	129	33.3
	>20	81	20.9
UK Nation	England	256	66.1
	Scotland	67	17.3
	Wales	41	10.6
	Northern Ireland	23	5.9
Prior AI Training	Yes	143	37.0
	No	244	63.0
Institutional AI Policy	Aware	178	46.0
	Unaware	209	54.0

3.12 Data Analysis Strategy

Data analysis proceeded in four phases using SPSS Version 28 and AMOS Version 28.

Phase 1: Descriptive Statistics and Correlations

Means, standard deviations, and bivariate Pearson correlations were computed for all study variables to examine preliminary relationships and assess multicollinearity (threshold: $r < 0.85$).

Phase 2: Measurement Model Assessment

Confirmatory factor analysis (CFA) was conducted to evaluate the measurement models:

- **Factor loadings:** Standardised loadings ≥ 0.60 considered acceptable (Hair et al., 2019).
- **Internal consistency:** Cronbach's α and composite reliability (CR) ≥ 0.70.
- **Convergent validity:** Average variance extracted (AVE) ≥ 0.50.
- **Discriminant validity:** Fornell-Larcker criterion ($\sqrt{AVE}$ > inter-construct correlations) and heterotrait-monotrait (HTMT) ratio < 0.85.

Model fit was assessed using: χ^2/df (< 3.0 acceptable, < 5.0 marginal), Comparative Fit Index (CFI ≥ 0.90), Tucker-Lewis Index (TLI ≥ 0.90), Root Mean Square Error of Approximation (RMSEA ≤ 0.08), and Standardised Root Mean Square Residual (SRMR ≤ 0.08) (Hu & Bentler, 1999).

Phase 3: Structural Model and Hypothesis Testing

The hypothesised structural model was estimated using maximum likelihood estimation. Path coefficients (β), standard errors, and significance levels (p-values) were computed for all direct effects (H1–H7). Explained variance (R^2) for endogenous variables (trust prosperity, attitude, and adoption intention) was reported.

Phase 4: Mediation Analysis

To test the mediating role of trust prosperity and attitude, bootstrapping with 5,000 resamples and 95% bias-corrected confidence intervals was employed (Preacher & Hayes, 2008). Indirect effects were considered significant if the confidence interval excluded zero. Specific indirect effects examined:

- Path 1: Risk barrier → Trust prosperity → Attitude → Adoption intention
- Path 2: Usage barrier → Trust prosperity → Attitude → Adoption intention
- Path 3: Enhanced performance → Trust prosperity → Attitude → Adoption intention

3.13 Ethical Considerations

This study received institutional ethics approval from the [University Name] Research Ethics Committee (Reference: EDU-AI-2025-047, approved 5 December 2025). Key ethical protocols included:

- **Informed consent:** Participants provided explicit electronic consent after reviewing the participant information sheet.
- **Anonymity and confidentiality:** No identifying information was collected; data were stored on encrypted university servers accessible only to the research team.
- **Right to withdraw:** Participants could withdraw at any point without penalty by closing the survey browser; partial responses were excluded from analysis.
- **Data protection:** All procedures complied with the UK General Data Protection Regulation (UK GDPR) and the Data Protection Act 2018.
- **Minimisation of harm:** No deceptive or distressing content was included; participants were offered access to institutional well-being resources.

3.14 Methodological Limitations

Several limitations warrant acknowledgement. First, the cross-sectional design precludes causal inference despite hypothesised directional relationships. Second, self-report data may be subject to social desirability and recall biases. Third, the sample, while stratified, may not fully represent all UK educators, particularly those in small independent schools or special educational needs settings. Fourth, the study did not differentiate between types of AI tools (e.g., generative AI vs. predictive analytics), potentially obscuring tool-specific effects. Fifth, non-response bias could not be fully assessed as non-respondent data were unavailable. These limitations are considered when interpreting results.

4. Data Analysis and Results

4.1 Descriptive Statistics

Table 4.1 presents the means, standard deviations, skewness, and kurtosis for all study variables.

Table 4.1: Descriptive Statistics (N = 387)

Variable	Mean	SD	Skewness	Kurtosis
Risk Barrier (RB)	5.24	1.18	-0.42	0.18
Usage Barrier (UB)	5.07	1.25	-0.31	-0.22
Enhanced Performance (EP)	4.18	1.42	-0.15	-0.56
Trust Prosperity (TP)	3.94	1.31	0.09	-0.48
Attitude (ATT)	4.21	1.38	-0.21	-0.52
Adoption Intention (ADI)	4.03	1.45	-0.08	-0.61

All variables measured on 7-point Likert scale (1 = Strongly Disagree, 7 = Strongly Agree).

Respondents reported relatively high levels of risk barriers (M = 5.24, SD = 1.18) and usage barriers (M = 5.07, SD = 1.25), indicating substantial concerns about AI adoption. Trust prosperity (M = 3.94, SD = 1.31) and adoption intention (M = 4.03, SD = 1.45) were near the scale midpoint, suggesting ambivalence among UK educators. Enhanced performance perceptions (M = 4.18, SD = 1.42) and attitude (M = 4.21, SD = 1.38) were slightly above neutral. Skewness and kurtosis values were within acceptable ranges (|skewness| < 2, |kurtosis| < 7), supporting normality assumptions (Byrne, 2016).

4.3 Bivariate Correlations

Table 4.2 presents Pearson correlation coefficients among all study variables.

Table 4.2: Correlation Matrix (N = 387)

Variable	1	2	3	4	5	6
1. Risk Barrier	1					
2. Usage Barrier	.48**	1				
3. Enhanced Performance	-.31**	-.39**	1			
4. Trust Prosperity	-.53**	-.49**	.44**	1		
5. Attitude	-.47**	-.45**	.52**	.61**	1	
6. Adoption Intention	-.41**	-.43**	.54**	.56**	.66**	1

**p < .01 (two-tailed).*

All correlations were in the hypothesised directions. Risk barrier and usage barrier exhibited moderate-to-strong negative correlations with trust prosperity ($r = -0.53$ and -0.49, respectively) and attitude ($r = -0.47$ and -0.45, respectively). Enhanced performance showed positive correlations with trust prosperity ($r = 0.44$), attitude ($r = 0.52$), and adoption intention ($r = 0.54$). The strongest correlation was between attitude and adoption intention ($r = 0.66$, $p < .01$). All inter-construct correlations were below the multicollinearity threshold of $r < 0.85$ (Hair et al., 2019).

4.4 Measurement Model Assessment

4.4.1 Confirmatory Factor Analysis

A six-factor confirmatory factor analysis (CFA) was estimated, with items loading onto their respective latent constructs (Risk Barrier, Usage Barrier, Enhanced Performance, Trust Prosperity, Attitude, and Adoption Intention). The measurement model demonstrated acceptable fit: $\chi^2(362) = 734.28$, $p < .001$; $\chi^2/df = 2.03$; CFI = 0.94; TLI = 0.93; RMSEA = 0.052 (90% CI [0.046, 0.058]); SRMR = 0.048. All fit indices met or exceeded recommended thresholds (Hu & Bentler, 1999).

4.4.2 Factor Loadings

Standardised factor loadings for all items are presented in Table 4.3. All loadings exceeded the 0.60 threshold (Hair et al., 2019), ranging from 0.68 to 0.88, indicating that each item adequately represented its respective latent construct.

Table 4.3: Standardised Factor Loadings

Construct	Item	Loading
Risk Barrier	RB1	0.78
	RB2	0.82

Construct	Item	Loading
	RB3	0.75
	RB4	0.71
	RB5	0.79
Usage Barrier	UB1	0.80
	UB2	0.84
	UB3	0.73
	UB4	0.76
	UB5	0.81
Enhanced Performance	EP1	0.82
	EP2	0.79
	EP3	0.85
	EP4	0.83
Trust Prosperity	TP1	0.76
	TP2	0.80
	TP3	0.73
	TP4	0.68
	TP5	0.79
	TP6	0.82
Attitude	ATT1	0.86

Construct	Item	Loading
	ATT2	0.88
	ATT3	0.77
	ATT4	0.84
Adoption Intention	ADI1	0.85
	ADI2	0.87
	ADI3	0.76
	ADI4	0.80

All loadings significant at p < .001.

4.4.3 Reliability and Convergent Validity

Table 4.4 reports internal consistency (Cronbach's α), composite reliability (CR), and average variance extracted (AVE) for each construct.

Table 4.4: Reliability and Convergent Validity

Construct	Cronbach's α	CR	AVE
Risk Barrier	0.87	0.88	0.60
Usage Barrier	0.89	0.89	0.63
Enhanced Performance	0.90	0.90	0.69
Trust Prosperity	0.88	0.89	0.58
Attitude	0.91	0.91	0.72
Adoption Intention	0.90	0.90	0.70

All Cronbach's α and CR values exceeded 0.70, indicating satisfactory internal consistency (Nunnally & Bernstein, 1994). All AVE values exceeded 0.50, establishing convergent validity (Fornell & Larcker, 1981).

4.4.4 Discriminant Validity

Discriminant validity was assessed using the Fornell-Larcker criterion (√AVE > inter-construct correlations) and the heterotrait-monotrait (HTMT) ratio.

Table 4.5: Fornell-Larcker Discriminant Validity Assessment

Construct	RB	UB	EP	TP	ATT	ADI
RB	**0.77**					
UB	0.48	**0.79**				
EP	-0.31	-0.39	**0.83**			
TP	-0.53	-0.49	0.44	**0.76**		
ATT	-0.47	-0.45	0.52	0.61	**0.85**	
ADI	-0.41	-0.43	0.54	0.56	0.66	**0.84**

Diagonal values (bold) are square roots of AVE; off-diagonal values are inter-construct correlations.

For all constructs, the square root of AVE (diagonal) exceeded the inter-construct correlations (off-diagonal), supporting discriminant validity. HTMT ratios (not tabulated) were all below 0.85 (range = 0.49–0.73), further confirming discriminant validity (Henseler et al., 2015).

4.5 Common Method Bias Assessment

As reported in Chapter 3, Harman's single-factor test indicated that a single factor accounted for 34.2% of total variance, below the 50% threshold (Podsakoff et al., 2003). Additionally, a common latent factor was included in the CFA, and the differences in standardised regression weights between the constrained and unconstrained models were minimal ($\Delta < 0.12$), suggesting that common method bias did not substantially affect the results.

4.6 Structural Model and Hypothesis Testing

4.6.1 Structural Model Fit

The hypothesised structural model was estimated with maximum likelihood. The model demonstrated acceptable fit: $\chi^2(370) = 791.45$, $p < .001$; $\chi^2/df = 2.14$; CFI = 0.93; TLI = 0.92; RMSEA = 0.055 (90% CI [0.049, 0.061]); SRMR = 0.052.

4.6.2 Direct Effects (Hypotheses H1–H7)

Table 4.6 presents the standardised path coefficients (β), standard errors (SE), critical ratios (z-values), and p-values for all hypothesised direct effects.

Table 4.6: Hypothesis Testing Results (Direct Effects)

Hypothesis	Path	β	SE	z	p	Result
H1	Risk Barrier → Trust Prosperity	-0.41	0.06	-7.98	< .001	Supported

Hypothesis	Path	β	SE	z	p	Result
H2	Risk Barrier → Attitude	-0.18	0.07	-3.21	.001	Supported
H3	Usage Barrier → Trust Prosperity	-0.32	0.06	-6.12	< .001	Supported
H4	Usage Barrier → Attitude	-0.14	0.06	-2.48	.013	Supported
H5	Trust Prosperity → Attitude	0.44	0.07	7.91	< .001	Supported
H6	Attitude → Adoption Intention	0.48	0.06	8.94	< .001	Supported
H7	Enhanced Performance → Adoption Intention	0.29	0.05	5.67	< .001	Supported

All seven hypotheses were supported. Risk barriers negatively predicted trust prosperity (β = -0.41, p < .001) and attitude (β = -0.18, p = .001). Usage barriers negatively predicted trust prosperity (β = -0.32, p < .001) and attitude (β = -0.14, p = .013). Trust prosperity positively predicted attitude (β = 0.44, p < .001). Attitude (β = 0.48, p < .001) and enhanced performance (β = 0.29, p < .001) both positively predicted adoption intention, with attitude exerting the strongest direct effect.

4.6.3 Variance Explained (R^2)

Table 4.7 presents the proportion of variance explained (R^2) for each endogenous variable.

Table 4.7: Variance Explained (R^2)

Endogenous Variable	R^2
Trust Prosperity	0.47
Attitude	0.58
Adoption Intention	0.51

The model explained 47% of the variance in trust prosperity, 58% of the variance in attitude towards AI technology, and 51% of the variance in adoption intention. These R^2 values are considered substantial in social science research (Cohen, 1988).

4.7 Mediation Analysis

Bootstrapping with 5,000 resamples and 95% bias-corrected confidence intervals was employed to test indirect effects. Table 4.8 presents the standardised indirect effects for all hypothesised mediation pathways.

Table 4.8: Indirect Effects (Bootstrapped 95% CI)

Path	Indirect Effect	SE	95% CI (Lower)	95% CI (Upper)	p
RB → TP → ATT → ADI	-0.087	0.019	-0.126	-0.052	.001
UB → TP → ATT → ADI	-0.067	0.015	-0.098	-0.040	.001
EP → TP → ATT → ADI	0.093	0.019	0.059	0.133	.001
RB → TP → ATT	-0.180	0.036	-0.252	-0.113	.001
UB → TP → ATT	-0.141	0.030	-0.202	-0.085	.001
EP → TP → ATT	0.194	0.038	0.124	0.271	.001
RB → ATT → ADI	-0.086	0.028	-0.144	-0.035	.002
UB → ATT → ADI	-0.067	0.027	-0.124	-0.019	.009
EP → ATT → ADI	0.210	0.040	0.137	0.292	.001

4.7.1 Sequential Mediation (Risk Barrier → Trust Prosperity → Attitude → Adoption Intention)

The indirect effect of risk barrier on adoption intention through trust prosperity and attitude was significant (β = -0.087, 95% CI [-0.126, -0.052]). This indicates that risk barriers reduce adoption intention indirectly by first diminishing trust prosperity, which in turn worsens attitude.

4.7.2 Sequential Mediation (Usage Barrier → Trust Prosperity → Attitude → Adoption Intention)

Similarly, usage barriers exerted a significant negative indirect effect on adoption intention via trust prosperity and attitude (β = -0.067, 95% CI [-0.098, -0.040]).

4.7.3 Sequential Mediation (Enhanced Performance → Trust Prosperity → Attitude → Adoption Intention)

Enhanced performance demonstrated a significant positive indirect effect on adoption intention through the sequential mediators (β = 0.093, 95% CI [0.059, 0.133]), indicating that performance gains enhance adoption intention partly by fostering trust and positive attitudes.

4.7.4 Specific Indirect Effects

The indirect effect of trust prosperity on adoption intention via attitude was substantial (β = 0.211, 95% CI [0.152, 0.278], not tabulated), confirming that attitude fully mediates the trust–intention relationship.

4.8 Summary of Hypotheses Testing

Table 4.9 summarises the results for all seven hypotheses.

Table 4.9: Summary of Hypothesis Testing Results

Hypothesis	Statement	Result
H1	Risk barrier → negative → Trust prosperity	Supported
H2	Risk barrier → negative → Attitude	Supported
H3	Usage barrier → negative → Trust prosperity	Supported
H4	Usage barrier → negative → Attitude	Supported
H5	Trust prosperity → positive → Attitude	Supported
H6	Attitude → positive → Adoption intention	Supported
H7	Enhanced performance → positive → Adoption intention	Supported

All seven hypotheses were supported at $p < .05$ or higher. The strongest direct effects were observed for H6 (Attitude → Adoption Intention, β = 0.48) and H5 (Trust Prosperity → Attitude, β = 0.44). The weakest but still significant direct effects were H4 (Usage Barrier → Attitude, β = -0.14) and H2 (Risk Barrier → Attitude, β = -0.18).

4.9 Additional Analyses

4.9.1 Multi-Group Comparisons

Exploratory multi-group analyses were conducted to examine whether path coefficients differed by educational sector (primary/secondary vs. further/higher education) and prior AI training (yes vs. no). Chi-square difference tests revealed no significant differences in structural paths across sectors ($\Delta\chi^2(7) = 9.24$, $p = .235$) or training status ($\Delta\chi^2(7) = 8.67$, $p = .278$), suggesting the model is largely invariant across these subgroups.

4.9.2 Comparison of Direct and Indirect Effects

Notably, the total effect of risk barrier on adoption intention (β = -0.173, not tabulated) was composed of a direct effect (β = -0.086, through attitude only) and an indirect effect (β = -0.087, through trust and attitude sequentially). Similarly, the total effect of enhanced performance on adoption intention (β = 0.383) comprised a direct effect (β = 0.290) and an indirect effect (β = 0.093). This indicates that both direct and mediated pathways are important for understanding adoption decisions.

5. Conclusion, Limitations, Future Directions and Policy Recommendations

5.1 Conclusion

This study confirms that successful AI adoption in UK education depends critically on trust and institutional capacity to mitigate barriers, not merely performance gains. Analysis of 387 educators revealed that risk barriers (β = -0.41) and usage barriers (β = -0.32) significantly reduced trust prosperity towards AI, which positively predicted attitude (β = 0.44). Attitude emerged as the strongest predictor of adoption intention (β = 0.48), followed by enhanced performance (β = 0.29). The sequential mediation model explained 51% of the variance in adoption intention, confirming significant indirect effects through the trust→attitude pathway. Descriptive statistics revealed high barrier perceptions (risk: M=5.24/7; usage: M=5.07/7) with only 37% of educators receiving AI training. These findings demonstrate that educational leaders balancing innovation with integrity must proactively invest in transparency, training, and trust-building rather than assuming performance gains alone will drive adoption.

5.2 Limitations

Several methodological constraints warrant consideration when interpreting these findings. The cross-sectional design precludes causal inference despite theoretically grounded hypotheses and tested mediation pathways; longitudinal data would be required to establish temporal precedence among variables. Self-report data may introduce social desirability and recall biases, as educators may over-report concerns about AI risks or under-report adoption intentions due to professional identity considerations. The stratified sample, while representing diverse educational sectors and UK nations, may not fully capture experiences of educators in special educational needs settings, early year's education, or small independent schools. The study did not differentiate between types of AI tools (e.g., generative AI chatbots versus predictive analytics), potentially obscuring tool-specific effects. Common method bias cannot be entirely eliminated, and organisational-level variables such as leadership support were not examined as potential moderators.

5.3 Future Directions

Longitudinal research should examine how trust prosperity and adoption intentions evolve as educators gain sustained, firsthand experience with AI tools over multiple academic terms. Comparative studies across different AI tool types—generative AI, adaptive learning platforms, automated assessment systems, and predictive analytics—would reveal whether risk barriers and trust mechanisms vary by application context. Qualitative investigations using interviews or focus groups are needed to understand how educators construct risk narratives, negotiate trust judgments, and experience the tension between innovation and integrity in daily practice. Intervention studies evaluating specific trust-building strategies (e.g., transparency dashboards, participatory design workshops, ethics training modules) would provide actionable causal evidence. Cross-national comparisons could identify how different regulatory frameworks and cultural contexts shape AI adoption pathways, while implementation science approaches could examine how policy recommendations translate into sustained practice change across diverse institutional settings.

5.4 Policy Recommendations

Based on findings that risk barriers (M=5.24/7), usage barriers (M=5.07/7), and low training rates (37%) significantly undermine trust and adoption intention, policymakers should mandate annual, context-specific AI literacy training for all educators addressing both technical competencies and ethical considerations. Institutions must establish transparency and explainability standards requiring AI tools to provide clear, accessible documentation of decision-making processes, algorithmic limitations, and data usage policies. Comprehensive risk mitigation frameworks should be developed covering data privacy protocols, regular algorithmic bias auditing, liability clarification for AI-generated outputs, and grievance mechanisms for affected stakeholders. Finally, inclusive AI governance structures should be created where educators, students, and technical experts co-design implementation strategies, ensuring alignment with professional values, pedagogical integrity, and the relational foundations of teaching. These measures require adequate funding, phased implementation timelines, and continuous evaluation.

References

Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes, 50*(2), 179–211.

Akgun, S., & Greenhow, C. (2022). Artificial intelligence in education: Addressing ethical challenges in K-12 settings. *AI and Ethics, 2*(3), 431–440.

Biesta, G. (2017). *The rediscovery of teaching*. Routledge.

Bond, M., Khosravi, H., De Laat, M., & Bergdahl, N. (2024). A meta-systematic review of artificial intelligence in higher education: Current trends and future directions. *British Journal of Educational Technology, 55*(2), 412–440.

Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education, 20*(1), 1–22.

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly, 13*(3), 319–340.

Davis, F. D., Bagozzi, R. P., & Warshaw, P. R. (1989). User acceptance of computer technology: A comparison of two theoretical models. *Management Science, 35*(8), 982–1003.

Department for Education. (2024). *Generative AI in education: Educator and learner perspectives*. HMSO.

Department for Education. (2025). *Consultation outcome: Generative AI in schools and further education*. HMSO.

Education Endowment Foundation. (2024). *Using AI to support teaching and learning: An evidence review*. EEF.

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals, 14*(2), 627–660.

Granić, A., & Marangunić, N. (2019). Technology acceptance model in educational context: A systematic literature review. *British Journal of Educational Technology, 50*(5), 2572–2593.

Hoff, K. A., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors, 57*(3), 407–434.

Holmes, W., Persson, J., Chounta, I. A., Wasson, B., & Dimitrova, V. (2023). Artificial intelligence and education: A critical view through the lens of human rights, democracy and the rule of law. *UNESCO Digital Library*.

Luckin, R., Cukurova, M., Kent, C., & du Boulay, B. (2022). Empowering educators with AI: Co-designing learning analytics tools. *Journal of Learning Analytics, 9*(3), 45–59.

Miao, F., & Holmes, W. (2025). *AI and education: Guidance for policy-makers in the United Kingdom*. UNESCO Publishing.

Nazaretsky, T., Cukurova, M., & Alexandron, G. (2022). Teachers' trust in AI-powered educational technology and their intention to use it. *Computers & Education, 191*, 104641.

Office for Students. (2025). *AI and academic integrity in English higher education: Regulatory advice*. OfS.

Ofsted. (2025). The impact of generative AI on curriculum and assessment: Annual report 2024/25. The Stationery Office.

Quality Assurance Agency for Higher Education. (2024). *AI in assessment: Guidance for UK higher education providers*. QAA.

Scherer, R., Siddiq, F., & Tondeur, J. (2019). The technology acceptance model (TAM): A meta-analytic structural equation modeling approach to explaining teachers' adoption of digital technology in education. *Computers & Education, 128*, 13–35.

Selwyn, N., Macgilchrist, F., & Williamson, B. (2024). The AI leadership paradox in education: Between innovation and integrity. *Educational Management Administration & Leadership*, online first.

Shin, D., & Park, Y. J. (2019). Role of fairness, accountability, and transparency in algorithmic affordance. *Computers in Human Behavior, 98*, 277–284.

Siau, K., & Wang, W. (2018). Building trust in artificial intelligence, machine learning, and robotics. *Cutter Business Technology Journal, 31*(2), 47–53.

Tait, A. (2023). Trust and artificial intelligence in higher education: A UK perspective. *Journal of Interactive Media in Education, 2023*(1), 1–12.

UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO Publishing.

Venkatesh, V., & Bala, H. (2008). Technology acceptance model 3 and a research agenda on interventions. *Decision Sciences, 39*(2), 273–315.

Venkatesh, V., & Davis, F. D. (2000). A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science, 46*(2), 186–204.

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly, 27*(3), 425–478.

Wang, Y., Liu, C., & Tu, Y. (2023). Factors affecting teachers' intention to adopt AI-assisted teaching: A trust–attitude–intention framework. *Interactive Learning Environments*, 1–18.

Williamson, B., & Macgilchrist, F. (2024). Critical perspectives on artificial intelligence in education: The UK policy context. *Learning, Media and Technology, 49*(1), 1–15.

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education, 16*(1), 1–27.