

Trading privacy for personalisation: Ethical trust and high school students' intention to use AI in education

Dr Sowmya Eashwaran^{*1}, Dr Manojkrishnan Champettil² and Ms Jiya Vivek³

^{1 2}EM Normandie Business School, Metis Lab, Dubai Knowledge Park, United Arab Emirates

³GEMS Wellington International, United Arab Emirates

Email: mchampettil@em-normandie.fr²; jiya.vivekv@gmail.com³

*Correspondence: svivek@em-normandie.fr

Abstract

Educational AI tools personalise learning by drawing on sensitive student data, yet the privacy calculus assumes that people can freely weigh benefits against risks and decline when risks dominate. This assumption is doubtful for adolescents, who have limited ability to scrutinise institutional data practices and little scope to opt out when AI is embedded in schooling. This study examines how perceived personalisation benefits, perceived data sensitivity and privacy concerns relate to high school students' intention to use AI, and tests whether ethical trust, defined as confidence that an AI system will handle personal data responsibly, transparently and in the user's interest, links privacy concerns to intention. Survey data from 323 students across IB, A Level, CBSE, American and other curriculums were analysed using partial least squares structural equation modelling. Perceived benefits strongly predicted ethical trust ($\beta = .605$, $p < .001$) and intention to use AI directly ($\beta = .454$, $p < .001$). Perceived data sensitivity raised privacy concerns ($\beta = .563$, $p < .001$), and privacy concerns modestly reduced ethical trust (β of negative .185, $p = .002$). Ethical trust in turn predicted intention to use AI ($\beta = .453$, $p < .001$), whereas the direct effect of privacy concerns on intention was negligible ($\beta = .023$, $p = .507$). The model explained 65.9% of the variance in intention. The pattern indicates that privacy concern shapes adoption through eroded trust rather than through avoidance, and that usefulness can build trust without proof of trustworthy data practices. Limitations include the cross sectional design, self reported intentions and weaker measurement of data sensitivity. Schools and providers should therefore demonstrate responsible data governance rather than simply reassure students.

Keywords: Privacy Calculus, Ethical Trust, Personalisation Privacy Paradox, Educational AI, Adolescent Privacy, Student Data, Technology Adoption, Structural Equation Modelling

JEL Classification: I21, D83, L86

1. Introduction

A student starts an AI tutor prior to an exam. Tutor asks students what they find difficult. It queries her about her previous test mark. It promises to keep in mind her weak points so that it can assist her out next time. She hesitates. Then she types the answer regardless. This paper addresses the question of that hesitation. It's tiny, it lasts a fraction of a second and it repeats millions of times per day. It is also the place where the ethics of AI-assisted education are being determined, not through policy documents, but when a young person decides to give up information about themselves to receive assistance (Kelso et al., 2024). Education Personalisation is based on personal data. As more information about the learner emerges, adaptive learning systems and intelligent tutors and generative AI assistants get better (Holmes & Tuomi, 2022; Zawacki-Richter et al., 2019; Akgun & Greenhow, 2022; Markauskaite et al., 2022). That data that makes them useful is the same data that is most sensitive: academic performance, learning difficulties, patterns of struggle, sometimes emotional state. This engine is the main force of the central tension of the field. The more you can disclose, the better the help will be. The more someone is disclosed, the greater the risk of misuse (Lee et al., 2024). This has traditionally been portrayed as the Personalisation–privacy paradox (Awad and Krishnan, 2006; Sutanto et al., 2013). The prevailing one is the privacy calculus, which suggests that people evaluate expected benefits and costs and disclose when benefits are greater than cost (Culnan and Armstrong, 1999; Dinev & Hart, 2006; Smith et al., 2011). There are three things that make the adolescent educational case unique. The disclosure is not fully voluntary, first. Where the AI tool is integrated into the school's learning platform, opting out has academic consequences (Regan and Jesse, 2019). Secondly, the data is institutionally stored and persistent. Academic record follows student, shopping history does not (Prinsloo and Slade, 2017). Third, adolescents are amid maturing their very skills, those used in the calculus model the ability to assess risk and engage in deliberative reasoning about consequences (Steinberg, 2008; Livingstone et al., 2019).

However, most of the literature on the privacy calculus is derived from adults, in business contexts and where there is real alternative (Kokolakis, 2017). AI systems that progressively instruct, evaluate, monitor, and profile pupils have received very little attention in the literature on children's privacy, in contrast, which has mostly concentrated on social media use and parental mediation (Stoilova et al., 2020). Considering how quickly AI technologies are being incorporated into educational contexts, this constitutes a substantial gap in the body of existing work. It is therefore becoming more crucial to comprehend how young people see these systems' acquisition and usage of their personal data. This study fills this gap. It poses a specific question: What factors influence high school students' decision to adopt AI Personalisation, given their sensitivity to their educational data?

Our answer is that it is not privacy concern that is the deciding variable. It is ethical trust, meaning that the student will feel the AI system is going to treat his/her information responsibly,

transparently and in his/her interest. The research proposes and tests a model where perceived Personalisation benefits and perceived data sensitivity act as two mediators between the development of privacy concerns and ethical trust and intention to use AI. The model was developed and estimated based on 323 secondary school students' responses using the PLS-SEM approach. The findings show a distinct pattern. Students in the sample had comparatively significant privacy concerns, especially with regard to the storage and use of personal data. Nevertheless, students' inclination to adopt AI was not significantly impacted by these worries ($\beta = .023$, $p = .507$). Rather, ethical trust in AI was considerably diminished by privacy concerns ($\beta = -.185$, $p < .01$). The intention to use AI was found to be strongly positively correlated with ethical trust ($\beta = .453$, $p < .001$). These results imply that students' direct use of AI is not always deterred by privacy concerns. Instead, it seems that their impact works via undermining ethical trust. To put it another way, students may still utilize AI in spite of worries about their personal information, but those worries may cause them to lose faith in the technology.

These results have significant applications for educational institutions and AI providers. Reassuring students or giving them more privacy information is a frequent way to address their privacy concerns. The results of this study, however, indicate that lowering privacy concerns might not be enough to affect students' intentions to utilize AI. Rather, ethical trust seems to be the more significant mechanism. This implies that schools and AI suppliers should concentrate on fostering trust thru open data practices, unambiguous accountability procedures, and verifiable proof that AI systems are created and applied in ways that safeguard and assist kids. These regulations require that institutions demonstrate proper data practices – not just provide a privacy notice – in ways that students can understand and evaluate. The paper makes three contributions. Theoretically, it adds ethical trust as the mechanism that balances benefit and risk, and it demonstrates that in the adolescent educational context the risk path is completely indirect. Empirically, it is one of the first to test the Personalisation-privacy trade-off using high school students as respondents, not adults. In essence, it shifts the design dialogue from concern-suppression to the building of trust (Bozkurt et al., 2024).

2. Theoretical Background and Hypothesis Development

2.1 The Privacy Calculus and its Boundaries

The privacy calculus states that disclosure is a rational exchange. People make cost-benefit analysis of giving information against the expected cost, and they provide information if the benefit is greater than cost (Culnan and Armstrong, 1999; Laufer and Wolfe, 1977). This reasoning was Formalised by Dinev and Hart (2006) in the context of e-commerce and it was concluded that trust and personal interest can override risk perceptions. In an overview of privacy research spanning the last 30 years, Smith et al. (2011) located the calculus as the prevailing explanatory model. Meta-analytic studies demonstrate that both perceived benefits and perceived risks are significant predictors of individuals' willingness to disclose personal information across a range of contexts

(Yun et al., 2019; Bandara et al., 2020). However, there are limitations to this reasoning and that has been an area of research of late.

The first is the privacy paradox itself the fact that, despite expressed concerns, Behavioural outcomes are hard to predict (Norberg et al., 2007; Kokolakis, 2017; Trepte et al., 2020). Acquisti et al. (2015) explained this by bounded rationality, meaning that individuals are incapable of calculating the long-term consequences of data disclosure and thus resort to heuristics, context cues and affect. Deliberative decision-making, however, does not always occur in practice. Dinev et al. (2015), in their “outside the APCO box” perspective, similarly challenge the assumption that privacy-related behaviour results from a straightforward and rational evaluation of risks and benefits. Writing from within the information systems (IS) tradition, they argue for a broader understanding of privacy decision-making that recognises the influence of contextual, psychological, and behavioural factors beyond deliberate privacy calculations.

The second edge is contextual. Nissenbaum (2010) theorizes privacy expectations as norm-governed and domain specific, which is termed contextual integrity. What is good data flow in one context is bad data flow in another. Empirically, Martin and Nissenbaum (2016) demonstrate that individuals' privacy judgments better reflect the context norms than the general concern trait.

The third, and most pertinent to the current study, is the presumption that individuals are the fully informed and autonomous decision makers who are able to assess the potential benefits against the potential privacy risks. This assumption is complicated in adolescence, when decision making can be influenced by emerging cognitive abilities, social factors, a lack of experience with the implications of privacy, and design of the digital environments where decisions are made.

2.2 Why Adolescents are a Different Case

Adolescents are not little adults who generate less data. Developmental studies indicate that the ability to consider delayed and possible costs develops later than the ability to take in immediate rewards (Steinberg, 2008; Casey et al., 2008). The advantage of an AI tutor is instant and colourfully apparent. The price of a data breach is far in the future and far away. If the two are assumed to be evaluated symmetrically as in a calculus model, that is the developmental fact that is being assumed away.

Livingstone et al. (2019), in their analysis of children's data literacy, found that although young people demonstrate a general awareness of privacy risks, they often struggle to translate this awareness into an understanding of the data practices embedded within the digital platforms they use in everyday life. Stoilova et al. (2020) further distinguish between interpersonal, institutional, and commercial dimensions of privacy, suggesting that children's understanding tends to be stronger in relation to interpersonal privacy and weaker in relation to institutional and commercial data practices. This distinction is especially pertinent in the context of educational AI, where the data is often gathered and processed within the confines of educational institutions and potentially

involving commercial technology vendors. Educational AI is therefore exactly in the zone of privacy which young people may be least aware of, and in this zone it is questionable whether adolescents can meaningfully assess the privacy risks of educational AI.

Another restriction relates to the assumption of meaningful consent. The privacy calculus makes the implicit assumption that people have a true choice: if the perceived risk of disclosure is greater than the perceived benefit, they can choose not to participate. In some commercial situations, this assumption might be true, but, in educational settings, opting out can have academic or practical implications. As digital technologies become increasingly integrated into education, Regan and Jesse (2019) point out some of the ethical issues that arise. In these situations, it may also be necessary to forgo access to course learning activities or experiences in order to refuse the use of an AI-based tool. This is an issue especially relevant to learning analytics. Prinsloo and Slade (2017) ask the question, can consent be truly informed and voluntary when there is significant power imbalance between the institution collecting the data and the institution analysing the data and the student whose data are being collected? This does not mean that the privacy calculus stops working under these circumstances; on the contrary, the explanatory power of the privacy calculus is more limited. There are times when a student feels that the privacy sacrifice is worthwhile because the apparent value of the alternative is too great. Consequently, the relationship between perceived privacy risk and technology adoption among students may differ from that predicted by the classical privacy calculus model. This provides an important rationale for placing ethical trust at the centre of the present study. Rather than assuming that adoption reflects a straightforward balancing of benefits and risks, the study examines whether students' willingness to use educational AI depends on their confidence that these systems, and the institutions deploying them, will collect and use personal data in ways that are transparent, responsible, and aligned with students' interests.

2.3 Ethical Trust as the Reconciling Mechanism

Trust generally is assumed to be a belief in competence and benevolence (Mayer et al., 1995; McKnight et al., 2002; Lee & See, 2004; Kumar et al., 2019). That framing has been expanded in the AI space to include beliefs about process integrity; transparency, accountability and non-exploitative use of personal data (Glikson and Woolley, 2020; Thiebes et al., 2021; Jobin et al., 2019; Mehrotra et al., 2024). The paper refers to this as ethical trust: a student's confidence that the AI system they are using will act responsibly, transparently, with respect to privacy and in user's best interest. It's not as broad as general trust, and not quite as abstract as "trust in technology". It is not about whether the system works, but the belief about how the system will treat user and his/her data.

The research proposes that the personalization–privacy trade-off is realized through ethical trust. Students can directly experience the benefits of personalization, but they are rarely able to verify the data practices that generate its privacy costs. They cannot normally observe data retention,

storage arrangements, third-party sharing or subsequent uses of their information. Under these conditions, students must instead form a judgement about whether the AI system is likely to treat them and their data fairly, transparently and responsibly. Ethical trust therefore acts as the mechanism through which an otherwise uncertain privacy calculus is resolved. Rather than simply supplementing the privacy calculus, trust may represent the form that the calculus takes when important risks and consequences cannot be directly verified. Based on this account, a specific prediction is made regarding how privacy concern relates to ethical trust and the adoption of AI. Rather than adoption being curtailed directly by privacy concern, it is assumed that privacy concern will impact the adoption indirectly. Rather, higher privacy concern should decrease students' ethical trust in the AI system and decreased ethical trust should, in turn, lower students' intention to use the system. Thus, privacy concern is an upstream/distant antecedent and ethical trust is a proximal mechanism that makes privacy concerns behaviourally relevant. The study now moves on to formalise these relationships as hypotheses.

2.4 Hypotheses

2.4.1 Perceived Personalisation benefits and ethical trust

Perceived Personalization Benefits (PPB) is the benefit that students perceive from giving an AI system personal information that allows the system to provide more suited and personally relevant learning support. These benefits are expected not only to correlate with ethical trust, but also to have a positive impact on it. When personalisation is actually seen to be enhancing the student's learning experience, it offers concrete examples of the use of information gathered by the system. When a system indicates that a misconception has been detected, an explanation has been adapted or a targeted piece of support has been provided and that this has helped the individual learner, this is evidence that personal data has been used in a beneficial way. Thus, perceived personalisation could act as a cue to the orientation of the system towards the student. This is in line with the results of personalised advertising, which found that personalisation was more likely to be trusted when it was seen as helpful rather than intrusive (Aguirre et al., 2015). The mechanism might be more important in an education context, where the gain is not promotional, but pedagogical; this is not about an advert, this could be to address a misconception or to enhance learning support. The evidence of the adoption of AI in higher education is also related, whereby usefulness leads to trust which leads to adoption (Chatterjee & Bhattacharjee, 2020).

H1. Perceived Personalisation benefits positively influence ethical trust in AI.

2.4.2 Perceived data sensitivity and privacy concerns

Perceived data sensitivity (PDS) refers to the perception that the data that the AI systems gather are personal, such as grades, difficulties, family context, mental health, etc. and that they should be protected.

The classic antecedent of concern is sensitivity. Malhotra et al. (2004) in creating the instrument IUIPC, assume sensitivity of the information class as one of the key determinants of concern over collection. The effect is demonstrated by Bansal et al. (2016): the same disclosure request leads to significantly different concern levels in case of financial or health information or trivial. The information in education isn't simple. Learning difficulties and mental health information is on the sensitive end of any reasonable taxonomy.

H2. Perceived data sensitivity positively influences privacy concerns.

2.4.3 Privacy concerns and ethical trust

Privacy concerns (PC) demonstrate that students are concerned about the way AI systems collect, store, share and utilize student information.

As opposed forces acting in parallel on behaviour, concern and trust are typically treated separately, they are modelled serially that is concern is an antecedent of trust. The justification is the verification problem set out above. Concern is about a suspicion about process it might over-collect, it might share without permission, it might repurpose. Trust is belief about process it will not. These are not independent input to a decision, the first is the negation of the second. If a student thinks data might be shared inappropriately, he or she is by construction a student who does not fully trust the system will be responsible with data.

Existing empirical research also indicates the proposed ordering. The paths from privacy concern to trust are significant (Bansal et al. 2016; Wu et al. 2012), thus suggesting that privacy concern is not just parallel to predicted behavior but also an antecedent of trust. Beldad et al. (2011) also propose that privacy concern can be viewed in its ability to undermine trust as a behavioural consequence. Taken together, these findings support the proposed sequence in which privacy concern operates upstream, weakening ethical trust, while ethical trust acts as the more proximal determinant of behavioral intention.

H3. Privacy concerns negatively influence ethical trust in AI.

2.4.4 Ethical trust and intention to use

Trust is a strong predictor of technology adoption in the IS literature (Gefen et al., 2003; Pavlou, 2003; McKnight et al., 2002). It plays a more focused role when uncertainty exists. Trust replaces verification when users are unable to verify systems behaviour, and adoption follows trust.

The Educational AI case is a case with high uncertainty. The student will not have the ability to view the model, the data pipeline or the vendors' agreements. Nazaretsky et al. (2022) identify trust as the key element influencing teachers' readiness to use AI-supported educational technology. Similarly, Choung et al. (2023) find that trust plays a pivotal role in the connection between perception and usage of AI voice assistants. Ethical trust is expected to be a proximal determinant of the intention to use AI by the students.

H4. Ethical trust in AI positively influences intention to use AI.

2.4.5 Perceived Personalisation benefits and intention to use

Perceived personalisation benefits are expected to influence intention to use AI through both direct and indirect pathways. First, benefit perceptions should exert a direct positive effect on intention to use. This relationship is central to technology acceptance models, beginning with the Technology Acceptance Model (TAM), in which perceived usefulness is a key determinant of behavioural intention (Davis, 1989), and continuing in later acceptance frameworks such as UTAUT (Venkatesh et al., 2003). Applied to AI-mediated education, students who perceive that personalisation provides more relevant feedback, better explanations, or learning support tailored to their individual needs should therefore be more willing to use the system, independently of their level of ethical trust.

Meanwhile, the impression of personalisation benefits can have an indirect impact via ethical trust. If learners feel that the information they are giving about themselves is being applied in a way that enriches learning, this can enhance their sense of the system's working on their behalf, lead to greater ethical trust and in turn greater intention to use. This model thus distinguishes between instrumental and signalling values of personalisation, with the latter contributing to ethical trust in a way that serves to facilitate the adoption of the beneficial personalisation.

This is also in line with the idea of the privacy-calculus approach. The calculus encompasses perceived personalisation benefits, which should also be incorporated into the adoption decision in their own right and not only as a function of trust. Ethical trust does not replace the role of perceived benefits; instead, it provides an additional mechanism through which those benefits can influence adoption under conditions of uncertainty. Accordingly, perceived personalisation benefits are expected both to directly increase intention to use AI and to indirectly increase intention through greater ethical trust.

H5. Perceived Personalisation benefits positively influence intention to use AI.

2.4.6 Privacy concerns and intention to use

A direct negative relationship between privacy concern and behavioural intention is a central prediction of the classical privacy calculus, in which greater perceived privacy risk or concern is expected to reduce willingness to disclose information or engage with a technology (Dinev & Hart, 2006). Accordingly, the present model retains a direct path from privacy concern to intention to use AI. However, there are three reasons to expect this relationship to be weak or potentially non-significant in the context of AI-mediated education.

First, research on the privacy paradox has repeatedly demonstrated that expressed privacy concerns do not consistently translate into corresponding behaviour. Individuals may report substantial concern about their privacy while continuing to disclose personal information and use data-

intensive technologies (Norberg et al., 2007; Kokolakis, 2017; Barth & de Jong, 2017). Privacy concern, therefore, does not necessarily produce behavioural avoidance.

Second, this concern–behaviour gap may be particularly pronounced in educational settings because students may have limited opportunities to avoid AI-mediated technologies. Even where use is not formally mandatory, students may depend on AI tools for explanations, feedback, homework assistance or other forms of learning support. The practical cost of non-use may therefore be sufficiently high that privacy concern cannot readily be translated into non-adoption. In this context, continued use should not necessarily be interpreted as evidence that privacy concerns are absent; rather, students may continue using a system despite those concerns.


Third, and most importantly, the theoretical model developed in this study proposes that the principal behavioural effect of privacy concern operates through ethical trust. Privacy concern is conceptualised as an upstream antecedent that weakens students' beliefs that the AI system will handle their data fairly, transparently and responsibly. Ethical trust, in turn, functions as the more proximal determinant of intention to use. If ethical trust substantially or fully mediates this relationship, little or no residual direct effect of privacy concern on intention should remain once ethical trust is included in the model.

Importantly, a failure to support H6 would not necessarily constitute a failure of the proposed theoretical model. Rather, when considered alongside a significant negative relationship between privacy concern and ethical trust and a significant positive relationship between ethical trust and intention to use, a non-significant direct effect would provide evidence consistent with the proposed trust-mediated account. It would suggest that privacy concern influences adoption primarily by altering ethical trust rather than by directly producing behavioural avoidance.

H6. Privacy concerns negatively influence intention to use AI.

Figure 1 presents the full model.

Figure 1. Conceptual model and structural results



3. Methodology

3.1 Research Design and Sample

A cross-sectional survey design was used for the study. The intended audience were high school students who had a certain level of interaction with AI tools in an education setting. Data was gathered from an online questionnaire (google form). No personally identifying information was recorded and participation was voluntary and anonymous. Participants were informed about the purpose of the study. Participants had the right to withdraw at any time and they were informed that their responses would not be viewed by teachers. The final sample consisted of 323 usable responses. The demographic profile of the 323 responses is reported in table 1.

Table 1. Sample Profile (n = 323)

Characteristic	Category	n	%
Age	12	206	63.8
	17	45	13.9
	16	34	10.5



		15	18	5.6
		14	12	3.7
		19	3	0.9
		13	3	0.9
		18	2	0.6
Gender	Female	184		57.0
	Male	139		43.0
Curriculum	Other	200		61.9
	IB	67		20.7
	A-Level	42		13.0
	CBSE	9		2.8
	American	5		1.5
AI use frequency	Sometimes	128		39.6
	Often	71		22.0
	Rarely	65		20.1
	Very Often	43		13.3
	Never	16		5.0
Primary AI tool	ChatGPT	175		54.2
	Gemini	62		19.2
	Other	40		12.4
	Claude	35		10.8
	Copilot	11		3.4

The sample has a gender balance with 43% of respondents being male and 57% of the respondents being female. All the different school curriculums were well represented in the sample. 20.7% of the respondents were from IB (20.7 percent) followed by 13% from A-Level (13.0 percent). Nearly all respondents use AI in some way with the maximum stating that they use AI sometimes. Only a small proportion (3.6 percent) say they never use AI for learning. ChatGPT is the biggest tool (54.2 percent) used by respondents followed by Gemini at 19.2 percent. There is one aspect of the age variable that needs to be explained. A large portion of the respondents ($n = 206$) checked the youngest age in the instrument, "12". The rest of the respondents fall in the usual 14-19 range. We view this as a characteristic of the response scale (probably a default choice or a misread response) and not as a real sample of 12-year-olds making up 2/3 of a high school sample. This is not an issue since age is not a construct in the structural model but is only used for descriptive purposes in the profiling of the sample. We mark it as a data-quality limitation and revisit it in Section 6.

3.2 Measures

All constructs were measured reflectively on multi-item scales, adapted from existing constructs and reworded for an adolescent education setting. A five-point Likert scale ranging from 1 (strongly disagree) to 5 (strongly agree) was used for all items.

Perceived Personalisation benefits (PPB) were adapted from Awad and Krishnan (2006) and Chellappa and Sin (2005) and consisted of 6 items that measured the belief that data sharing enhances the quality of AI educational support.

Perceived data sensitivity (PDS) was created based on Malhotra et al. (2004) and Bansal et al. (2016) and consisted of six items namely perceived sensitivity of grades, perceived sensitivity of learning difficulties, perceived sensitivity of family background, perceived sensitivity of mental health information and perceived sensitivity of personal goals.

Privacy concerns (PC): Seven items were used to measure privacy concerns (PC), drawing on the Internet Users' Information Privacy Concerns (IUIPC) scale developed by Malhotra et al. (2004) and the privacy-related dimensions identified by Smith et al. (1996). The items assessed students' concerns regarding data storage, unauthorized sharing of personal information, excessive data collection, inadequate protection of personal data, and potential misuse of their information.

Ethical trust in AI (ET) Seven items were used to measure ethical trust in AI (ET) based on the process-integrity principles highlighted in AI ethics frameworks (Jobin et al., 2019; Thiebes et al., 2021) and on McKnight et al. (2002). Students' perceptions of responsible data handling, ethical use of personal data, protection of privacy, whether AI systems act in their best interests, and confidence in AI when data practices are clear and transparent were assessed.

Intention to use AI (IU) was measured using six items adapted from Venkatesh et al. (2003); Pavlou (2003) which included continued use, willingness to use despite needing data, recommendation, and perceived future use.

Final loadings, list of items retained and dropped are given in the Appendix along with the complete wording for the items.

3.3 Instrument Refinement

There were two issues with the original instrument that were not included in the results, but were issues for design and had to be addressed before they could be estimated.

The first was that an intention was repeated within the intention block: the intention "AI will play an important role in my future learning" was repeated once. Before analysis, the duplicate item was removed resulting in 32 items.

Secondly, the original ethical trust block is a mix of two logically different sorts of statement. Four items measured belief (e.g., 'I trust AI systems to use student information responsibly'). The two were measured as widely held normative expectations: AI companies should be transparent about how they use student data" and "AI providers should be accountable for how they collect and use student information". This is not the same construct. Students might feel strongly about the need for AI companies to be transparent and accountable, but feel not very confident about their practices. The "should" items thus reflect expectations of ethical standards and not a judgement of ethical trust. Consistent with this conceptual distinction, the two normative items also performed as expected during measurement-model assessment (see Section 4.1) and were removed. Their removal was primarily a conceptual correction rather than a purely statistical decision; the items would have been excluded on theoretical grounds even if their psychometric performance had been satisfactory.

3.4 Analytical Approach

Partial least squares structural equation Modelling (PLS-SEM) was employed for analysis using the standard iterative outer-approximation algorithm with a Mode A (reflective) measurement scheme and a path-weighting inner scheme (Lohmöller, 1989; Wold, 1982). According to Hair et al. (2019), the decision of using PLS over covariance-based SEM was based on three reasons. First, the model is prediction-oriented, as the primary objective is to identify the factors that explain and predict students' intention to adopt AI. PLS-SEM is appropriate for this purpose because it focuses on maximising the explained variance in endogenous constructs, particularly the key outcome variable, intention to use AI. Second, the data are ordinal Likert responses that exhibit significant non-normality (several items have a skewness of greater than -1.0 , Table 3), and PLS makes no distributional assumption. Third, the model includes a mediation chain, in which the effect of one construct on another operates through an intervening variable. PLS-SEM is well suited to estimating such direct and indirect effects simultaneously and is less dependent on the strict distributional and identification assumptions associated with covariance-based maximum-likelihood (ML) estimation. The final sample size ($n = 323$) comfortably exceeded the minimum requirements for PLS-SEM. Applying the inverse square root method proposed by Kock

and Hadaya (2018), the minimum path coefficient of interest required a sample of approximately 155 observations. The available sample was therefore more than sufficient to achieve the statistical power required for estimating the proposed structural model.

Path coefficients were estimated and tested using bootstrapping. The analysis used 5000 bootstrap re-samples as in Streukens and Leroi-Werelds (2016) without sign changes. Results are given in terms of t-values, p-values and percentile bootstrap confidence interval which offer a more full evaluation of the precision and significance of the estimates. For the indirect effects, mediation was assessed using the bootstrap distribution of the product of the relevant path coefficients. This approach follows Preacher and Hayes (2008) and is recommended for mediation analysis in PLS-SEM by Nitzl et al. (2016). An indirect effect was considered statistically significant when its bootstrap confidence interval did not include zero. The traditional causal-steps procedure of Baron and Kenny (1986) was not used because subsequent research has shown that it can have low statistical power and does not require a significant total effect for mediation to exist (Zhao et al., 2010).

Q^2 based on ten-fold cross-validated redundancy was used to calculate the predictive relevance. Effect sizes are reported as f^2 . Multicollinearity was assessed at the item level (outer VIF) and the construct level (inner VIF) and used the conservative value of 3.3 established by Kock and Lynn (2012).

3.5 Common Method Bias

Common method variance is a potential concern since all data are gathered in the same manner at the same time using a single instrument (Podsakoff et al., 2003). Procedural remedies were incorporated into design. Anonymity of the respondents was assured and acknowledged. Items measuring the predictor and criterion constructs were separated in the instrument. No item was reverse coded and scale anchors were held constant to avoid introducing method factors that would not be part of the original method. Statistically, the single factor test, as developed by Harman, was used. The first unrotated factor explains only 31.95 per cent of the variance (which is below 50 per cent) and five factors have an eigenvalue of >1.0 . Following Podsakoff et al. (2003), we note that the Harman's test is a weak diagnostic and that its success is only necessary but not sufficient. Further evidence against common method bias is provided by the contrasting strength of relationships observed within the same survey instrument. The direct relationship between privacy concerns and intention to use AI was negligible and statistically non-significant ($\beta = .023$), whereas the relationship between perceived personalization benefits and intention to use AI was considerably stronger ($\beta = .454$). If a substantial common method effect were driving the results, a more consistent inflation of relationships among constructs measured using the same method might be expected. The presence of both negligible and strong path coefficients therefore provides some reassurance that the observed relationships are not simply artifacts of the common

measurement method. Inner VIFs were also checked for full collinearity, all of which are under the threshold of 3.3 that Kock (2015) considers a live issue of method bias (Table 6).

4. Results

4.1 Measurement Model: Initial Estimation and Purification

The results of the initial model estimation are reported before those of the purified model. This was a deliberate decision intended to ensure transparency in the measurement-model assessment. Reporting only the final purified model would prevent readers from evaluating the problems identified in the original specification and the basis for subsequent item removal. Importantly, these issues were not treated merely as statistical problems; they also revealed theoretically meaningful distinctions between the constructs and their measurement. Presenting both stages therefore makes the model-refinement process explicit and allows readers to assess the conceptual as well as the empirical justification for the changes. The results of the initial measurement model, comprising all 32 items, are presented in Table 2.

Table 2. Initial measurement model (all 32 items) items and constructs below threshold in bold

Construct	Cronbach's α	CR	AVE	Items below $\lambda = .70$
PPB	0.875	0.908	0.626	PPB6 (.584)
PDS	0.694	0.781	0.383	PDS1 (.550), PDS2 (.548), PDS3 (.686), PDS4 (.659), PDS5 (.399)
PC	0.904	0.925	0.640	
ET	0.761	0.828	0.419	ET2 (.354) , ET4 (.605), ET5 (.522) , ET1 (.668)
IU	0.897	0.921	0.661	

Two constructs failed convergent validity outright.

Ethical trust had an average variance extracted (AVE) of .419, which was below the recommended minimum threshold of .50 (Fornell and Larcker, 1981). The source of this problem was clearly identifiable. The two items framed as normative “should” statements—ET2, which assessed expectations of transparency, and ET5, which assessed expectations of accountability—had loadings of $\lambda = .354$ and $\lambda = .522$, respectively, the lowest loadings in the instrument. This is in

line with the conceptual approach taken in Section 3.3 where it is argued that these items represent expectations of ethics, but not the trust of ethics. Having faith that AI providers operate in a trustworthy and transparent way is different from expecting them to do so. A student might actually want more accountability if they don't have a lot of confidence in AI providers. The low loadings are consequently not simply psychometric "anomalies" but rather offer empirical evidence for the conceptual distinction between normative moral judgments and judgments of trust.

The perceived data sensitivity (PDS) also exhibited poor measurement properties, with an AVE of .383 and a Cronbach's α of .694, and three items had loadings of less than .60. This result has also a substantive interpretation. The PDS block requested students to rate the sensitivity of a variety of information, such as grades, learning difficulties, family background, mental health information, personal interests and goals, along with a behavioral item. These items may not function as interchangeable indicators of a single underlying perception of data sensitivity. Instead, students may reasonably regard different categories of information as having different levels of sensitivity. For example, a student may consider mental health information highly sensitive while viewing information about personal interests as considerably less sensitive, without any inconsistency in these judgments. This is particularly evident for PDS5, "*Information about personal interests and goals is private*," which had a loading of only $\lambda = .399$. The weak measurement performance therefore suggests that perceived data sensitivity may be type-specific rather than a single unidimensional reflective construct. Consequently, imposing a reflective specification may have required a degree of commonality among the indicators that the construct does not theoretically possess. The HTMT ratio for ethical trust and intention to use is .872 which is above the conservative HTMT value of .85 (Henseler et al., 2015), suggesting that the discriminant validity between the two constructs is questionable. Purification was conducted in two stages: first on conceptual grounds and then on psychometric grounds. Items that did not adequately represent the intended construct were removed based on theoretical considerations before the remaining items were evaluated using their statistical measurement properties. We removed:

- ET2 and ET5 normative expectation items that were error coded for a belief construct (Section 3.3);
- The weakest-loading sensitivity items are PDS1, PDS2 and PDS5 which focus on institutionally held, protection-warranting data (family background, mental health and Behavioural caution item);
- PPB6 ("*The benefits of personalised AI outweigh the inconvenience of sharing some personal information*", $\lambda = .584$) was also removed on conceptual grounds. On closer examination, the item does not measure perceived personalization benefits alone. Instead, it asks respondents to make an explicit trade-off between the benefits of personalization and the costs of sharing personal information. In this way, the item has actually implemented the privacy-calculus judgment that the structural model is meant to make

sense of. This would mean that a portion of the decision logic would be part of the model's predictor constructs, leading to a conceptual overlap between the measurement and structural aspects of the model for PPB6. In light of the need to maintain the construct validity and conceptual independence of perceived personalization benefits, the removal of this was thus necessary.

After the purification process, 26 items remained in the final measurement model. Elimination of 5 items on 2 constructs is not a minor statistical adjustment, but is regarded as a substantive change of the original measurement specification. Specifically, the PDS construct is measured with just three items, therefore not capturing the scope of perceived data sensitivity in full. This is not something that is ignored as insignificant, but it is explicitly acknowledged. Accordingly, the implications of the reduced measurement specification, particularly for PDS, are discussed as a study limitation in Section 6 and should be considered when interpreting the findings.

4.2 Measurement Model: Purified Solution

Table 3 reports item-level descriptives and final loadings. Table 4 reports construct-level reliability and convergent validity.

Table 3. Item descriptives and outer loadings (purified model)

Item	Mean	SD	Skew.	Kurt.	Loading
PPB1	3.53	0.95	-0.61	0.56	0.851
PPB2	3.80	1.00	-0.93	0.80	0.867
PPB3	3.27	1.01	-0.35	-0.20	0.830
PPB4	3.39	1.00	-0.51	0.10	0.800
PPB5	3.60	0.95	-0.59	0.38	0.814
PDS3	4.40	0.87	-1.41	1.53	0.729
PDS4	4.19	0.90	-0.77	-0.37	0.686
PDS6	4.48	0.77	-1.29	0.71	0.865
PC1	3.65	0.95	-0.15	-0.34	0.796
PC2	3.85	1.01	-0.64	-0.15	0.806

PC3	3.54	1.00	-0.05	-0.62	0.751
PC4	3.88	1.00	-0.68	0.11	0.734
PC5	3.72	0.88	-0.27	-0.24	0.778
PC6	3.87	0.89	-0.47	0.01	0.871
PC7	3.90	0.96	-0.60	-0.08	0.854
ET1	2.94	0.92	-0.06	0.18	0.726
ET3	3.76	0.85	-0.09	-0.60	0.751
ET4	2.91	0.88	-0.07	0.40	0.651
ET6	3.31	0.88	-0.25	0.54	0.763
ET7	3.60	0.92	-0.37	0.25	0.751
IU1	3.59	1.03	-0.56	0.25	0.868
IU2	2.94	1.05	-0.09	-0.51	0.750
IU3	3.38	0.99	-0.50	0.37	0.845
IU4	3.48	1.01	-0.47	0.05	0.775
IU5	3.46	0.98	-0.47	0.28	0.830
IU6	3.85	1.04	-0.86	0.60	0.806

Three constructs are unambiguously sound. Perceived Personalisation benefits ($\alpha = .890$, CR = .919, AVE = .694), privacy concerns ($\alpha = .904$, CR = .925, AVE = .640) and intention to use AI ($\alpha = .897$, CR = .921, AVE = .661) all clear every threshold with margin.

Ethical trust does meet the criteria but barely with AVE = .532 just above the .50 floor. Both composite reliability and α are acceptable at .850 and .782, respectively. One item, ET4 ("I trust AI systems to protect student privacy", $\lambda = .651$), falls below .70 standard. However we decided to retain it. The rationale is that Hair et al. (2019) allow for keeping items in the .40–.70 range as they do not increase AVE or CR beyond threshold and because they represent conceptual content the construct demands.

Table 4. Construct reliability and convergent validity (purified model)

Construct	Items	Cronbach's α	CR	AVE
Perceived Personalisation benefits (PPB)	5	0.890	0.919	0.694
Perceived data sensitivity (PDS)	3	0.659	0.806	0.583
Privacy concerns (PC)	7	0.904	0.925	0.640
Ethical trust in AI (ET)	5	0.782	0.850	0.532
Intention to use AI (IU)	6	0.897	0.921	0.661

ET4 was retained because it was the only remaining item that directly captured privacy protection within the ethical trust construct. Although removing the item might have produced a modest improvement in the construct's statistical indicators, doing so would have weakened its content validity by removing an important conceptual dimension of ethical trust. The decision therefore prioritized theoretical coverage over a purely statistical improvement in model fit. Accordingly, the convergent validity of ethical trust is considered acceptable, although relatively close to the recommended threshold and therefore interpreted with appropriate caution. As is explicitly stated, perceived data sensitivity is a weaker construct. AVE = .583 and CR = .806 are satisfactory. But Cronbach's α = .659 is still lower than the .7 criterion and PDS4 loads at .686, which is just below 0.70.

Rather than overlooking these results, they are reported transparently. Three considerations are important when interpreting this weaker measurement performance. First, α is a lower bound estimate and is always deflated by low numbers of indicators; the fact that there are only three indicators and α = .659 with CR = .806 is a familiar pattern; and composite reliability is the appropriate statistic in a PLS context precisely because it does not assume tau-equivalence (Hair et al., 2019; McNeish, 2018). Second, the substantive criteria for interpreting the structural interpretation are met AVE is greater than .50, indicating that more variance in the indicators is explained than error, and discriminate validity is obtained. The third and most important aspect is PDS in the structure model. It has the largest effect size in the study (f^2 = .464 on privacy concerns). If a construct is measured with some noise, and it still accounts for 31.7 per cent of the variance in

its outcome, then it isn't a construct that's failed, it's a construct whose true effect is, if anything, obscured by that noise. Nevertheless, the Cronbach's α of .659 remains a genuine measurement limitation, and conclusions regarding the magnitude and interpretation of the PDS effect should therefore be made with appropriate caution.

4.3 Discriminant Validity

The Fornell-Larcker criterion and the HTMT ratios are presented in Table 5.

Table 5. Discriminant validity: Fornell–Larcker (below diagonal, $\sqrt{\text{AVE}}$ on diagonal) and HTMT (above diagonal, italic)

	PPB	PDS	PC	ET	IU
PPB	0.833	<i>0.027</i>	<i>-0.072</i>	<i>0.720</i>	<i>0.808</i>
PDS	0.010	0.763	<i>0.660</i>	<i>-0.036</i>	<i>0.058</i>
PC	-0.066	0.563	0.800	<i>-0.292</i>	<i>-0.113</i>
ET	0.617	-0.047	-0.225	0.730	<i>0.844</i>
IU	0.732	0.015	-0.109	0.727	0.813

All constructs satisfied the Fornell–Larcker criterion, with the square root of the AVE for each construct exceeding its correlations with the other constructs. The highest inter-construct correlation was between perceived personalization benefits (PPB) and intention to use AI (IU), at .732, compared with square roots of AVE of .833 for PPB and .813 for IU. The results of the HTMT also confirmed the discriminant validity as all the values are less than the recommended .85 value. Notably, the ET–IU HTMT ratio was dropped from .872 in the initial model to .844 in the purified model, which is just below the .85 threshold. This finding must, however, be viewed with caution, and not held up as definitive evidence of separation. Ethical trust and intention to use AI are conceptually related: If a student thinks that an AI system is handling data in an ethical way, they may be more likely to use the system. The constructs, however, are still theoretically separable as ethical trust is a belief in the trustworthiness of the AI system, while intention to use is a behavioural disposition to use the AI system in the future. Thus, although the ET–IU value of .844 meets the adopted criterion, its proximity to the threshold indicates that the distinction is relatively close and should be considered when interpreting the ET → IU relationship.

Inspection of the cross-loadings provided additional evidence of discriminant validity. Each retained item loaded more strongly on its assigned construct than on any other construct in the model. No retained item exhibited a higher cross-loading on another construct, supporting the empirical distinctiveness of the measures.

4.4 Structural Model

Collinearity was diagnosed prior to the estimation of paths. The inner VIF values are all below the 3.3 threshold, ranging from 1.004 to 1.708. All outer VIFs are also OK. There is no collinearity affecting the coefficients.

The results of structural analysis are shown in Table 6.

Table 6. Structural Model: Direct Effects (5,000 Bootstrap Resamples)

H	Path	β	SE	t	p	95% CI	f ²	VIF	Result
H1	PPB → ET	0.605	0.041	14.90	<.001	[.524, .687]	0.621	1.004	Supported
H2	PDS → PC	0.563	0.037	15.39	<.001	[.493, .637]	0.464		Supported
H3	PC → ET	-0.185	0.060	-3.09	.002	[-.302, -.061]	0.058	1.004	Supported
H4	ET → IU	0.453	0.045	10.05	<.001	[.374, .551]	0.352	1.708	Supported
H5	PPB → IU	0.454	0.050	9.03	<.001	[.346, .546]	0.371	1.629	Supported
H6	PC → IU	0.023	0.035	0.66	.507	[-.043, .093]	0.001	1.063	Not supported

Five of the six hypothesised paths are supported. The sixth is not, and its failure is the most substantively interesting result in the study.

H1 is supported. Perceived Personalisation benefits raise ethical trust strongly ($\beta = .605$, $t = 14.90$, $p < .001$, 95% CI [.524, .687]). The effect size is large ($f^2 = .621$). Students who believe AI Personalisation helps them are markedly more likely to believe the AI will treat their data ethically.

H2 is supported. Perceived data sensitivity raises privacy concerns ($\beta = .563$, $t = 15.39$, $p < .001$, CI [.493, .637]), with a large effect ($f^2 = .464$), explaining 31.7 per cent of the variance in concern.

H3 is supported. Privacy concerns erode ethical trust ($\beta = -.185$, $t = -3.09$, $p = .002$, CI [-.302, -.061]). The direction is as predicted and the effect is significant. But the effect size is small ($f^2 = .058$) noticeably smaller than the benefit path acting on the same construct. This asymmetry matters, and we return to it in the discussion.

H4 is supported. Ethical trust raises intention to use ($\beta = .453$, $t = 10.05$, $p < .001$, CI [.374, .551]), with a large effect ($f^2 = .352$).

H5 is supported. Perceived Personalisation benefits raise intention to use directly ($\beta = .454$, $t = 9.03$, $p < .001$, CI [.346, .546]), also with a large effect ($f^2 = .371$).

H6 is not supported. Privacy concerns have no direct effect on intention to use AI ($\beta = .023$, $t = 0.66$, $p = .507$, CI [−.043, .093]). The effect size is practically nothing ($f^2 = .001$). The confidence interval is well above and well below zero and I should note that the point estimate is in the positive direction, not just that it is not significant in the direction it was hypothesised.

This should not be read as a 'null' result due to a lack of statistical power. The sample size $n = 323$, together with the bootstrap confidence interval [−.043, .093] indicates that the estimate is precise enough to rule out that there is a significant direct connection between privacy concerns and intention to use AI. Thus, the range indicates that the direct effect is unlikely to exceed about −.04. So, for this sample, privacy concerns seem to have no direct effect on decreasing the use intention of students. Their impact might rather manifest itself in other ways, namely by the means of ethical trust.

4.5 Mediation Analysis

Table 7 reports the indirect effects.

Table 7. Indirect Effects (5,000 Bootstrap Resamples)

Indirect path	Effect	SE	t	p	95% CI	Type
PPB → ET → IU	0.278	0.035	7.87	<.001	[.213, .352]	Complementary partial mediation
PC → ET → IU	−0.086	0.030	−2.86	.004	[−.147, −.027]	Full (indirect-only) mediation
PDS → PC → ET	−0.107	0.036	−2.94	.003	[−.180, −.033]	Significant
PDS → PC → ET → IU	−0.049	0.018	−2.71	.007	[−.087, −.014]	Significant serial mediation

Three mediating pathways are significant. The benefit route (PPB → ET → IU) carries an indirect effect of .278 ($t = 7.87$, $p < .001$, CI [.213, .352]). This combined with the direct effect of .454 gives a total effect of .732. Thus, the ethical trust is the carrier of 38.0 per cent of the total effect of benefit perception on adoption. The results showed that both direct and indirect effects were large and positive, which means that in the taxonomy of Zhao et al. (2010), this is a complementary

partial mediation: trust complements, but does not supplant, the utilitarian effect of perceived benefit. The concern route (PC → ET → IU) carries an indirect effect of $-.086$ ($t = -2.86$, $p = .004$, CI $[-.147, -.027]$). This direct effect is not significant. This is full (indirect only) mediation. There is no other channel that leads from privacy concern to adoption intention except for the ethical trust channel.

The serial route (PDS → PC → ET → IU) carries a significant indirect effect of $-.049$ ($t = -2.71$, $p = .007$, CI $[-.087, -.014]$). The complete causal pathway from data sensitivity to concern to trust to intention is present. It is indeed a road, but it's a simple road. The first-stage serial link (PDS → PC → ET) is likewise significant at $-.107$ ($t = -2.94$, $p = .003$). All together, these findings constitute the main empirical claim of this paper. Ethical trust is not one factor out of others. The only path through which privacy concern leads to Behaviour. If you take the trust away from the model, then privacy concern is causally dead in relation to adoption.

4.6 Explanatory and Predictive Power

Table 8 reports R^2 , adjusted R^2 and Q^2 .

Table 8. Explanatory and Predictive Power

Endogenous construct	R^2	Adjusted R^2	Q^2	Assessment
Privacy concerns (PC)	0.317	0.315	0.188	Moderate explanatory; predictive relevance established
Ethical trust in AI (ET)	0.414	0.411	0.214	Moderate explanatory; predictive relevance established
Intention to use AI (IU)	0.659	0.656	0.426	Substantial explanatory; medium-to-large predictive relevance

The model accounts for 65.9 per cent of the variance in intention to use AI (adjusted $R^2 = .656$). This is significant based on standards in the field (Hair et al., 2019) and is comparable to what has been found in the published literature on adoption in educational technology, which ranges from .40 to .60. Ethical trust is moderate with $R^2 = .414$ and privacy concerns $R^2 = .317$.

Ten folded cross-validated redundancy yields Q^2 values of .426, .214 and .188 for intention to use, ethical trust, and privacy concerns respectively, all of which are positive. All the values are positive indicating out-of-sample predictive relevance. The Q^2 for the dependent construct of .426 is a medium-to-large predictive relevance, meaning that the model does not just fit these data, it also predicts held out cases. The estimated model has an SRMR value of .097. This is within the .10 threshold value sometimes applied in the PLS approach but is higher than the .08 threshold value

that is conventionally applied in covariance-based SEM (Hu and Bentler, 1999; Henseler et al., 2014).

This issue is explicitly acknowledged as a limitation of the study. With PLS-SEM, the standardized root mean square residual (SRMR) is considered as an approximation to the model-fit measure and there is some debate on how to interpret it. For instance, Hair et al. (2019) warn that global fit indices are not necessarily reliable criteria for a composite-based method like PLS-SEM where the aim is prediction and variance explained and not covariance reproduction. Nevertheless, the observed SRMR of .097 is not particularly strong and should not be disregarded. It is also in line with the previously detected measurement limitations of the PDS and ET constructs. The study therefore cannot boast of a very good overall fit of the model. Instead, the results show good measures of predicted and explanatory power with the recognized weaknesses in certain aspects of the measurement model and these weaknesses are communicated in the interpretation of results.

4.7 Summary of Results Below Threshold

For transparency, Table 9 consolidates every criterion in this study that falls below its conventional threshold.

Table 9. Consolidated Record of Criteria Below Conventional Thresholds

Criterion	Value	Threshold	Status	Treatment
Cronbach's α PDS	0.659	≥ 0.70	Below	Retained. CR = .806 and AVE = .583 are met. α is deflated by the three-item count. Flagged as a limitation; PDS inferences carry this qualification.
AVE ET	0.532	≥ 0.50	Met, marginally	Retained. Clears the floor but with little margin. Flagged.
Loading ET4	0.651	≥ 0.70	Below	Retained per Hair et al. (2019): deletion would not raise AVE/CR above threshold, and ET4 is the only remaining item addressing privacy protection directly.
Loading PDS4	0.686	≥ 0.70	Below (marginal)	Retained. Within the .40–.70 acceptable band; deletion would reduce PDS to two items.

HTMT ET ↔ IU	0.844	< 0.85	Met, marginally	Clears the conservative threshold. Constructs are conceptually adjacent; readers should weigh the ET → IU path accordingly.
SRMR	0.097	< 0.08 (strict); < 0.10 (PLS)	Above strict threshold	Reported openly. Approximate fit is contested for composite-based PLS (Hair et al., 2019). We claim predictive performance, not excellent fit.
Age variable distribution	n = 206 at lowest option	Plausible cohort	Implausible	Treated as an instrument artefact. Age enters no structural path, so estimates are unaffected. Demographic profile should be read with caution.
H6 direct path	$\beta =$ 0.023, p = .507	Significant negative	Null	Not a defect this is the study's central substantive finding. Concern is fully mediated by ethical trust.

There are no conflicting major structural claims in the evidence reported in Table 9. But these diagnostic findings give important rationale for exercising due caution in interpreting those conclusions. Instead of providing a “pure” set of numbers while screening out areas of measurement weakness, the study explicitly describes these limitations, allowing readers to form their own judgments about the numbers in the context of all the evidence. Conclusions are thus appropriately qualified by the measurement and model-fit limitations described above and the structural results remain supported.

5. Discussion

5.1 The Finding That Matters

The students in this sample showed a high level of interest in the collection and use of their personal information. The results indicate concern regarding their data being shared without their permission or understanding ($M = 3.85$), misused by AI providers ($M = 3.87$), and used in ways they were not aware of ($M = 3.90$). In addition, the construct mean for perceived data sensitivity was the highest among the constructs in the study ($M = 4.36$). These findings suggest that the lack of a strong direct link between privacy concerns and intention to use AI does not mean that students are not concerned about privacy or are complacent. In contrast, students clearly view their information as personal and have substantive concerns about how this information might be used.

But these concerns do not impact students' AI usage plans directly. This aligns with the privacy paradox in an adolescent educational context: students express high levels of privacy concerns,

even as they continue to express intention to use AI. The results, however, are not merely a restatement of the privacy paradox, but an opportunity to see a potential mechanism through which privacy concerns function. The inaction of privacy concern, on the other hand, is not inaction but seems to be passed through ethical trust. Ethical trust is significantly negatively predicted by higher privacy concern ($\beta = -.185$, $p = .002$) and strongly and positively predicted by intention to use AI ($\beta = .453$, $p < .001$). Accordingly, the indirect effect of privacy concern on intention to use AI is significant ($\beta = -.086$, $p = .004$), but the direct effect is not ($\beta = .023$, $p = .507$). Therefore, privacy issues seem to impact adoption indirectly by reducing ethical trust of students in AI use, not by dissuading them from using it.

Based on this evidence, the privacy paradox cannot be reduced to a case of the lack of effects of privacy concern. Instead, privacy concern seems to impact behavioural intention by way of the belief that privacy concern negatively impacts—ethical trust. If a student is concerned about the misuse of their personal data they are not required to cease using an AI tutor. Rather, the student can persist in using the system and trust less and less. This erosion of trust is then reflected in the student's use intention of AI.

The interpretation provides an alternative account to that of either a lack of students' concern for privacy, or an inherently irrational and inconsistent student behaviour. The results indicate that there may be a more coherent process: If the decision to opt out has a significant impact on learning, students may feel very conscious of privacy and yet still use AI. In these situations, rejecting isn't always a viable technique for revealing privacy issues. For this reason, concern might take the form of a drop in trust in institutions and AI systems that manage students' data, instead of immediate disuse. This interpretation is supported by the results of this study, but the cross sectional design suggests that the proposed mechanism needs to be seen as an explanation, supported by the data, rather than proof of causality.

5.2 The Asymmetry of the Trust Equation

The second aspect of the findings relates to the two precursors of ethical trust: the impact of each is very different. The perceived personalization benefits are significantly related to ethical trust ($\beta = .605$, $f^2 = .621$), while privacy concerns are significantly but less related to ethical trust ($\beta = -.185$, $f^2 = .058$). Thus the standardized coefficient for perceived benefits is more than three times as big as the one for privacy concerns.

The pattern of this asymmetry indicates that students' ethical trust in AI is driven more by them perceiving it as helpful and useful, than by their privacy concerns. A helpful AI system can build significant trust in practice, even in the face of data practice concerns. This is a potentially significant and somewhat disconcerting discovery; because while there are privacy concerns that are not being resolved, it is possible that usefulness is outweighing that component in the creation of trust. So privacy concern seems less of a breaking force on trust and more of a dampening force in an otherwise powerful benefit-driven process.

But this is not a sign that privacy issues are not significant. Instead, it indicates that the functional value and ethical confidence may fuse in students' judgments; that is, students might believe in the system partly because it is working for them, even when issues of how their data are being used remain unresolved. This distinction matters a lot for schools and AI providers, as usefulness is not a justification for only those data practices that are demonstrably ethical.

If students' ethical trust in AI education is more about usefulness rather than a proxy for trustworthiness, then trust cannot be equated with trustworthiness. This is why an AI system could potentially gain the trust of students by offering them effective and personalised support in learning, even though practices of such AI systems have not been individually analysed or proven to be ethically correct. In this way, students can effectively give an “ethical” credit for pedagogical performance, even if there is no evidence that students' data are being collected, stored, shared, and handled responsibly.

This creates an important tension. The same perceived benefits that encourage students to adopt educational AI may also reduce the likelihood that its ethical practices are critically scrutinised. A system that is highly useful can therefore become highly trusted without necessarily becoming more trustworthy. The implication is that schools and AI providers should not rely on students expressed trust as evidence of ethical performance. Trustworthiness must instead be established independently through transparent data practices, meaningful accountability, privacy protection, and appropriate oversight.

5.3 Theoretical Implications

Three contributions follow:

First, the privacy calculus requires a mediator in constrained-choice settings. In the classical view, benefit and risk are two inputs to a disclosure decision that are parallel to each other (Dinev and Hart, 2006 and Culnan and Armstrong, 1999). In our data the parallel specification of the risk term is rejected. Risk does not directly come into the decision; it comes in through trust. This is a common pattern of low-exit privacy contexts (where not adopting is associated with a material cost). In such situations, the risk element cannot manifest itself behaviourally and therefore manifests itself attitudinally. One such context is education. Presumably other possibilities are healthcare and employment. In these settings the calculus does not break down, it reroutes.

Second, ethical trust is a distinguishable construct with distinct antecedents. The purification analysis also reveals an important conceptual distinction that has not always been made clearly in the literature: trust in AI and ethical expectations of AI are not the same construct. In the initial measurement model, the two normative items—ET2 and ET5—loaded weakly on ethical trust ($\lambda = .354$ and $\lambda = .522$, respectively). This suggests that students' expectations about how AI providers should behave are conceptually different from their beliefs about how AI providers actually behave. At the same time, students reported only moderate confidence in providers' actual

behaviour, as reflected in ET1 ($M = 2.94$) and ET4 ($M = 2.91$), while expressing stronger expectations regarding how AI providers ought to behave.

There is a substantive difference between the ethical expectation and the ethical belief. Students can make very strong demands for transparency, accountability, and privacy protection, which they don't feel AI providers are currently doing. This can therefore create an expectation–belief gap, which constitutes an ethical tension in education AI. Some measurement instruments that attempt to measure AI provider behavior and student expectations for such behavior may not discriminate between ethical trust and ethical expectations. Future research should therefore measure ethical expectations and ethical trust separately, as the difference between the two can then be an object of analysis, rather than merely a source of measurement error.

Third, the perception of data sensitivity is not a single dimension. This is not just a measurement issue – it is a finding of substance. The first reflective specification of the perceived data sensitivity was not very well ($AVE = .383$), which may indicate that students do not have an overall generic sensitivity perception that is manifested in the same way in each type of personal data. Their judgments, rather, seem to be typespecific. For instance, the item about mental health information was very close to the high sensitivity end of the scale ($M = 4.19$), while the item about interests and goals was not very common with the overall construct ($\lambda = .399$) and thus it was dropped.

In this way, Bansal et al. (2016) and Nissenbaum's (2010) theory of contextual integrity, which posits that privacy judgements vary based on information type and context, are supported. Thus, the poor performance of the initial PDS scale might more accurately be attributed to a problem in specification; strong assumptions are being made that different kinds of sensitivity judgments are equivalent expressions of the same underlying disposition. The results of the present study challenge this assumption and indicate that type specific or other measurement specifications of perceived data sensitivity should be taken into account for future research.

5.4 Practical Implications

The results of this research on this study have significant implications for schools and the providers of educational AI.

Reassurance is not an appropriate response. It seems the response to the student privacy issue by industry is communicative: privacy notices, FAQ pages, statements of commitment. These are intended to help ease the concern. We found that decreasing privacy concerns alone is not likely to yield a significant direct effect on students' intention to use AI. There was statistically weak direct association between privacy concerns and intention to use AI ($\beta = .023$, $p = .507$). This means that for schools and AI vendors, just trying to ease the minds of students about privacy will not be enough to impact adoption. The more significant pathway seems to be in relation to ethical trust and focus therefore should be placed on the practices that foster and enable ethical trust rather than just attempting to abate students' concerns. Concern-reduction is not the lever if the aim is

adoption. To the extent that the objective is ethical adoption, concern reduction efforts can backfire, as concern is not universally correlated with increased adoption and can also shield from the critical assessment of AI systems that students make.

That's the issue, trust is established through a proven benefit. In our model perceived usefulness ($\beta = .605$) is the strongest predictor. This discovery doesn't necessarily bode well for AI vendors. It indicates that learners might be ready to adopt systems that they find beneficial as long as there are some doubts on how their data is being managed. From a business standpoint, this provides an incentive to do whatever you can to appear helpful but not do the same with data governance. This means that securing ethical adoption will require regulatory protections and procurement requirements that ensure responsible data practices are a prerequisite for adoption.

Students are not held responsible; schools are held accountable. Concern is not refusal and responsibility for privacy should not be 100% with the student in a low-exit environment. Consent processes are not allowing for students' autonomy (Regan and Jesse, 2019; Prinsloo and Slade, 2017) where students have a nominal opt out. This task falls on the institutions at procurement, before the student ever sees the interface.

Design for verified not asserted stewardship. Clearly explained data practices (ET7: "I feel confident using AI tools that clearly explain how data is collected and used") had the highest mean score within the ethical trust construct ($M = 3.60$) compared to the mean score for the construct ($M = 3.30$). This indicates that transparency is a factor for ethical trust. But explanation is not enough. Transparency for adolescent users must be actionable, meaning it should be communicated in a way that is understandable to a teenager, it should be specific as to what data are collected and how they're used, and there must be some real control over this.

6. Limitations and Future Research

Measurement. There are two areas of measurement weakness. Following scale purification, Perceived Data Sensitivity was reduced to a three-item construct with Cronbach's $\alpha = .659$, below the conventional .70 threshold. Ethical Trust achieved an AVE of .532, only marginally above .50, and retained an item with a loading of .651. In addition, the ET-IU HTMT value of .844 approaches the .85 threshold for discriminant validity. Although each result is individually defensible, taken together they indicate an adequate rather than strong measurement model, a conclusion reinforced by the SRMR of .097. This is an area that can be explored with further research.

The PDS specification. As discussed in Section 5.3, this weakness may be diagnostic rather than merely psychometric. The poor performance of Perceived Data Sensitivity suggests that sensitivity may not be appropriately specified as a single reflective construct. A formative specification, in which sensitivity to different categories of data contributes to the overall construct, or separate type-specific constructs—for example, academic, health, and background data sensitivity—may

provide a more theoretically appropriate and informative representation. We did not respecify the construct in the present study because doing so would have altered the model originally specified for testing. Future research should therefore explicitly test and compare these alternative specifications rather than simply refine the existing reflective scale.

Causality. Because the study uses a cross-sectional design, the direction of the estimated relationships cannot be established empirically. Although we specify privacy concern as an antecedent of ethical trust on theoretical grounds (Section 2.4.3), the reverse ordering is also plausible: lower ethical trust may itself generate greater privacy concern. The fact that the proposed model is supported by the data does not resolve this issue, because model fit cannot establish causal or temporal direction. Experimental or longitudinal research is therefore required to distinguish between these competing explanations.

Common method variance. All measures were self-reported, collected at the same time, and obtained through a single questionnaire, creating the possibility of common method bias. Harman's single-factor test showed that the first factor accounted for 31.95% of the variance, providing some reassurance that a single common factor did not dominate the data; however, this test cannot rule out more selective forms of method-related inflation. Furthermore, no procedural measures to prevent the possibility of procedural effects were included in the study, such as reverse-coded items, a temporal separation of measures, or a marker variable. This is a limitation in method and not one that can be overcome by statistical reassurance at post hoc. These should be included in the design of instruments and in data collection in future studies.

Intention, not Behavior. The study is conducted as a measure of intention to use, not as a measure of actual use of AI. This distinction is important since the gap between intention and behaviour is well documented and may be more significant in privacy contexts, as reflected in the privacy paradox which motivates this study (Norberg et al., 2007). Therefore, the relationships found here should not be assumed to apply to students' usage of AI tools. A more robust test would connect the perceptual and attitudinal constructs addressed in this study with actual behavior, e.g., system usage logged. This would enable future studies to find out whether privacy concern, ethical trust, and perceived usefulness, predict the actual behavior of students, not simply what they say they are going to do.

Beyond these experimental research needs to be conducted to determine if trust can be increased when specific trust cues (transparency disclosures, accountability signals, meaningful control affordances, etc.) are manipulated apart from perceived usefulness, thereby increasing ethical trust. This is especially significant in the context of the dominant factor in the current model – perceived usefulness – because if trust can be gained through governance mechanisms rather than usefulness, this has significant implications for the design of responsible AI.

Also, there is a need for cross-cultural and cross-regulatory comparisons. Students who are subject to GDPR-equivalent data protections might have different values, risks, and trust issues concerning

pedagogical AI than students under a less explicit or less robust regulatory framework. Such comparisons would allow to determine if the privacy calculus observed here is domain-specific or not, and to determine its generalizability to other types of regulatory settings.

Future research may look at using longitudinally designed studies, where students are tracked as AI systems learn more about them over time. The present study captures privacy perceptions largely at one point in students' engagement with these technologies. Yet the more consequential question concerns what happens after repeated disclosure and prolonged data collection. Tracking students over several years would reveal whether continued exposure produces greater trust through familiarity, increasing concern as awareness develops, or gradual resignation to data collection. Understanding how this relationship evolves over the next five years may ultimately be more informative than any single cross-sectional estimate.

7. Conclusion

We started with a pupil who was reluctant to write the answer to a question in an AI tutor. We can now say something about what that hesitancy is composed of. It's not a computation of whether to use the tool. That in practice, is a question that has already been answered, homework must be done and the tool assists with that. It's a matter of whether to trust the instrument. We have demonstrated that belief is where the privacy concern resides.

This study revealed that high school students perceive data sensitivity and have high privacy concerns. Both do not directly decrease their intention to use AI. What they do is they undermine ethical trust that the system will deal with their information responsibly and it is ethical trust, along with perceived benefit, that drives adoption. The model accounts for 65.9 per cent of the variance in intention.

The bottom line is embarrassing and should not be sugar-coated. Students have no Behavioural outlet for concern, since they cannot refuse. Rather it is stored as a sense of loss of confidence in the institutions and systems educating them. Because usefulness ($\beta = .605$) is a much stronger driver of trust than suspicion ($\beta = -.185$), a very helpful AI system can continue to generate trust even when its data practices are not investigated.

The findings therefore do not necessarily describe a generation making a deliberate trade of privacy for personalisation. Rather, they point to students operating in contexts where the ability to refuse AI may be constrained, while the educational usefulness of these systems can generate trust despite unresolved questions about their data practices. The trade-off is real, but it may not always be freely chosen; the currency of that exchange is ethical trust, which students may have limited opportunity to withhold when AI becomes embedded in learning.

This extends the contribution beyond technology adoption to the ethics of AI in education. It raises a more fundamental question about what happens to privacy as a value when exercising it carries an educational cost. When saying no to data collection may also mean saying no to academic

support, convenience, or even opportunity, privacy is no longer simply a preference to be traded against personalisation. It becomes a right whose exercise has a price.

References

- Acquisti, A., Brandimarte, L. and Loewenstein, G. (2015), "Privacy and human behavior in the age of information", *Science*, Vol. 347 No. 6221, pp. 509-514.
<https://doi.org/10.1126/science.aaa1465>
- Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K. and Wetzels, M. (2015), "Unraveling the personalization paradox: the effect of information collection and trust-building strategies on online advertisement effectiveness", *Journal of Retailing*, Vol. 91 No. 1, pp. 34-49.
<https://doi.org/10.1016/j.jretai.2014.09.005>
- Akgun, S. and Greenhow, C. (2022), "Artificial intelligence in education: addressing ethical challenges in K-12 settings", *AI and Ethics*, Vol. 2 No. 3, pp. 431-440.
<https://doi.org/10.1007/s43681-021-00096-7>
- Awad, N.F. and Krishnan, M.S. (2006), "The personalization privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalization", *MIS Quarterly*, Vol. 30 No. 1, pp. 13-28.
<https://doi.org/10.2307/25148715>
- Bandara, R., Fernando, M. and Akter, S. (2020), "Explicating the privacy paradox: a qualitative inquiry of online shopping consumers", *Journal of Retailing and Consumer Services*, Vol. 52, 101947. <https://doi.org/10.1016/j.jretconser.2019.101947>
- Bansal, G., Zahedi, F.M. and Gefen, D. (2016), "Do context and personality matter? Trust and privacy concerns in disclosing private information online", *Information & Management*, Vol. 53 No. 1, pp. 1-21. <https://doi.org/10.1016/j.im.2015.08.001>
- Baron, R.M. and Kenny, D.A. (1986), "The moderator-mediator variable distinction in social psychological research: conceptual, strategic, and statistical considerations", *Journal of Personality and Social Psychology*, Vol. 51 No. 6, pp. 1173-1182.
<https://doi.org/10.1037//0022-3514.51.6.1173>
- Barth, S. and de Jong, M.D.T. (2017), "The privacy paradox – investigating discrepancies between expressed privacy concerns and actual online behavior – a systematic literature review", *Telematics and Informatics*, Vol. 34 No. 7, pp. 1038-1058.
<https://doi.org/10.1016/j.tele.2017.04.013>



- Beldad, A., de Jong, M. and Steehouder, M. (2011), "I trust not therefore it must be risky: determinants of the perceived risks of disclosing personal data for e-government transactions", *Computers in Human Behavior*, Vol. 27 No. 6, pp. 2233-2242. <https://doi.org/10.1016/j.chb.2011.07.002>
- Bozkurt, A., Xiao, J., Farrow, R., Bai, J.Y.H., Nerantzi, C., Moore, S., Dron, J., et al. (2024), "The manifesto for teaching and learning in a time of generative AI: a critical collective stance to better navigate the future", *Open Praxis*, Vol. 16 No. 4, pp. 487-513. <https://doi.org/10.55982/openpraxis.16.4.777>
- Casey, B.J., Getz, S. and Galvan, A. (2008), "The adolescent brain", *Developmental Review*, Vol. 28 No. 1, pp. 62-77. <https://doi.org/10.1016/j.dr.2007.08.003>
- Chatterjee, S. and Bhattacharjee, K.K. (2020), "Adoption of artificial intelligence in higher education: a quantitative analysis using structural equation modelling", *Education and Information Technologies*, Vol. 25 No. 3 pp. 3443-3463. <https://doi.org/10.1007/s10639-020-10159-7>
- Chellappa, R.K. and Sin, R.G. (2005), "Personalization versus privacy: an empirical examination of the online consumer's dilemma", *Information Technology and Management*, Vol. 6 Nos 2-3, pp. 181-202. <https://doi.org/10.1007/s10799-005-5879-y>
- Choung, H., David, P. and Ross, A. (2023), "Trust in AI and its role in the acceptance of AI technologies", *International Journal of Human-Computer Interaction*, Vol. 39 No. 3, pp. 1727-1739. <https://doi.org/10.1080/10447318.2022.2050543>
- Culnan, M.J. and Armstrong, P.K. (1999), "Information privacy concerns, procedural fairness, and impersonal trust: an empirical investigation", *Organization Science*, Vol. 10 No. 1, pp. 104-115. <https://doi.org/10.1287/orsc.10.1.104>
- Culnan, M.J. and Bies, R.J. (2003), "Consumer privacy: balancing economic and justice considerations", *Journal of Social Issues*, Vol. 59 No. 2, pp. 323-342. <https://doi.org/10.1111/1540-4560.00067>
- Davis, F.D. (1989), "Perceived usefulness, perceived ease of use, and user acceptance of information technology", *MIS Quarterly*, Vol. 13 No. 3, pp. 319-340. <https://doi.org/10.2307/249008>
- Dinev, T. and Hart, P. (2006), "An extended privacy calculus model for e-commerce transactions", *Information Systems Research*, Vol. 17 No. 1, pp. 61-80. <http://www.jstor.org/stable/23015781>
- Dinev, T., McConnell, A.R. and Smith, H.J. (2015), "Informing privacy research through information systems, psychology, and behavioral economics: thinking outside the 'APCO'

- box", *Information Systems Research*, Vol. 26 No. 4, pp. 639-655. <https://doi.org/10.1287/isre.2015.0600>
- Fornell, C. and Larcker, D.F. (1981), "Evaluating structural equation models with unobservable variables and measurement error", *Journal of Marketing Research*, Vol. 18 No. 1, pp. 39-50. <https://doi.org/10.2307/3151312>
- Gefen, D., Karahanna, E. and Straub, D.W. (2003), "Trust and TAM in online shopping: an integrated model", *MIS Quarterly*, Vol. 27 No. 1, pp. 51-90. <https://doi.org/10.2307/30036519>
- Glikson, E. and Woolley, A.W. (2020), "Human trust in artificial intelligence: review of empirical research", *Academy of Management Annals*, Vol. 14 No. 2, pp. 627-660. <https://doi.org/10.5465/annals.2018.0057>
- Hair, J.F., Risher, J.J., Sarstedt, M. and Ringle, C.M. (2019), "When to use and how to report the results of PLS-SEM", *European Business Review*, Vol. 31 No. 1, pp. 2-24. <https://doi.org/10.1108/EBR-11-2018-0203>
- Henseler, J., Dijkstra, T.K., Sarstedt, M., Ringle, C.M., Diamantopoulos, A., Straub, D.W., Ketchen, D.J., Hair, J.F., Hult, G.T.M. and Calantone, R.J. (2014), "Common beliefs and reality about PLS: comments on Rönkkö and Evermann (2013)", *Organizational Research Methods*, Vol. 17 No. 2, pp. 182-209. <https://psycnet.apa.org/doi/10.1177/1094428114526928>
- Henseler, J., Ringle, C.M. and Sarstedt, M. (2015), "A new criterion for assessing discriminant validity in variance-based structural equation modeling", *Journal of the Academy of Marketing Science*, Vol. 43 No. 1, pp. 115-135. <https://doi.org/10.1007/s11747-014-0403-8>
- Holmes, W. and Tuomi, I. (2022), "State of the art and practice in AI in education", *European Journal of Education*, Vol. 57 No. 4, pp. 542-570. <https://doi.org/10.1111/ejed.12533>
- Hu, L. and Bentler, P.M. (1999), "Cutoff criteria for fit indexes in covariance structure analysis: conventional criteria versus new alternatives", *Structural Equation Modeling*, Vol. 6 No. 1, pp. 1-55. <https://doi.org/10.1080/10705519909540118>
- Jobin, A., Ienca, M. and Vayena, E. (2019), "The global landscape of AI ethics guidelines", *Nature Machine Intelligence*, Vol. 1 No. 9, pp. 389-399.
- Kelso, E., Soneji, A., Rahaman, S., Shoshitaishvili, Y. and Hasan, R. (2024), "Trust, because you can't verify: privacy and security hurdles in education technology acquisition practices", *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security*, ACM, New York, NY, pp. 1656-1670. <https://doi.org/10.1145/3658644.3690353>



- Kock, N. (2015), "Common method bias in PLS-SEM: a full collinearity assessment approach", *International Journal of e-Collaboration*, Vol. 11 No. 4, pp. 1-10. <https://doi.org/10.4018/IJEC.2015100101>
- Kock, N. and Hadaya, P. (2018), "Minimum sample size estimation in PLS-SEM: the inverse square root and gamma-exponential methods", *Information Systems Journal*, Vol. 28 No. 1, pp. 227-261. <https://doi.org/10.1111/isj.12131>
- Kock, N. and Lynn, G.S. (2012), "Lateral collinearity and misleading results in variance-based SEM: an illustration and recommendations", *Journal of the Association for Information Systems*, Vol. 13 No. 7, pp. 546-580. <https://doi.org/10.17705/1jais.00302>
- Kokolakis, S. (2017), "Privacy attitudes and privacy behaviour: a review of current research on the privacy paradox phenomenon", *Computers & Security*, Vol. 64, pp. 122-134. <https://doi.org/10.1016/j.cose.2015.07.002>
- Kumar, P.C., Chetty, M., Clegg, T.L. and Vitak, J. (2019), "Privacy and security considerations for digital technology use in elementary schools", *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, pp. 1-13. <https://doi.org/10.1145/3290605.3300537>
- Laufer, R.S. and Wolfe, M. (1977), "Privacy as a concept and a social issue: a multidimensional developmental theory", *Journal of Social Issues*, Vol. 33 No. 3, pp. 22-42. <https://doi.org/10.1111/j.1540-4560.1977.tb01880.x>
- Lee, H.-P., Yang, Y.-J., Serban von Davier, T., Forlizzi, J. and Das, S. (2024), "Deepfakes, phrenology, surveillance, and more! A taxonomy of AI privacy risks", *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, pp. 1-19. <https://doi.org/10.1145/3613904.3642116>
- Livingstone, S., Stoilova, M. and Nandagiri, R. (2019), "Talking to children about data and privacy online: research methodology", Department of Media and Communications, London School of Economics and Political Science, London. <https://researchonline.lse.ac.uk/id/eprint/101284>
- Lohmöller, J.-B. (1989), *Latent Variable Path Modeling with Partial Least Squares*, Physica-Verlag, Heidelberg. <https://doi.org/10.1007/978-3-642-52512-4>
- Malhotra, N.K., Kim, S.S. and Agarwal, J. (2004), "Internet users' information privacy concerns (IUIPC): the construct, the scale, and a causal model", *Information Systems Research*, Vol. 15 No. 4, pp. 336-355. <https://doi.org/10.1287/isre.1040.0032>
- Markauskaite, L., Marrone, R., Poquet, O., Knight, S., Martinez-Maldonado, R., Howard, S., Tondeur, J., De Laat, M., Buckingham Shum, S., Gašević, D. and Siemens, G. (2022), "Rethinking the entwinement between artificial intelligence and human learning: what

- capabilities do learners need for a world with AI?", *Computers and Education: Artificial Intelligence*, Vol. 3, 100056. <https://doi.org/10.1016/j.caeai.2022.100056>
- Martin, K. and Nissenbaum, H. (2016), "Measuring privacy: an empirical test using context to expose confounding variables", *Columbia Science and Technology Law Review*, Vol. 18 No.1 pp. 176-218. <https://doi.org/10.7916/stlr.v18i1.4015>
- Martin, K.D., Borah, A. and Palmatier, R.W. (2017), "Data privacy: effects on customer and firm performance", *Journal of Marketing*, Vol. 81 No. 1, pp. 36-58. <https://doi.org/10.1509/jm.15.0497>
- Mayer, R.C., Davis, J.H. and Schoorman, F.D. (1995), "An integrative model of organizational trust", *Academy of Management Review*, Vol. 20 No. 3, pp. 709-734. <https://doi.org/10.2307/258792>
- McKnight, D.H., Choudhury, V. and Kacmar, C. (2002), "Developing and validating trust measures for e-commerce: an integrative typology", *Information Systems Research*, Vol. 13 No. 3, pp. 334-359. <https://doi.org/10.1287/isre.13.3.334.81>
- McNeish, D. (2018), "Thanks coefficient alpha, we'll take it from here", *Psychological Methods*, Vol. 23 No. 3, pp. 412-433. <https://doi.org/10.1037/met0000144>
- Mehrotra, S., Degachi, C., Vereschak, O., Jonker, C.M. and Tielman, M.L. (2024), "A systematic review on fostering appropriate trust in human-AI interaction: trends, opportunities and challenges", *ACM Journal on Responsible Computing*, Vol. 1 No. 4, pp. 1-45. <https://doi.org/10.1145/3696449>
- Nazaretsky, T., Ariely, M., Cukurova, M. and Alexandron, G. (2022), "Teachers' trust in AI-powered educational technology and a professional development program to improve it", *British Journal of Educational Technology*, Vol. 53 No. 4, pp. 914-931. <https://doi.org/10.1111/bjet.13232>
- Nissenbaum, H. (2010), *Privacy in Context: Technology, Policy, and the Integrity of Social Life*, Stanford University Press, Stanford, CA.
- Nitzl, C., Roldán, J.L. and Cepeda, G. (2016), "Mediation analysis in partial least squares path modeling: helping researchers discuss more sophisticated models", *Industrial Management & Data Systems*, Vol. 116 No. 9, pp. 1849-1864. <https://doi.org/10.1108/IMDS-07-2015-0302>
- Norberg, P.A., Horne, D.R. and Horne, D.A. (2007), "The privacy paradox: personal information disclosure intentions versus behaviors", *Journal of Consumer Affairs*, Vol. 41 No. 1, pp. 100-126. <https://doi.org/10.1111/j.1745-6606.2006.00070.x>



- Pavlou, P.A. (2003), "Consumer acceptance of electronic commerce: integrating trust and risk with the technology acceptance model", *International Journal of Electronic Commerce*, Vol. 7 No. 3, pp. 101-134. <https://doi.org/10.1080/10864415.2003.11044275>
- Podsakoff, P.M., MacKenzie, S.B., Lee, J.-Y. and Podsakoff, N.P. (2003), "Common method biases in behavioral research: a critical review of the literature and recommended remedies", *Journal of Applied Psychology*, Vol. 88 No. 5, pp. 879-903. <https://doi.org/10.1037/0021-9010.88.5.879>
- Preacher, K.J. and Hayes, A.F. (2008), "Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models", *Behavior Research Methods*, Vol. 40 No. 3, pp. 879-891. <https://doi.org/10.3758/BRM.40.3.879>
- Prinsloo, P. and Slade, S. (2017), "An elephant in the learning analytics room: the obligation to act", *Proceedings of the Seventh International Learning Analytics & Knowledge Conference, ACM*, New York, NY, pp. 46-55. <https://doi.org/10.1145/3027385.3027406>
- Regan, P.M. and Jesse, J. (2019), "Ethical challenges of edtech, big data and personalized learning: twenty-first century student sorting and tracking", *Ethics and Information Technology*, Vol. 21 No. 3, pp. 167-179. <https://doi.org/10.1007/s10676-018-9492-2>
- Slade, S. and Prinsloo, P. (2013), "Learning analytics: ethical issues and dilemmas", *American Behavioral Scientist*, Vol. 57 No. 10, pp. 1510-1529. <https://doi.org/10.1177/0002764213479366>
- Smith, H.J., Dinev, T. and Xu, H. (2011), "Information privacy research: an interdisciplinary review", *MIS Quarterly*, Vol. 35 No. 4, pp. 989-1015. <https://doi.org/10.2307/41409970>
- Smith, H.J., Milberg, S.J. and Burke, S.J. (1996), "Information privacy: measuring individuals' concerns about organizational practices", *MIS Quarterly*, Vol. 20 No. 2, pp. 167-196. <https://doi.org/10.2307/249477>
- Steinberg, L. (2008), "A social neuroscience perspective on adolescent risk-taking", *Developmental Review*, Vol. 28 No. 1, pp. 78-106. <https://doi.org/10.1016/j.dr.2007.08.002>
- Stoilova, M., Livingstone, S. and Nandagiri, R. (2020), "Digital by default: children's capacity to understand and manage online data and privacy", *Media and Communication*, Vol. 8 No. 4, pp. 197-207. <https://doi.org/10.17645/mac.v8i4.3407>
- Streukens, S. and Leroi-Werelds, S. (2016), "Bootstrapping and PLS-SEM: a step-by-step guide to get more out of your bootstrap results", *European Management Journal*, Vol. 34 No. 6, pp. 618-632. <https://doi.org/10.1016/j.emj.2016.06.003>



- Sutanto, J., Palme, E., Tan, C.-H. and Phang, C.W. (2013), "Addressing the personalization–privacy paradox: an empirical assessment from a field experiment on smartphone users", *MIS Quarterly*, Vol. 37 No. 4, pp. 1141-1164. <https://doi.org/10.25300/MISQ/2013/37.4.07>
- Thiebes, S., Lins, S. and Sunyaev, A. (2021), "Trustworthy artificial intelligence", *Electronic Markets*, Vol. 31 No. 2, pp. 447-464. <https://doi.org/10.1007/s12525-020-00441-4>
- Trepte, S., Scharkow, M. and Dienlin, T. (2020), "The privacy calculus contextualized: the influence of affordances", *Computers in Human Behavior*, Vol. 104, 106115. <https://doi.org/10.1016/j.chb.2019.08.022>
- Venkatesh, V., Morris, M.G., Davis, G.B. and Davis, F.D. (2003), "User acceptance of information technology: toward a unified view", *MIS Quarterly*, Vol. 27 No. 3, pp. 425-478. <https://doi.org/10.2307/30036540>
- Wold, H. (1982). "Soft modeling: the basic design and some extensions", *Systems under indirect observation, Part II*, 2, 36-37.
- Wu, K.-W., Huang, S.Y., Yen, D.C. and Popova, I. (2012), "The effect of online privacy policy on consumer privacy concern and trust", *Computers in Human Behavior*, Vol. 28 No. 3, pp. 889-897. <https://doi.org/10.1016/j.chb.2011.12.008>
- Yun, H., Lee, G. and Kim, D.J. (2019), "A chronological review of empirical research on personal information privacy concerns: an analysis of contexts and research constructs", *Information & Management*, Vol. 56 No. 4, pp. 570-601. <https://doi.org/10.1016/J.IM.2018.10.001>
- Zawacki-Richter, O., Marín, V.I., Bond, M. and Gouverneur, F. (2019), "Systematic review of research on artificial intelligence applications in higher education – where are the educators?", *International Journal of Educational Technology in Higher Education*, Vol. 16 No. 1, 39. <https://doi.org/10.1186/s41239-019-0171-0>
- Zhao, X., Lynch, J.G. and Chen, Q. (2010), "Reconsidering Baron and Kenny: myths and truths about mediation analysis", *Journal of Consumer Research*, Vol. 37 No. 2, pp. 197-206. <https://doi.org/10.1086/651257>

Appendix A. Measurement instrument, initial and final loadings

All items measured on a five-point Likert scale (1 = strongly disagree, 5 = strongly agree). "Initial λ " is the outer loading in the full 32-item solution; "Final λ " is the loading in the purified 26-item model. Every dropped item and the reason for its removal is recorded.

Table A1. Full instrument with initial and final outer loadings

Code	Item wording	Initial λ	Final λ	Status
PPB1	AI provides more useful educational support when it knows more about students.	0.847	0.851	Retained
PPB2	Personalised AI tools can better meet my learning needs.	0.854	0.867	Retained
PPB3	Sharing information with AI can improve the quality of educational assistance.	0.828	0.830	Retained
PPB4	AI can provide more accurate learning recommendations when it has access to student data.	0.791	0.800	Retained
PPB5	Personalised AI systems can improve student learning experiences.	0.811	0.814	Retained
PPB6	The benefits of Personalised AI outweigh the inconvenience of sharing some personal information.	0.584	DROPPED measures the trade-off itself, not benefit perception	
PDS1	Academic grades are sensitive personal information.	0.550	DROPPED weak loading; sensitivity is type-specific	

PDS2	Learning difficulties should be considered sensitive information.	0.548		DROPPED weak loading; sensitivity is type-specific
PDS3	Family background information is highly personal.	0.686	0.729	Retained
PDS4	Information about mental health should be treated as highly sensitive.	0.659	0.686	Retained
PDS5	Information about personal interests and goals is private.	0.399		DROPPED weakest loading in instrument
PDS6	Students should be careful about sharing personal information with AI systems.	0.796	0.865	Retained
PC1	I worry about how AI systems store student data.	0.798	0.796	Retained
PC2	I am concerned that my information could be shared without my permission.	0.809	0.806	Retained
PC3	AI systems collect too much personal information.	0.750	0.751	Retained
PC4	I am uncomfortable sharing certain personal information with AI tools.	0.733	0.734	Retained
PC5	Educational AI systems may not adequately protect student privacy.	0.775	0.778	Retained
PC6	I worry that AI providers could misuse student data.	0.870	0.871	Retained

PC7	I am concerned that my personal information may be used for purposes I am unaware of.	0.856	0.854	Retained
ET1	I trust AI systems to handle student information responsibly.	0.668	0.726	Retained
ET2	AI companies should be transparent about how they use student data.	0.354		DROPPED normative expectation, not trust belief
ET3	I believe educational AI systems can be used ethically.	0.776	0.751	Retained
ET4	I trust AI systems to protect student privacy.	0.605	0.651	Retained
ET5	AI providers should be accountable for how they collect and use student information.	0.522		DROPPED normative expectation, not trust belief
ET6	I believe educational AI systems generally act in students' best interests.	0.741	0.763	Retained
ET7	I feel confident using AI tools that clearly explain how data is collected and used.	0.755	0.751	Retained
IU1	I intend to continue using AI tools for learning.	0.868	0.868	Retained
IU2	I would use AI even if it requires some personal information.	0.751	0.750	Retained

IU3	I would recommend educational AI tools to other students.	0.844	0.845	Retained
IU4	AI will play an important role in my future learning.	0.775	0.775	Retained
IU5	I am willing to use AI systems that provide Personalised educational support.	0.829	0.830	Retained
IU6	I would continue using AI if I trusted how my data was handled.	0.806	0.806	Retained
