

Risk Perception as a Moderating Variable in Generative AI Adoption Among Business School Students

AIVARS SPILBERGS

Economics and Finance Department

BA School of Business and Finance, University of Latvia

Kr. Valdemara 161

Riga

LV1013,

LATVIA

Abstract: - Generative artificial intelligence (AI) tools are spreading fast through higher education, reshaping how students learn. Yet what actually drives students to adopt these tools - especially in business schools, where ethical sensitivity and interpretive judgment matter as much as technical skill - is still not well understood. This study looks at how risk perception shapes generative AI adoption, extending the Unified Theory of Acceptance and Use of Technology 2 (UTAUT2) model to do so. Survey data collected from 441 business students in the Baltic Sea region were analyzed using partial least squares structural equation modeling (PLS-SEM) to evaluate the effects of performance expectancy, effort expectancy, social influence, study value, habit, and risk perception on AI usage. The model demonstrates substantial explanatory power ($R^2 = 0.654$), with habit identified as the most significant predictor, followed by risk perception and social influence. Students rated usefulness, ease of use, and study value highly - yet these factors mattered less than behavioral ones in predicting actual use. The findings point to a clear need: ethical awareness and AI literacy should be built directly into technology-adoption frameworks in business education, not treated as an afterthought.

Key-Words: - Business education, Generative AI, Moderating impact, PLS-SEM, Risk perception, UTAUT2.

Received: June 26, 2026. Revised: July 29, 2026. Accepted: August 19, 2026. Published: October 1, 2026.

1 Introduction

The rapid emergence of generative artificial intelligence (AI) technologies has initiated significant transformations across various sectors, with a pronounced impact on educational paradigms in business schools [1]. Advancements such as sophisticated language models and image generators are altering pedagogical strategies and influencing student engagement with academic tasks [2]. The widespread adoption of these tools presents substantial challenges concerning academic integrity, the cultivation of critical thinking, and the preservation of rigorous academic standards [3]. Consequently, it is essential to identify the factors that affect the adoption and integration of generative AI in academic environments, especially within business schools, to inform the development of effective guidelines and curricula. Addressing this issue requires a focused analysis of how perceived risks, including academic fraud, data inaccuracy, and ethical concerns, influence business students' intentions to use generative AI tools. Although these technologies are increasingly prevalent, empirical research on business students' perceptions and use of

generative AI remains limited, particularly regarding satisfaction, trust, and perceived risks.

This study addresses a significant research gap by investigating how risk perception moderates business school students' intentions to adopt generative AI, employing established technology acceptance frameworks. We build on this by adding perceived risk as a moderating variable to the UTAUT model, asking what really drives generative AI adoption among business students once ethical concerns are factored in. Untangling how these factors interact matters beyond theory - it can inform policy and teaching practice as AI becomes a fixture in academic life. UTAUT already offers a solid base for studying adoption behavior; adding perceived risk sharpens that lens further.

Perceived risk, in particular, still needs closer scrutiny - its role in technology acceptance research remains murky, especially regarding how it moderates adoption behavior. Here, we ask a more pointed question: does risk perception push business students away from generative AI, or does it make them more careful, deliberate adopters rather than outright resistant ones? Answering this matters for

anyone trying to reconcile AI's disruptive potential with the academic rigor and ethical standards business education is supposed to uphold. With that in mind, the study offers practical guidance for educators, policymakers, and developers working to integrate AI responsibly into educational and professional settings. Much of the existing literature has stopped at students' first encounters with generative AI and their initial intentions to use it - far less attention has gone to what happens after adoption, once use becomes routine and satisfaction (or dissatisfaction) sets in.

2 Literature review

2.1 Generative AI in higher education

Generative AI tools represent a significant technological advancement with wide-ranging implications for higher education, enabling enhanced content creation and fostering critical analysis [4].

Although generative AI tools have attracted substantial attention, previous research on their adoption in higher education has primarily focused on perceived ease of use and usefulness, often neglecting important moderating factors such as risk. This oversight is problematic because, although ethical considerations and potential misuses of AI tools are widely recognized, they remain largely underexplored in empirical research, creating a significant gap in understanding their interactive dynamics [5]. Consequently, empirical studies are necessary to investigate students' motivations for both the initial and sustained adoption of generative AI, as the determinants of continued use differ markedly from those that influence initial engagement [6]. Identifying factors that drive ongoing use, such as perceived performance expectancy and ethical judgment, is essential for evaluating the long-term integration of generative AI into academic practices. Moreover, the rapid proliferation of generative AI has generated considerable interest in understanding how students perceive and intend to use these tools for academic purposes [7].

Although interest in technology adoption is increasing, comprehensive studies of continued usage intentions remain limited, especially in specialized academic contexts such as business schools. This limitation highlights the need to expand current technology adoption models to include constructs specifically relevant to generative AI, such as ethical judgment, perceived risk, and disciplinary norms. Examining how disciplinary differences

affect students' perceptions and adoption behaviors is crucial for establishing a conceptual foundation for the subsequent analysis.

2.2 Disciplinary distinctions between the social sciences and STEM fields

Generative AI adoption is often discussed as if it were one uniform phenomenon, but it isn't. Recent scholarship shows disciplinary epistemology shapes how students perceive, judge, and actually use these tools. Business and management studies, like other social sciences, lean on interpretive, context-sensitive, ethically attentive inquiry, and that orientation creates its own friction points when generative AI enters the picture. STEM fields work differently: their methodological traditions favor quantification, replicability, and statistical rigor, so AI adoption there tends to slot more naturally into existing analytical routines [8]. These aren't just stylistic differences - they shape the kinds of data students work with and what they expect in terms of accuracy, transparency, and academic integrity.

Worries about misconduct and slipping scholarly rigor hit hardest in fields where knowledge is built through interpretation rather than measured directly. Some scholars go further, arguing that generative AI muddies academic integrity standards in interpretive disciplines specifically: AI-generated text can blur who actually authored an idea, flatten interpretive diversity, and quietly reinforce whatever cultural narrative is already dominant [9]. Others warn that adopting these tools uncritically risks spreading a myth of technological neutrality, a problem particularly acute where cultural nuance and context are central to how knowledge gets produced in the first place [10]. STEM disciplines aren't immune either. AI there is often framed as a natural extension of computation, yet educators still worry about falsified data, AI doing the problem-solving instead of the student, and a gradual erosion of foundational skills [11].

Ethical judgments about AI-assisted writing and analysis, in other words, look quite different depending on which discipline you're standing in. Students frequently struggle to tell the difference between using AI as legitimate cognitive support and simply delegating the thinking to it, a distinction that gets especially blurry in qualitative work requiring interpretation, reflexivity, or engagement with lived experience [12]. Taken together, these patterns suggest generative AI adoption across higher education is anything but homogeneous: each field's epistemic norms, integrity standards, and methodological habits shape both what opportunities AI offers and what risks it brings with it. Table 1

presents the primary differences between social sciences and STEM disciplines regarding data types, integrity priorities, and challenges related to adopting generative AI.

Table 1. Epistemic and methodological distinctions between social sciences and STEM disciplines and their implications for generative AI adoption

Dimension	Social Sciences	STEM Disciplines
Data type	•Predominantly qualitative (interviews, ethnographies, case studies) •Context-dependent and shaped by cultural/social settings •Multiple valid interpretations possible	•Primarily quantitative (experiments, measurements, standardized instruments) •Objective, replicable, generalizable •Emphasis on statistical rigor
Integ-rity focus	• Transparency in interpretive choices and coding schemes •Ethical sensitivity to participants and researcher positionality •Risk of misrepresentation or oversimplification of lived experiences	•Accuracy, reproducibility, and error minimization •Misconduct tied to fabrication, falsification, or statistical manipulation •Integrity enforced through replication and methodological protocols
Gene-rative AI challenges	•Difficulty handling nuanced, context-rich qualitative data •Risk of flattening interpretive diversity or privileging dominant narratives •Outputs may lack cultural sensitivity or oversimplify complex narratives	•More straightforward adoption in pattern recognition and quantitative modeling •Integrity concerns focus on algorithmic accuracy and reproducibility •AI can accelerate analysis, simulation, and hypothesis testing if transparency is ensured

Source: created by the authors

This comparative framework offers critical context for understanding the unique tensions that business school students, as part of the broader social sciences, encounter between the perceived utility of

generative AI and concerns regarding ethical judgment, interpretive accuracy, and academic integrity.

The distinctions summarized in Table 1 provide a conceptual foundation for analyzing how business school students, whose studies integrate social science reasoning with applied analytical tasks, evaluate generative AI tools. These disciplinary nuances directly inform the factors discussed in the following section, which examines the specific constructs influencing generative AI adoption among social science students.

2.3 Factors influencing generative AI adoption by social science students

Building on the previously established disciplinary context, this section analyzes the primary psychological, social, and ethical factors influencing business school students' adoption of generative AI. Within a social science paradigm that emphasizes interpretive reasoning, ethical sensitivity, and contextual judgment, these students' adoption decisions are shaped by both traditional technology acceptance constructs and concerns about academic integrity, data privacy, and the reliability of AI-generated information [13].

Performance expectancy, effort expectancy, social influence, habit, study value, and risk perception all play significant roles in shaping adoption intentions. As demonstrated in Section 2.2, the interpretive nature of social science inquiry increases the importance of ethical judgment and perceived risk, rendering these constructs especially influential for business school students. The following subsections examine how these factors interact within a social science disciplinary context and influence students' willingness to adopt generative AI tools.

This study draws on the Unified Theory of Acceptance and Use of Technology 2 (UTAUT2) framework [14] to examine what shapes business school students' intentions to adopt generative AI tools in a social science context. UTAUT2 identifies performance expectancy, effort expectancy, social influence, facilitating conditions, hedonic motivation, price value, and habit as core drivers of technology adoption. Given the academic setting and the specific way generative AI gets used in coursework, we focus on the constructs most relevant to learning environments: performance expectancy, effort expectancy, social influence, habit, and perceived value.

Risk perception is also incorporated as a key moderating variable. This extension is theoretically justified, as generative AI introduces distinct ethical,

epistemic, and integrity-related uncertainties that the original UTAUT2 model does not address.

Within social science disciplines, where interpretive reasoning, ethical sensitivity, and contextual judgment are central, students' adoption decisions are especially influenced by concerns regarding academic integrity, data privacy, and the reliability of AI-generated content. Section 2.2 and Table 1 demonstrate that these disciplinary characteristics heighten perceived risk, thereby establishing it as a theoretically significant moderator of UTAUT2 relationships. Therefore, the following hypotheses are proposed.

2.3.1 Performance expectancy

Performance expectations (PER) refer to the extent to which individuals believe that using a particular system will enhance their job performance [14]. In the context of generative AI in higher education, students who perceive AI tools as improving efficiency, quality, or learning outcomes are more likely to adopt these technologies.

H1: Performance expectancy positively influences business school students' adoption of generative AI.

2.3.2 Effort expectancy

Effort expectancy (EFF) refers to the perceived ease of using a particular technology; for generative AI, this concerns how readily students can learn and apply these tools without significant cognitive or temporal burden [15].

Perceptions of AI-generated content tools as intuitive and functionally simplified increase the likelihood of adoption among students, thereby enhancing their intention to continue using these tools. Lower effort expectancy reduces psychological barriers to adoption and increases the probability of integrating generative AI into academic workflows. Furthermore, the technical performance of generative AI, especially regarding generation quality and context-awareness, directly affects both performance expectancy and effort expectancy. High-quality, relevant outputs minimize the need for user intervention and correction, which in turn reduces perceived cognitive costs. When students perceive AI tools as intuitive and cognitively manageable, their intention to adopt these technologies increases.

H2: Effort expectancy positively influences the adoption of generative AI among business school students.

2.3.3 Social influence

Social influence (SOC) is defined as the extent to which an individual's perception of others'

expectations regarding the use of a new method shapes their behavioral intention to adopt a technology [16]. Primary sources of social influence, such as peer recommendations, faculty endorsements, and institutional mandates, contribute to normative pressures that may facilitate or impede the adoption of generative AI tools. Empirical research consistently identifies social influence as a significant determinant of behavioral intention, with substantial evidence supporting its positive impact on technology adoption [6]. In academic settings, endorsements from peers, faculty, or industry leaders increase the likelihood that students will incorporate generative AI into their learning activities. Conversely, limited social endorsement or prevailing negative societal attitudes toward AI may substantially inhibit adoption, even when individual benefits are acknowledged.

H3: Social influence positively affects the adoption of generative AI among business school students

2.3.4 Study value

Study value (STV) refers to the perceived benefits and advantages that students attribute to using generative AI in academic contexts, including improved learning outcomes, greater productivity, and expanded research capabilities.

Students' perception of the "humanness" of AI tools is closely associated with their willingness to adopt these technologies. This perception may serve as either an incentive or a barrier, depending on students' attitudes toward novelty and ethical considerations [17]. The perceived usefulness of AI tools is a critical determinant of students' intention to use them, particularly when clear academic benefits are evident. Students are more likely to adopt generative AI when they view it as valuable for completing assignments efficiently, searching for information, and obtaining personalized responses rapidly [18].

H4: Study value positively influences business school students' adoption of generative AI.

2.3.5 Habit

Habit (HAB) is defined as the automaticity with which individuals use a technology, resulting from prior learning and repeated successful interactions, leading to ingrained usage patterns [19]. Within the context of generative AI, habit formation is crucial for sustained usage. What seems to matter most is repetition paired with success: when students keep getting good results from generative AI, it works its way into their regular study routines [7]. Early wins with the tool matter disproportionately - a few

positive experiences early on tend to snowball into a self-reinforcing cycle where use breeds benefit, which breeds more use. That cycle holds together as long as the tool keeps delivering: when students see it saving time or improving their work in measurable ways, they keep coming back to it.

H5: Habit positively influences the adoption of generative AI among business school students.

2.3.6 Risk perception

Risk perception (RIS) in the context of generative AI adoption by business school students refers to the subjective assessment of potential negative consequences, such as concerns about academic integrity, data privacy breaches, or the erosion of critical thinking skills, that could arise from the utilization of these tools [4].

H6: Risk perception positively influences business school students' adoption of generative AI.

2.3.7 Integration of risk perception as a moderator

Perceived risks significantly moderate the positive effects of performance expectancy and effort expectancy on adoption intention, even when students acknowledge the potential utility of generative AI. Concerns related to accuracy, transparency, privacy, ethical issues, and potential impacts on holistic competencies and career prospects often deter students from fully adopting generative AI tools, despite their recognized benefits [4]. Worries about AI's effect on future job prospects and on human values more broadly feed directly into this risk calculus, and that in turn can slow adoption in academic settings generally. The pattern seems fairly consistent: the more risk a student perceives, the less likely they are to adopt generative AI in any sustained way. That points to something worth digging into further - how students' willingness to adopt these tools is shaped not just by the tools themselves, but by the surrounding learning environment, their own sense of competence, and how instructors frame AI use.

Social influence, as noted, is a well-established driver of technology adoption generally. However, its effect diminishes when individuals perceive considerable risks related to generative AI tools. Negative social feedback or awareness of peers' adverse experiences can reinforce apprehensions, thereby reducing the impact of social endorsement. Even when peers or instructors advocate for AI use, elevated perceptions of risk may override these social pressures and impede adoption. While social influence can facilitate AI adoption, concerns about academic integrity or privacy may attenuate this

effect, leading to a more cautious approach despite peer recommendations [20].

Concerns about data accuracy and ethics can significantly dampen the link between social endorsement and behavioral intention - peer or faculty approval only goes so far if students still have doubts. So while positive social signals might get students to try generative AI in the first place, keeping them engaged long-term depends on whether their risk concerns get addressed [21]. Open, honest communication about a tool's limitations and safeguards seems to help here [22]. Understanding how social influence and risk perception interact, then, isn't just an academic exercise - it's central to encouraging responsible, sustained adoption.

Habitual use normally lowers perceived effort and smooths the path to acceptance, but that can break down fast. A single high-profile security breach or ethical controversy involving AI is often enough to disrupt an otherwise settled routine, prompting users to stop and reconsider whether they should keep engaging with the tool at all. This effect is amplified when students have already had a bad experience themselves, or have simply heard enough anecdotes about AI misuse, that accumulated unease can outweigh the sheer convenience of habitual use [5].

There's also what might be called an expectation-threat paradox: fear that AI will eventually replace human judgment or labor breeds a deeper uncertainty that undercuts habitual adoption altogether, particularly in cultures where professional identity is closely tied to job security [23]. Habit usually lowers the mental effort of using a tool, but a serious ethical concern or a new risk disclosure can override that, pushing users to break from routine and re-evaluate their relationship with generative AI even at the cost of disrupting their own workflow [24]. This dynamic highlights the vulnerability of habitual technology use to evolving risk landscapes, especially in advanced AI applications.

Given the ethical and epistemic complexities of generative AI, particularly in social science contexts, risk perception is expected to influence how students perceive the benefits and ease of use of AI tools [20], see the conceptual model in Figure 1.

H7: Risk perception moderates the relationship between performance expectancy and generative AI adoption, such that the relationship is weaker when perceived risk is high.

H8: Risk perception moderates the relationship between effort expectancy and generative AI adoption, such that the relationship is weaker when perceived risk is high.

H9: Risk perception moderates the relationship between social influence and generative AI adoption,

such that the relationship is weaker when perceived risk is high.

H10: Risk perception moderates the association between study value and generative AI adoption, such that this association weakens under high perceived risk.

H11: Risk perception moderates the association between habit and generative AI adoption, such that this association weakens when perceived risk is high.




Fig. 1: Conceptual model for testing business school students' adoption of GenAI in the study process

Source: created by the authors

3 Data and methods

The proposed hypotheses were evaluated using a survey instrument developed after an extensive review of scholarly literature to identify the primary factors influencing generative AI adoption (see Table A1 in the Appendix). Independent variables were measured on a 10-point Likert scale with response options from "disagree" to "strongly agree." Generative AI usage (USE) was measured using a five-point frequency scale with response options from "never" to "daily." Before full deployment, the questionnaire was pilot-tested in Autumn 2024 with a subset of the target population to evaluate item clarity and estimate completion time. Minor adjustments were implemented based on participant feedback. The questionnaire also collected demographic information, such as age, study level, and geographic location. The finalized questionnaire was distributed via academic networks associated with business schools in the Baltic Sea region.

Data collection occurred between January and March 2025, yielding 452 responses. After excluding 11 incomplete submissions, 441 valid cases remained for analysis. The age distribution was as follows: 43.3% of respondents were younger than 23 years, 32.7% were aged 23–27, 11.8% were aged 28–32, 5.2% were aged 33–37, 2.9% were aged 38–42, 2.7% were aged 43–47, 0.2% were aged 48–52, and 1.1%

were aged 53–57. Females represented 61.2% of the sample, while males accounted

for 38.8%. Most participants were enrolled in bachelor's programs (51.8%), with additional representation from master's (18.3%) and doctoral (9.5%) programs. A total of 20.4% did not report their study level. Geographically, most respondents resided in Latvia (59.2%), followed by Lithuania (12.7%), Denmark (7.9%), Estonia (7.5%), and other regions (12.7%). The convenience sample closely reflected the demographic characteristics of business school students at Baltic higher education institutions in terms of age, gender, and study level.

Respondents consistently agreed with the effort expectancy (EE1–EE3) items, with mean scores ranging from 8.01 to 8.66 and low standard deviations ($SD = 1.33$ – 1.41).

Habit (HA1–HA3) items demonstrated considerable variability ($SD = 1.61$ – 2.32). HA1 and HA3 exhibited high mean scores (7.49 and 8.93, respectively), while HA2 showed a moderate mean score (6.06). Performance expectancy (PE1–PE4) items were strongly endorsed, with mean scores ranging from 8.04 to 8.71 and moderate standard deviations ($SD = 1.37$ – 1.61). Risk perception (RP1–RP3) items produced mid-range mean scores (6.18–6.80), reflecting moderate concern among respondents.

Social influence (SI1–SI3) items elicited mixed but generally positive responses, with mean scores ranging from 6.49 to 7.59. Study value (SV1–SV4) items received high ratings (means = 7.58–8.65), reflecting a substantial perceived learning benefit of generative AI. The usage frequency items (US1, US2, US3, US4), assessed using a 5-point scale, indicated moderate to frequent utilization, with mean scores ranging from 3.03 to 4.03. Self-study (US4) had the highest mean score (4.03), suggesting

consistent reliance on generative AI for academic support. Table A2 in the Appendix provides item-level descriptive statistics for the survey constructs.

All item distributions demonstrated pronounced negative skewness and elevated excess kurtosis, reflecting a concentration of responses at the upper end of the scale. The Cramér–von Mises tests ($p < 0.001$ for all items) revealed significant deviations from normality. These findings support using non-parametric or robust estimation methods in subsequent analyses.

The conceptual model was evaluated using a two-stage analytical procedure that employed partial least squares structural equation modeling (PLS-SEM) with SmartPLS version 4.1.1.6 [25]. This method was selected based on the characteristics of the collected data and is supported by evidence of

effectiveness in comparable studies [17], [21], [24], [26], [27].

4 Results

4.1 Measurement model assessment results

Construct validity was evaluated through measurement model analysis, which included indicator loadings, internal consistency (as measured by Cronbach's alpha and composite reliability), and validity (convergent and discriminant).

The outer loadings demonstrate strong indicator reliability across all reflective constructs (see Table A3 and Table A4 in the Appendix). Effort expectancy, Habit, Performance expectancy, Social influence, Study value, and Usage all exhibit loadings predominantly above the recommended 0.70 threshold, with high t-statistics and narrow confidence intervals confirming statistical significance and measurement stability. Risk perception shows two powerful indicators (RP1, RP2) and one acceptable loading (RP3 = 0.657), which remains within acceptable bounds for exploratory PLS-SEM. Interaction terms modeled as single-item constructs load as expected at 1.000.

The results indicate that all constructs are measured with sufficient reliability, and no indicators require removal. Table 2 indicates that the USE construct exhibits strong internal consistency and convergent validity. Cronbach's alpha (0.793) and composite reliability ($\rho_a = 0.802$; $\rho_c = 0.794$) both exceed the recommended threshold of 0.70, thereby confirming the reliability of the construct. The average variance extracted (AVE = 0.523) also surpasses the 0.50 benchmark, demonstrating that the indicators account for more than half of the variance in the latent construct. Therefore, the USE construct satisfies established criteria for reliability and convergent validity within the framework of partial least squares structural equation modeling (PLS-SEM).

Table 2. Construct reliability and validity

Item	Cronbach's alpha	ρ_a	ρ_c	AVE
USE	0.793	0.802	0.794	0.523

Source: created by the authors

Discriminant validity was assessed using the Heterotrait–Monotrait (HTMT) ratio. All construct pairs met the recommended thresholds (see Table 3). All ratio values were substantially below the conservative threshold, indicating adequate separation between constructs. The HTMT ratios between USE and the interaction terms ranged from

0.191 to 0.429, indicating robust discriminant validity. Interaction terms involving Risk Perception demonstrate moderate associations with Habit ($r = 0.601$), Effort Expectancy ($r = 0.587$), Social Influence ($r = 0.560$), Study Value ($r = 0.647$), and Performance Expectancy ($r = 0.799$). All values remained within the acceptable upper limit. The highest observed value (HTMT = 0.799 for RIS $\times$ PER) was within acceptable bounds, confirming that no construct pair exhibited problematic overlap.

Overall, the HTMT results indicate satisfactory discriminant validity across the measurement model.

Table 3. Heterotrait-monotrait ratios

Item	USE	RIS $\times$ HAB	RIS $\times$ EFF	RIS $\times$ SOC	RIS $\times$ STV
RIS $\times$ HAB	0.429				
RIS $\times$ EFF	0.191	0.601			
RIS $\times$ SOC	0.261	0.593	0.354		
RIS $\times$ STV	0.359	0.738	0.466	0.440	
RIS $\times$ PER	0.357	0.799	0.587	0.560	0.647

Source: created by the authors

Variance inflation factor (vif) values for all predictors of USE fall within acceptable limits, indicating no critical multicollinearity issues in the structural model (see Table 4). All primary-effect constructs (EFF, HAB, PER, RIS, SOC, STV) show vif values between 1.410 and 3.299, well below the commonly accepted threshold of 5. Interaction terms exhibit slightly higher vif's, with RIS $\times$ HAB (4.658) and RIS $\times$ PER (4.330) being the largest, but still within acceptable bounds for models that include product indicators.

The remaining interaction terms range from 1.722 to 2.982, further supporting the absence of problematic collinearity. Variance inflation factor (vif) diagnostics confirm that multicollinearity is not present, thereby validating the inclusion of all predictors in the model.

Table 4. Inner model collinearity

Item	Vif
EFF $\rightarrow$ USE	2.032
HAB $\rightarrow$ USE	3.299
PER $\rightarrow$ USE	3.279
RIS $\rightarrow$ USE	1.877
RIS $\times$ EFF $\rightarrow$ USE	2.321
RIS $\times$ HAB $\rightarrow$ USE	4.658
RIS $\times$ PER $\rightarrow$ USE	4.330
RIS $\times$ SOC $\rightarrow$ USE	1.722
RIS $\times$ STV $\rightarrow$ USE	2.982
SOC $\rightarrow$ USE	1.410
STV $\rightarrow$ USE	2.981

Source: created by the authors

4.2 Structural model assessment results

The structural model demonstrated substantial explanatory power ($R^2 = 0.654$; adjusted $R^2 = 0.645$). These findings suggest that the predictors collectively explain approximately two-thirds of the variance in generative AI usage, indicating strong predictive accuracy in the context of behavioral research. Among the predictors, habit had the greatest effect on generative AI use ($\beta = 0.385$), followed by risk perception ($\beta = 0.202$) and social influence ($\beta = 0.172$).

Model fit indices provide additional evidence of the model's adequacy. The SRMR values of 0.060 for both the saturated and estimated models fall below the recommended threshold of 0.08, indicating a strong overall fit. The d_ULS (1.070 vs. 1.079) and d_G (0.390 vs. 0.389) values indicate minimal discrepancy between the saturated and estimated models, which supports the stability of the model estimation. The chi-square values for the saturated (894.821) and estimated (897.443) models are closely aligned. Additionally, the NFI values (0.876 and 0.875) exceed the commonly accepted threshold of 0.80, indicating an acceptable incremental fit (see Table 5). These results collectively demonstrate that the model exhibits strong explanatory power and satisfactory global fit.

Table 5. Model fit summary

Ratio	Saturated model	Estimated model
SRMR	0.060	0.060
d_ULS	1.070	1.079
d_G	0.390	0.389
Chi-square	894.821	897.443
NFI	0.876	0.875

Source: created by the authors

The bootstrapped structural paths indicate that several predictors exert significant direct effects on USE (see Table 6).

Performance expectancy ($\beta = 0.138$, $t = 1.966$) and effort expectancy ($\beta = 0.132$, $t = 2.093$) show positive, statistically significant effects, supporting H1 and H2. Social influence exhibits a significant positive effect on usage ($\beta = 0.172$, $t = 3.318$), confirming H3. In contrast, study value does not significantly predict usage ($\beta = 0.086$, $t = 1.153$), and thus H4 is not supported. Habit is identified as the strongest predictor of generative AI usage ($\beta = 0.385$, $t = 4.977$), providing robust support for H5. Risk perception also demonstrates a significant positive effect ($\beta = 0.202$, $t = 3.037$), supporting H6.

Regarding moderation effects, only the interaction between effort expectancy and risk perception is statistically significant ($\beta = 0.104$, $t =$

2.194), supporting H8. The remaining interaction terms, specifically $PER \times RIS$, $SOC \times RIS$, $STV \times RIS$, and $HAB \times RIS$, do not achieve statistical significance at the 95% confidence level, as their confidence intervals include zero. Therefore, hypotheses H7, H9, H10, and H11 are not supported.

Table 6. Bootstrapped paths statistics

Paths	Orig. Est.	Boot. Mean	Boot. SD	t-stat	H tests
PER→USE	0.138	0.141	0.070	1.966	H1 Yes
EFF→USE	0.132	0.130	0.063	2.093	H2 Yes
SOC→USE	0.172	0.173	0.052	3.318	H3 Yes
STV→USE	0.086	0.090	0.075	1.153	H4 No
HAB→USE	0.385	0.384	0.077	4.977	H5 Yes
RIS→USE	0.202	0.209	0.067	3.037	H6 Yes
PER*RIS→USE	-0.035	-0.031	0.057	0.615	H7 No
EFF*RIS→USE	0.104	0.097	0.048	2.194	H8 Yes
SOC*RIS→USE	0.051	0.051	0.041	1.242	H9 No
STV*RIS→USE	0.007	0.002	0.052	0.135	H10 No
HAB*RIS→USE	0.005	0.015	0.068	0.081	H11 No

Source: created by the authors

The conditional direct effects analysis shows how each predictor's impact on USE varies across levels of risk perception (see Table 7). When risk perception is high ($RIS = 1$), effort expectancy ($\beta = 0.236$, $t = 2.583$), habit ($\beta = 0.391$, $t = 3.230$), and social influence ($\beta = 0.224$, $t = 3.258$) each demonstrate significant positive effects on generative AI usage, as evidenced by confidence intervals that exclude zero. In contrast, performance expectancy and study value are not significant at high RIS , since their confidence intervals include zero. At low risk perception ($RIS = 0$), the pattern is similar, although effect sizes are slightly reduced.

Effort expectancy ($\beta = 0.132$, $t = 2.093$), habit ($\beta = 0.385$, $t = 4.977$), performance expectancy ($\beta = 0.138$, $t = 1.966$), and social influence ($\beta = 0.172$, $t = 3.318$) all demonstrate statistically significant positive effects. Study value again remains non-significant.

Overall, these results indicate that Habit and Social influence consistently predict usage across both low and high levels of risk perception. At the

same time, Effort expectancy becomes more influential when perceived risk is high. Performance expectancy predicts usage only under low-risk conditions, and Study value does not significantly influence usage at either risk level.

Table 7. Conditional direct effects

Paths	Orig. Est.	Boot. Mean	Boot. SD	2.5% CI	97.5% CI
EFF→USE RIS at one	0.236	0.227	0.091	0.072	0.430
HAB→USE RIS at one	0.391	0.399	0.121	0.152	0.629
PER→USE RIS at one	0.103	0.113	0.104	- 0.138	0.324
SOC→USE RIS at one	0.224	0.224	0.069	0.083	0.353
STV→USE RIS at one	0.093	0.092	0.107	- 0.120	0.301
EFF→USE RIS at zero	0.132	0.130	0.063	0.016	0.263
HAB→USE RIS at zero	0.385	0.384	0.077	0.240	0.543
PER→USE RIS at zero	0.138	0.145	0.070	0.005	0.285
SOC→USE RIS at zero	0.172	0.173	0.052	0.069	0.272
STV→USE RIS at zero	0.086	0.090	0.075	- 0.074	0.222

Source: created by the authors

The effect size analysis reveals that most predictors make small but meaningful contributions to explaining USE, as shown in Table 8.

Table 8. f-square

Item	f-square
EFF → USE	0.025
HAB → USE	0.130
PER → USE	0.017
RIS → USE	0.063
RIS x EFF→ USE	0.032
RIS x HAB→ USE	0.000
RIS x PER -> USE	0.003
RIS x SOC→ USE	0.009
RIS x STV → USE	0.000
SOC → USE	0.061
SVT → USE	0.007

Source: created by the authors

Habit ($f^2 = 0.130$) exhibits the largest effect among all predictors, approaching a medium effect size. This finding confirms habit's central role in shaping generative AI usage.

Risk perception ($f^2 = 0.063$) and social influence ($f^2 = 0.061$) exhibit modest yet significant effects,

indicating that both perceived risks and social context affect usage behaviour. Effort expectancy ($f^2 = 0.025$) exerts a relatively smaller influence. In contrast, performance expectancy ($f^2 = 0.017$) and study value ($f^2 = 0.007$) present minimal effects, suggesting limited additional explanatory power relative to other predictors.

Among the interaction terms, only $RIS \times EFF$ ($f^2 = 0.032$) shows a small effect, consistent with its significant moderation. All other interaction effects ($RIS \times HAB$, $RIS \times PER$, $RIS \times SOC$, $RIS \times STV$) exhibit negligible effect sizes ($f^2 = 0.000$ – 0.009), indicating minimal moderating influence on usage.

Overall, the pattern of effect sizes indicates that Habit is the primary predictor, whereas Risk perception, Social influence, and Effort expectancy make minor yet meaningful contributions. Other predictors and interactions provide limited additional explanatory value.

4.3 Estimated model predictive power

The PLS_Predict results demonstrate positive predictive relevance for all usage indicators, as shown by $Q^2_{predict}$ values ranging from 0.253 to 0.391 (see Table 9). Since all values are greater than zero, the model exhibits substantial out-of-sample predictive power.

Table 9. PLS_predict MV summary

Item	$Q^2_{predict}$	PLS-SEM RMSE	PLS-SEM MAE	LM RMSE	LM MAE	IA RMSE	IA MAE
US1	0.255	0.966	0.760	0.980	0.768	1.135	0.869
US2	0.253	0.800	0.619	0.825	0.629	0.926	0.684
US3	0.257	0.711	0.545	0.718	0.549	0.824	0.685
US4	0.391	0.760	0.569	0.788	0.577	1.010	0.740

Source: created by the authors

US4 exhibits the highest predictive relevance ($Q^2_{predict} = 0.391$), followed by US3, US1, and US2. Partial least squares structural equation modeling (PLS-SEM) provides superior or equivalent predictive performance compared to the linear benchmark (LM) across all items. For each indicator, the root mean square error (RMSE) and mean absolute error (MAE) values for PLS-SEM are consistently lower than or similar to those of LM.

This finding demonstrates improved or comparable predictive accuracy. The improvement is most evident for USS and UCL, where PLS-SEM produces notably lower error metrics. Compared with the naïve indicator average (IA) benchmark, the PLS-SEM model yields substantially lower RMSE and MAE across all indicators, confirming meaningful predictive value beyond simple mean-based prediction.

The MV-level PLS-Predict results demonstrate that the model possesses moderate predictive power. PLS-SEM consistently outperforms both LM and IA benchmarks, supporting the model's robustness in predicting generative AI usage behaviors.

A comparison of PLS-SEM prediction errors with the indicator-average (IA) benchmark shows that the PLS-SEM model achieves substantially greater predictive accuracy across all usage indicators. For each item (US1, US2, US3, US4), PLS loss values are consistently lower than the corresponding IA loss values. This finding indicates reduced prediction error with the structural model compared with mean-based prediction. The average loss differences are negative for all indicators (−0.174 to −0.398), as presented in Table 10, confirming that PLS-SEM outperforms the IA benchmark. The observed differences are statistically significant, as indicated by t-values between 4.747 and 6.236 and p-values less than 0.001 for all items.

Table 10. PLS-SEM vs indicator average statistics

Item	PLS loss	IA loss	Average loss difference	t value	p value
US1	0.960	1.289	-0.329	6.008	0.000
US2	0.640	0.857	-0.216	4.972	0.000
US3	0.505	0.680	-0.174	4.747	0.000
US4	0.621	1.020	-0.398	6.236	0.000

Source: created by the authors

The most significant improvement is observed for USS, where PLS-SEM reduces prediction error by nearly 0.40 compared to the IA approach. Overall, the results demonstrate that the PLS-SEM model offers meaningful and statistically significant predictive improvements over the naïve indicator-average method, reinforcing the model's robustness and practical predictive utility.

The comparison between PLS-SEM and the linear model (LM) shows (see Table 11) that PLS-SEM yields slightly lower prediction error across all usage indicators (US1, US2, US3, US4), as reflected by consistently negative average loss differences. However, these improvements do not reach statistical significance, as indicated by t-values ranging from 0.567 to 1.754 and all p-values greater than 0.05.

Table 11. PLS-SEM vs. Linear model (LM)

Item	PLS loss	LM loss	Average loss diff.	t-value	p-value
UCL	0.640	0.680	-0.040	1.754	0.080
UHW	0.505	0.515	-0.010	0.567	0.571
USS	0.621	0.665	-0.044	1.360	0.174
UTW	0.960	0.988	-0.028	0.681	0.496

Source: created by the authors

The largest reduction in prediction error is observed for US2 (difference = −0.040); however, this effect is marginal ($t = 1.754$, $p = 0.080$). These results suggest that although PLS-SEM exhibits performance comparable to or slightly better than the LM benchmark, the differences are insufficient to establish superior predictive performance.

5 Discussions

This study offers key insights into the adoption of generative AI tools among business school students and investigates the impact of perceived risk on their behavior. Analysis using the UTAUT2 framework shows that habit is the strongest predictor of generative AI use. Once students integrate AI tools into their study routines, continued use tends to become automatic. This observation aligns with previous research suggesting that repeated exposure and positive reinforcement promote sustained engagement with digital tools in educational settings [7], [19]. The pronounced effect of habit further demonstrates the rapid normalization of generative AI within students' academic workflows, especially for self-study, summarization, and exploratory writing, which exhibited the highest usage frequencies in the dataset.

Another key finding is the significant direct effect of risk perception on generative AI use. Unlike traditional technologies examined in previous UTAUT2 research, generative AI introduces unique uncertainties related to academic integrity, data accuracy, and ethical boundaries. The positive coefficient for risk perception indicates that students who recognize these risks may, contrary to initial assumptions, become more deliberate and engaged users. This ultimately stems from heightened critical awareness during interactions with AI tools.

This interpretation aligns with recent scholarship suggesting that risk-aware students use AI more strategically, balancing innovation with caution [3], [5]. It also reflects the disciplinary context of business schools, where students are trained to evaluate risk-reward trade-offs and may therefore integrate AI tools while simultaneously monitoring their limitations.

Social influence emerged as a significant predictor of generative AI adoption. Peers, instructors, and institutional norms all seem to play a real part in shaping how students approach these tools. That tracks with earlier research showing that social endorsement helps legitimize new technologies in academic settings [6], [16]. Given the ongoing debate over academic integrity and AI-assisted writing, positive social cues may do more than just encourage use - they can help students work through uncertainty and feel more confident about using AI responsibly. This seems especially true in social science disciplines, where interpretive norms and ethical expectations vary considerably, and students often look to instructors for guidance on what counts as acceptable AI use.

This dynamic is especially pertinent within social science disciplines, where interpretive norms and ethical expectations differ considerably, and students frequently seek guidance from instructors regarding acceptable AI practices [9], [10].

Although students rated performance expectancy, effort expectancy, and study value highly, these factors did not exert effects as strong as habit or risk perception. This indicates that recognition of generative AI's usefulness, ease of use, and academic value is insufficient to explain actual usage. Instead, behavioral patterns (habit) and ethical-epistemic considerations (risk perception) play a more decisive role. These results support arguments that generative AI adoption cannot be fully explained by traditional utilitarian models and must include constructs that address the moral, interpretive, and disciplinary complexities of AI-mediated learning [8], [12].

The results indicate that risk perception only weakly moderates the relationships between UTAUT2 predictors and generative AI usage. While risk remains an important factor, it does not substantially alter the influence of performance expectancy, effort expectancy, social influence, study value, or habit on usage. One explanation is that students perceive generative AI risks as manageable or context-dependent, rather than as barriers that fundamentally reshape adoption decisions.

Alternatively, students may have established stable usage patterns, as indicated by the strong habit effect, which appear less susceptible to risk-based fluctuations. This observation aligns with recent research showing that students often continue using AI tools despite concerns about accuracy, ethics, or academic integrity [4], [28].

The findings underscore the importance of disciplinary context in influencing generative AI adoption. As detailed in Section 2.2 and presented in Table 1, social science students engage with

epistemic traditions that prioritize interpretation, reflexivity, and ethical sensitivity. These disciplinary norms likely contribute to the moderate levels of risk perception observed in the sample and may explain why risk perception plays a more central role than in typical UTAUT2 applications. The findings therefore support calls for discipline-specific frameworks for AI integration in higher education [10], [11], [29], particularly in fields where qualitative reasoning and academic integrity concerns are paramount.

Overall, this study extends the theoretical understanding of generative AI adoption by demonstrating that habit and risk perception are central drivers of use. Traditional UTAUT2 predictors serve a supplementary function. The findings highlight the need to include ethical and disciplinary considerations in technology adoption models as generative AI becomes more common in academic settings.

6 Practical implications

This study's findings offer practical recommendations for business schools seeking to integrate generative AI into teaching, learning, and assessment. Habit's strong pull here suggests something useful: students settle into AI-supported study routines quickly once early experiences go well. Business schools can use that. Introducing structured, low-stakes opportunities for AI engagement early in the curriculum, guided exercises, reflective prompts, and supervised experimentation gives students a chance to build good habits before bad ones take hold. Scaffolding of this kind matters precisely because it steers habitual use toward academic integrity rather than leaving it to develop unsupervised.

Risk perception matters just as much, but differently. Students aren't adopting AI tools blindly; they're weighing concerns about accuracy, ethics, and academic integrity as they go. That means business schools need to be upfront: clear communication about acceptable AI use, data privacy, and where these models fall short. Institutional guidelines, discipline-specific case studies, and direct discussion of authorship and accountability all help students manage that uncertainty and make better-informed choices. Folding AI literacy modules, covering bias, hallucination, and responsible use, into the core curriculum would reinforce this further.

Social influence points to a similar lesson about faculty and peers. Faculty development programs matter here, equipping instructors to model responsible AI use, design assessments that hold up

against AI intervention, and keep expectations consistent from course to course. Peer-led workshops, student ambassador programs, and collaborative learning activities can reinforce the right norms too, making clear what appropriate AI use actually looks like. Meanwhile, the high ratings students gave performance expectancy, effort expectancy, and study value suggest they already see generative AI as both useful and easy to access, a kind of readiness business schools could tap into more directly by weaving AI into experiential learning, analytics coursework, and applied projects. Doing so gives students room to build not just technical proficiency, but the strategic and ethical judgment that AI-driven workplaces will demand of them.

7 Limitations and future research

This study has some real limitations worth flagging. The sample was a convenience sample drawn from business schools in the Baltic Sea region, which limits how far the results can be generalized: cultural, institutional, and regulatory differences elsewhere could shift both risk perceptions and adoption patterns in ways this study can't capture. Cross-national or multi-institutional samples would help future work account for that variability.

There's also the usual caveat with self-reported measures: social desirability bias is a real risk here, especially given how sensitive academic integrity and AI-assisted work have become as topics. Pulling in behavioral data, system logs, assignment metadata, and observational methods, could give a more objective read on actual usage than self-report alone.

The cross-sectional design is another constraint: it can't establish causality, nor can it track how perceptions shift over time. Given how fast generative AI itself is changing, longitudinal work would help clarify how habits form, how risk perception evolves with more exposure, and how institutional policy shapes adoption as it unfolds.

A fourth limitation concerns how usage was measured. This study looked at actual usage rather than behavioral intention, which is a step forward, but it measured usage through frequency alone: how often students turned to generative AI for coursework, homework, or self-study. That doesn't capture depth. A student who uses AI for quick, surface-level help and one who uses it for sustained conceptual exploration or iterative reasoning would look identical on a frequency scale, even though their engagement is quite different. That distinction is only going to matter more as generative AI becomes further embedded in how students actually learn.

Therefore, the present study offers only a partial perspective on qualitative differences in AI engagement. Future studies should include measures that assess the depth, intentionality, and cognitive engagement involved in AI use.

While this study extends the UTAUT2 by incorporating risk perception as a moderator, additional theoretically relevant constructs, including trust in AI, ethical judgment, AI literacy, and perceived academic pressure, could further enhance the model. Future research should investigate these variables or explore alternative frameworks that more comprehensively address the ethical and epistemic complexities of generative AI in the social sciences.

Future research should take this further with qualitative or mixed-methods approaches, looking at reflective use, critical evaluation of AI output, and AI literacy as their own dimensions of adoption, distinct from simple frequency. Cross-cultural comparative work would also help establish whether risk perception moderates adoption consistently across educational systems with different integrity norms, regulatory environments, and cultural attitudes toward AI. Longitudinal studies, similarly, could show how habitual use shifts as institutional policy and the technology itself keep evolving. Longitudinal studies may also clarify how habitual use changes as institutional policies and technology advance.

8 Conclusion

This study advances understanding of generative AI adoption among business school students by integrating disciplinary context into an extended UTAUT2 framework. The findings suggest that habit is the principal determinant of generative AI use, highlighting the importance of early experiences and routine formation in promoting sustained engagement. Risk perception is also a key factor, directly affecting adoption and reflecting the ethical and epistemic uncertainties associated with AI use in social science disciplines. Furthermore, social influence also plays a role here, underscoring how much faculty guidance and peer norms do to legitimize responsible AI use.

Traditional UTAUT2 predictors, performance expectancy, effort expectancy, and study value, tell a different story: students rate them highly, but they carry less weight in actually predicting usage. That points to something worth taking seriously: generative AI adoption seems to hinge more on behavioral patterns and ethical considerations than on straightforward utility. Which means AI integration strategies probably need to be discipline-specific,

particularly in fields where interpretation, reflexivity, and academic integrity sit at the core of the work.

This study advances theory by emphasizing the importance of including risk-related and ethical factors in technology adoption models. It also offers practical guidance for business schools navigating what generative AI both offers and threatens. As AI continues to reshape higher education, understanding how students actually adopt and engage with it will be essential to building curricula that support innovation without giving up on academic integrity.

Declaration of Generative AI and AI-assisted Technologies in the Writing Process

During the preparation of this work, the author(s) utilized artificial intelligence (AI) technologies solely for the purpose of linguistic refinement and grammatical correction. Following the application of these tools, the author(s) comprehensively reviewed, verified, and edited the resulting text. The author(s) maintain full accountability and responsibility for the accuracy, integrity, and final content of the publication.

References:

- [1] P. Shrivastava, "Understanding acceptance and resistance toward generative AI technologies: A multi-theoretical framework integrating functional, risk, and sociolegal factors," *Frontiers in Artificial Intelligence*, vol. 8, pp. 1565927, 2025. <https://doi.org/10.3389/frai.2025.1565927>
- [2] A. E. Sousa and P. Cardoso, "Use of generative AI by higher education students," *Electronics*, vol. 14, no. 7, pp. 1258, 2025. <https://doi.org/10.3390/electronics14071258>
- [3] L. Hughes, T. Malik, S. Dettmer, A. S. Al-Busaidi, and Y. K. Dwivedi, "Reimagining higher education: Navigating the challenges of generative AI adoption," *Information Systems Frontiers*, 2025. <https://doi.org/10.1007/s10796-025-10582-6>
- [4] C. K. Y. Chan and W. Hu, "Students' voices on generative AI: Perceptions, benefits, and challenges in higher education," *International Journal of Educational Technology in Higher Education*, vol. 20, no. 43, pp. 1-18, 2023. <https://doi.org/10.1186/s41239-023-00411-8>
- [5] S. Balaskas, V. Tsiantos, S. Chatzifotiou, and M. Rigou, "Determinants of ChatGPT adoption intention in higher education: Expanding on TAM with the mediating roles of trust and risk," *Information*, vol. 16, no. 2, p. 82, 2025. <https://doi.org/10.3390/info16020082>
- [6] P. Holzmann, P. Gregori, and E. J. Schwarz, "Students' little helper: Investigating continuous-use determinants of generative AI and ethical judgment," *Education and Information Technologies*, vol. 30, pp. 24991–25011, 2025. <https://doi.org/10.1007/s10639-025-13708-0>
- [7] M. Faraon, K. Rönkkö, M. Milrad, and E. Tsui, "International perspectives on artificial intelligence in higher education: An explorative study of students' intention to use ChatGPT across the Nordic countries and the USA," *Education and Information Technologies*, vol. 30, pp. 17835–17880, 2025. <https://doi.org/10.1007/s10639-025-13492-x>
- [8] K. B. Ooi, G. Tan, M. Al-Emran, M. Al-Sharafi, A. Capatina, A. Chakraborty, Y. Dwivedi, T.-L. Huang, A. Kar, V.-H. Lee, X.-M. Loh, A. Micu, P. Mikalef, E. Mogaji, N. Pandey, R. Raman, N. Rana, P. Sarker, A. Sharma, C. Teng, S. Wamba, L.-W. Wong, "The potential of generative artificial intelligence across disciplines: Perspectives and future directions," *Journal of Computer Information Systems*, vol. 65, no. 1, pp. 76–107, 2025. <https://doi.org/10.1080/08874417.2023.2261010>
- [9] D. R. Cotton, P. A. Cotton, and J. R. Shipway, "Chatting and cheating: Ensuring academic integrity in the era of ChatGPT," *Innovations in Education and Teaching International*, vol. 61, no. 2, pp. 228–239, 2024. <https://doi.org/10.1080/14703297.2023.2190148>
- [10] J. Rudolph, F. Ismail, S. Tan, and P. Seah, "Don't believe the hype: AI myths and the need for a critical approach in higher education," *Journal of Applied Learning and Teaching*, vol. 8, no. 1, pp. 6–27, 2025. <https://search.informit.org/doi/abs/10.3316/informit.T2025091200002501930224078> [Access Date: 01/01/2026]
- [11] N. Werse and J. Smith, "Nothing new under the sun: Generative AI and educator anxiety over academic dishonesty," *Impacting Education: Journal on Transforming Professional Practice*, vol. 10, no. 1, pp. 90–95, 2025. <https://doi.org/10.5195/ie.2025.484>

- [12] I. J. Yusof, Z. M. Ashari, L. H. Ismail, and M. Panadi, “‘We don’t plagiarise, we parrot’: Cognitive load and ethical perceptions in higher education written assessment,” *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, vol. 5, no. 4, p. 100254, 2025.
<https://doi.org/10.1016/j.tbench.2025.100254>
- [13] X. Bai and L. F. Yang, “Research on the influencing factors of generative artificial intelligence usage intent in post-secondary education: An empirical analysis based on the AIDUA extended model,” *Frontiers in Psychology*, vol. 16, pp. 1644209, 2025.
<https://doi.org/10.3389/fpsyg.2025.1644209>
- [14] V. Venkatesh, “Adoption and use of AI tools: A research agenda grounded in UTAUT,” *Annals of Operations Research*, vol. 308, pp. 641–652, 2021.
<https://doi.org/10.1007/s10479-020-03918-9>
- [15] A. Gulati, H. Saini, S. Singh, and V. Kumar, “Enhancing learning potential: Investigating marketing students’ behavioral intention to adopt ChatGPT,” *Marketing Education Review*, vol. 34, no. 3, pp. 201–234, 2024.
<https://doi.org/10.1080/10528008.2023.2300139>
- [16] C. H. Sundkvist and E. M. Kulset, “Teaching accounting in the era of ChatGPT – The student perspective,” *Journal of Accounting Education*, vol. 69, p. 100932, 2024.
<https://doi.org/10.1016/j.jaccedu.2024.100932>
- [17] Y. He, X. Liu, and X. Liao, “Do artificial intelligence ethical anxiety, perceived ethical risks and ethical awareness affect college students’ use of generative artificial intelligence products? Research from an ethical perspective,” *International Business Research*, vol. 18, no. 5, pp. 15–23, 2025.
<https://doi.org/10.5539/ibr.v18n5p15>
- [18] R. Lenart, B. A. Sypniewska, J. Chen, and K. Janowski, “Antecedents of students’ behavioural intention to use generative artificial intelligence: Quantitative research,” *Entrepreneurial Business and Economics Review*, vol. 13, no. 1, pp. 7–21, 2025.
<https://doi.org/10.15678/EBER.2025.130101>
- [19] K. Lavidas, I. Voulgari, S. Papadakis, S. Athanassopoulos, A. Anastasiou, A. Filippidi, V. Komis, N. Karacapilidis, “Determinants of humanities and social sciences students’ intentions to use artificial intelligence applications for academic purposes,” *Information*, vol. 15, no. 6, p. 314, 2024.
<https://doi.org/10.3390/info15060314>
- [20] N. Mat, S. Yusuf, N. Ramli, M. Mamat, M. Setapa, Y. Yusoff, A. Shari, “The relationship between performance expectancy, effort expectancy, and perceived risk on AI acceptance in higher education,” *International Journal of Research and Innovation in Social Science*, vol. IX, issue VIII, pp. 1017–1026, 2025.
<https://dx.doi.org/10.47772/IJRIS.2025.908000087>
- [21] Y. Tang and L. Su, “Graduate education in China meets AI: Key factors for adopting AI-generated content tools,” *Libri*, vol. 75, issue 1, pp. 81–96, 2024.
<https://doi.org/10.1515/libri-2024-0079>
- [22] A. Johri, A. Hingle, and J. Schleiss, “Misconceptions, pragmatism, and value tensions: Evaluating students’ understanding and perception of generative AI for education,” in *Proc. IEEE Frontiers in Education Conference*, Washington, DC, 2024, pp. 1–9.
<https://doi.org/10.1109/FIE61694.2024.10893017>
- [23] T. Yu, Y. Tian, Y. Chen, Y. Huang, Y. Pan, W. Jang, “How do ethical factors affect user trust and adoption intentions of AI-generated content tools? Evidence from a risk-trust perspective,” *Systems*, vol. 13, no. 6, p. 461, 2025.
<https://doi.org/10.3390/systems13060461>
- [24] C. Lee, J. Kim, J. S. Lim, and D. Shin, “Generative AI risks and resilience: How users adapt to hallucination and privacy challenges,” *Telematics and Informatics Reports*, vol. 19, pp. 100221, 2025.
<https://doi.org/10.1016/j.teler.2025.100221>
- [25] C. M. Ringle, S. Wende, and J.-M. Becker, *SmartPLS 4*. Bönningstedt: SmartPLS, 2024. [Online]. Available: <https://www.smartpls.com> [Access Data: 01/01/2026]
- [26] A. Strzelecki, “Students’ acceptance of ChatGPT in higher education: An extended unified theory of acceptance and use of technology,” *Innovative Higher Education*, vol. 49, pp. 223–245, 2023.
<https://doi.org/10.1007/s10755-023-09686-1>
- [27] M. Soliman, R. A. Ali, J. Khalid, I. Mahmud, and W. B. Ali, “Modelling continuous

intention to use generative artificial intelligence as an educational tool among university students: Findings from PLS-SEM and ANN,” *Journal of Computers in Education*, vol.12, pp. 897-928, 2024. <https://doi.org/10.1007/s40692-024-00333-y>

- [28] L. Reiter, M. Jörling, C. Fuchs, and R. Böhm, “Student (mis)use of generative AI tools for university-related tasks,” *International Journal of Human-Computer Interaction*, vol. 41, issue 19, pp. 12390–12403, 2025. <https://doi.org/10.1080/10447318.2025.2462083>

- [29] C. Stöhr, A. W. Ou, and H. Malmström, “Perceptions and usage of AI chatbots among students in higher education across genders, academic levels and fields of study,” *Computers and Education: Artificial Intelligence*, vol. 7, p. 100259, 2024. <https://doi.org/10.1016/j.caeai.2024.100259>

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

The authors equally contributed to the present research, at all stages from the formulation of the problem to the final findings and solution.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

This research was funded by Project No.5.2.1.1.i.0/2/24/I/CFLA/007 "Internal and external consolidation of the University of Latvia", Project No. LU-BAPA-2024/1–0052.

Conflict of Interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.

Creative Commons Attribution License 4.0 (Attribution 4.0 International, CC BY 4.0)

This article is published under the terms of the Creative Commons Attribution License 4.0

APPENDIX

Table A1. Measurement model variables

Construct	Items
PER	PE1: GAI is useful for my studies/research
	PE2: GAI enhances the quality of my studies/research
	PE3: GAI allows me to accomplish study/research activities more quickly
	PE4: GAI increases my study/research productivity
EFF	EE1: Learning how to use GAI is easy for me
	EE2: My interaction with GAI is clear and understandable
	EE3: GAI is easy to use
SOC	SI1: My fellow students/colleagues who influence my behavior recommend that I use GAI
	SI2: My lecturers, who are important to me, recommend that I use GAI
	SI3: Employers whose views I value recommend that I use GAI
	SI4: Public opinion leaders, who are important to me, recommend that I use GAI
STV	SV1: Application of GAI in the study process is worth more than the time and effort given to it
	SV2: GAI allows me to quickly and easily share my knowledge with others
	SV3: GAI allows me to decide on the pace of my own study
	SV4: GAI allows me to increase my knowledge and to control my success
HAB	HA1: The use of GAI has become a habit for me
	HA2: I am addicted to using GAI to accomplish my study/research tasks
	HA3: I intend to continue using GAI in the future
RIS	RI1: The information obtained from GAI is correct and accurate
	RI2: The information obtained from GAI is reliable
	RIS3: Verification mechanisms to identify the use of GAI are sufficient
USE	US1: I am using GAI to write my thesis (undergraduate/master's/PhD)
	US2: I am using GAI while in classes at the university
	US3: I use GAI for doing homework/preparing projects/and writing papers
	US4: I use GAI for self-study, gaining knowledge and skills on an interesting topic related to my studies/research

Source: created by the authors

Table A2. Descriptive statistics of survey responses by items

Item	Mean	Median	Observation		Standard deviation	Excess kurtosis	Skewness	Cramér-von Mises	
			min	max				test stat.	p-value
EE1	8.009	8.000	1.000	10.000	1.326	3.167	-1.116	3.270	0.000
EE2	8.095	8.000	1.000	10.000	1.406	4.738	-1.500	2.835	0.000
EE3	8.664	9.000	1.000	10.000	1.340	5.057	-1.780	4.199	0.000
HA1	7.492	8.000	1.000	10.000	2.010	2.362	-1.526	4.682	0.000
HA2	6.063	7.000	1.000	10.000	2.318	-0.106	-0.774	3.442	0.000
HA3	8.925	9.000	1.000	10.000	1.607	6.689	-2.325	7.987	0.000
PE1	8.281	8.000	1.000	10.000	1.373	5.701	-1.656	3.507	0.000
PE2	8.043	8.000	1.000	10.000	1.518	5.203	-1.752	3.933	0.000
PE3	8.705	9.000	1.000	10.000	1.547	9.058	-2.609	5.974	0.000
PE4	8.548	9.000	1.000	10.000	1.614	5.647	-2.049	4.647	0.000
RP1	6.735	7.000	1.000	10.000	1.720	1.483	-1.330	5.918	0.000
RP2	6.797	7.000	1.000	10.000	1.654	1.671	-1.356	5.700	0.000
RP3	6.177	6.000	1.000	10.000	1.553	1.942	-0.915	3.870	0.000
SI1	7.588	8.000	1.000	10.000	1.717	3.892	-1.634	4.749	0.000
SI2	6.487	7.000	1.000	10.000	1.970	1.138	-1.214	4.192	0.000
SI3	6.510	7.000	1.000	10.000	1.907	1.639	-1.055	2.694	0.000
SV1	7.707	8.000	1.000	10.000	1.696	3.013	-1.599	5.588	0.000
SV2	7.580	8.000	1.000	10.000	1.737	3.897	-1.817	4.965	0.000
SV3	8.166	9.000	1.000	10.000	1.918	3.553	-1.813	5.333	0.000
SV4	8.649	9.000	1.000	10.000	1.516	6.096	-2.017	4.144	0.000
US1	3.032	3.000	1.000	5.000	1.132	-0.604	-0.673	5.172	0.000
US2	3.125	3.000	1.000	5.000	0.924	0.264	-0.614	5.302	0.000
US3	3.508	4.000	1.000	5.000	0.822	1.065	-0.664	5.937	0.000
US4	4.032	4.000	1.000	5.000	1.007	1.161	-1.132	4.778	0.000

Source: created by the authors

Table A3 Cross loadings

Item	EFF	HAB	PER	RIS	SOC	STV	USE	RxHAB	RxEFF	RxSOC	RxSTV	RxPER
EE1	0.636	0.311	0.323	0.179	0.113	0.225	0.297	-0.212	-0.445	-0.095	-0.177	-0.216
EE2	0.704	0.382	0.408	0.326	0.187	0.308	0.329	-0.316	-0.478	-0.210	-0.220	-0.333
EE3	0.996	0.458	0.489	0.381	0.298	0.453	0.465	-0.282	-0.517	-0.185	-0.221	-0.268
HA1	0.410	0.806	0.518	0.492	0.402	0.455	0.593	-0.552	-0.307	-0.322	-0.384	-0.373
HA2	0.288	0.801	0.432	0.530	0.380	0.526	0.590	-0.416	-0.146	-0.276	-0.324	-0.264
HA3	0.457	0.878	0.673	0.490	0.374	0.622	0.647	-0.579	-0.341	-0.314	-0.433	-0.486
PE1	0.367	0.544	0.731	0.196	0.260	0.418	0.452	-0.417	-0.279	-0.287	-0.311	-0.461
PE2	0.374	0.616	0.843	0.314	0.279	0.539	0.521	-0.435	-0.262	-0.269	-0.321	-0.502
PE3	0.426	0.615	0.855	0.353	0.309	0.527	0.528	-0.471	-0.293	-0.318	-0.337	-0.548
PE4	0.482	0.610	0.963	0.425	0.393	0.660	0.595	-0.443	-0.297	-0.319	-0.361	-0.543
RP1	0.393	0.583	0.397	0.980	0.368	0.494	0.574	-0.498	-0.331	-0.326	-0.469	-0.350
RP2	0.380	0.590	0.409	0.963	0.357	0.513	0.564	-0.524	-0.319	-0.347	-0.489	-0.362
RP3	0.243	0.392	0.204	0.657	0.230	0.372	0.385	-0.308	-0.239	-0.258	-0.392	-0.202
SI1	0.263	0.403	0.387	0.222	0.742	0.267	0.374	-0.302	-0.181	-0.357	-0.224	-0.291
SI2	0.267	0.389	0.301	0.380	0.912	0.356	0.460	-0.247	-0.153	-0.296	-0.187	-0.208
SI3	0.208	0.304	0.274	0.265	0.748	0.301	0.377	-0.226	-0.119	-0.282	-0.264	-0.208
SV1	0.374	0.558	0.589	0.423	0.259	0.688	0.429	-0.331	-0.253	-0.125	-0.361	-0.320
SV2	0.343	0.548	0.494	0.493	0.366	0.742	0.463	-0.348	-0.227	-0.208	-0.459	-0.261
SV3	0.437	0.631	0.540	0.534	0.357	0.887	0.553	-0.385	-0.237	-0.244	-0.472	-0.255
SV4	0.402	0.564	0.620	0.391	0.324	0.919	0.573	-0.313	-0.125	-0.217	-0.416	-0.274
US1	0.225	0.451	0.394	0.447	0.403	0.328	0.647	-0.240	-0.081	-0.183	-0.159	-0.195
US2	0.338	0.476	0.438	0.340	0.391	0.373	0.659	-0.287	-0.102	-0.199	-0.226	-0.245
US3	0.341	0.530	0.416	0.366	0.289	0.384	0.669	-0.331	-0.186	-0.117	-0.294	-0.281
US4	0.396	0.601	0.485	0.484	0.344	0.627	0.819	-0.343	-0.165	-0.232	-0.326	-0.278
RxEFF	-0.515	-0.317	-0.323	-0.335	-0.187	-0.206	-0.192	0.601	1.000	0.354	0.466	0.587
RxHAB	-0.292	-0.616	-0.492	-0.522	-0.314	-0.390	-0.430	1.000	0.601	0.593	0.738	0.799
RxPER	-0.279	-0.461	-0.587	-0.364	-0.282	-0.306	-0.357	0.799	0.587	0.560	0.647	1.000
RxSOC	-0.198	-0.360	-0.345	-0.345	-0.377	-0.252	-0.263	0.593	0.354	1.000	0.440	0.560
RxSTV	-0.224	-0.460	-0.385	-0.492	-0.259	-0.498	-0.363	0.738	0.466	0.440	1.000	0.647

Source: created by the authors

Table A4 Outer loadings confidence intervals

Item	Original sample (O)	Sample mean (M)	Stand. dev. (STDEV)	T statistics (O/STDEV)	2.5% CI	97.5% CI
EE1 → EFF	0.636	0.621	0.096	6.596	0.404	0.783
EE2 → EFF	0.704	0.692	0.099	7.146	0.476	0.856
EE3 → EFF	0.996	0.980	0.020	49.215	0.928	1.000
HA1 → HAB	0.806	0.802	0.048	16.628	0.696	0.884
HA2 → HAB	0.801	0.800	0.041	19.685	0.713	0.871
HA3 → HAB	0.878	0.872	0.040	22.199	0.786	0.940
PE1 → PER	0.731	0.719	0.067	10.991	0.568	0.828
PE2 → PER	0.843	0.834	0.047	17.936	0.727	0.911
PE3 → PER	0.855	0.844	0.045	18.910	0.747	0.921
PE4 → PER	0.963	0.955	0.025	38.856	0.896	0.991
RIS x EFF → RIS x EFF	1.000	1.000	0.000	n/a	1.000	1.000
RIS x HAB → RIS x HAB	1.000	1.000	0.000	n/a	1.000	1.000
RIS x PER → RIS x PER	1.000	1.000	0.000	n/a	1.000	1.000
RIS x SOC → RIS x SOC	1.000	1.000	0.000	n/a	1.000	1.000
RIS x STV → RIS x STV	1.000	1.000	0.000	n/a	1.000	1.000
RP1 → RIS	0.980	0.972	0.022	45.472	0.916	0.998
RP2 → RIS	0.963	0.954	0.028	34.520	0.889	0.994
RP3 → RIS	0.657	0.652	0.079	8.273	0.481	0.797
SI1 → SOC	0.742	0.732	0.075	9.904	0.576	0.863
SI2 → SOC	0.912	0.901	0.047	19.349	0.795	0.975
SI3 → SOC	0.748	0.739	0.073	10.245	0.583	0.865
SV1 → STV	0.688	0.680	0.067	10.219	0.539	0.802
SV2 → STV	0.742	0.734	0.060	12.435	0.606	0.836
SV3 → STV	0.887	0.879	0.042	21.040	0.784	0.949
SV4 → STV	0.919	0.907	0.040	22.839	0.813	0.970
US1 ← USE	0.647	0.645	0.040	16.352	0.563	0.721
US2 ← USE	0.659	0.662	0.039	16.728	0.582	0.736
US3 ← USE	0.669	0.668	0.043	15.571	0.581	0.749
US4 ← USE	0.819	0.815	0.037	22.144	0.735	0.881

Source: created by the authors