JURNAL LOCUS: Penelitian & Pengabdian
Volume 5 No. 11 November 2026, pp 10427-10443
E-ISSN 2829-7334
DOI: 10.58344/locus. v5i11.6531

Privacy Calculus in Agentic AI Adoption: The Roles of Perceived Human Oversight, Perceived Explainability, and Trust in Users' Willingness to Delegate Tasks

Vebrita Rizky Auliasari*, Teguh Raharjo, Farizky
Universitas Indonesia, Indonesia
Email : vebrita.rizky@ui.ac.id*

Keywords:
agentic ai; privacy calculus; human oversight; explainability; task delegation.

This is an open access article under the CC BY-SA license.

Copyright © 2026 by Author. Published by Riviera Publishing

ABSTRACT

Agentic artificial intelligence (agentic AI) enables systems to plan and execute tasks with increasing autonomy, requiring users to determine the extent of access and decision-making authority they are willing to delegate. Although privacy calculus theory explains the trade-offs between perceived benefits and privacy risks, limited empirical research has examined the roles of perceived human oversight, perceived explainability, and trust in users' willingness to delegate tasks to agentic AI systems. This study investigated the relationships between perceived benefits, perceived privacy risks, perceived human oversight, perceived explainability, trust, and willingness to delegate tasks, with trust examined as a mediating variable. A cross-sectional online survey was conducted among 500 eligible AI users using a standardized task delegation scenario and five-point Likert-scale measures. Data were analyzed using partial least squares structural equation modeling (PLS-SEM). The findings revealed that perceived benefits, perceived human oversight, perceived explainability, and trust were positively associated with willingness to delegate tasks, whereas perceived privacy risks exhibited a negative association. Perceived human oversight and perceived explainability were also positively associated with trust, which partially mediated their respective relationships with willingness to delegate tasks. These findings extend the application of privacy calculus theory beyond information disclosure to the delegation of decision-making authority and task execution. The study further highlights the importance of integrating clearly defined permission boundaries, meaningful human oversight and intervention mechanisms, understandable AI-generated explanations, and robust privacy safeguards into agentic AI systems to facilitate informed, trustworthy, and responsible task delegation.

INTRODUCTION

The development of agentic artificial intelligence (agentic AI) represents a shift from systems that primarily provide information or recommendations toward systems capable of interpreting goals, planning multistep processes, accessing data and external tools, and executing actions with varying degrees of autonomy. Consequently, users must determine not only whether to adopt AI technologies but also the extent of decision-making authority they are willing to delegate. Task delegation is not necessarily a binary decision, as users may authorize different levels of autonomy depending on task characteristics, perceived benefits, and associated risks (Mihale-Wilson, 2025a). Industry reports highlight the growing practical significance of this issue. According to Google Cloud (2025), 52% of surveyed executives reported deploying AI agents in production environments, while Gartner (2025) identified substantial concerns regarding AI agent-related risks and organizational governance readiness.

These concerns are further supported by empirical research documenting privacy leakage and potentially harmful actions performed by autonomous AI agents (Ruan et al., 2024; Shao et al., 2024).

These delegation decisions can be examined through privacy calculus theory, which conceptualizes privacy-related behavior as a trade-off between perceived benefits and perceived privacy risks (Meier & Krämer, 2024). However, previous research applying privacy calculus theory has predominantly examined information disclosure and technology acceptance, particularly individuals' willingness to share personal information when anticipated benefits are perceived to outweigh or justify potential privacy risks (Dey & Kumar, 2026; Meier & Krämer, 2024; Wang et al., 2024; Zhou & Tang, 2025). Agentic AI introduces a more complex decision-making context because users must evaluate not only whether to disclose information but also whether to authorize systems to access data and resources, exercise discretion in task execution, and perform actions on their behalf (Gor et al., 2026; Liu et al., 2025; Mihale-Wilson, 2025a; Nguyen & Nallaperuma, 2026). Accordingly, the behavioral decision extends beyond information disclosure to delegated agency, whereby users grant AI systems varying degrees of access, operational discretion, and authority to execute tasks (Rodea, 2025; Wagner, 2026).

Beyond the conventional benefit–risk trade-off, users' willingness to delegate tasks may also depend on perceived human oversight, perceived explainability, and trust. Perceived human oversight refers to users' subjective assessments of their ability to monitor, constrain, intervene in, override, or terminate agentic AI operations. Perceived explainability concerns the extent to which users consider an AI system's decision-making rationale, operational processes, use of information, and potential consequences understandable. Previous research suggests that the perceived adequacy of human oversight may be more influential than the mere presence of human involvement (Nagpal & Cotte, 2026). Similarly, perceived explainability has been positively associated with trust and AI acceptance (Cheung & Ho, 2025; Zhang et al., 2024). However, providing explanations does not consistently enhance users' subjective trust or promote appropriately calibrated reliance on AI systems (Leichtmann et al., 2023; Naiseh et al., 2023; Scharowski et al., 2023). These mixed findings underscore the importance of distinguishing users' subjective perceptions of oversight and explainability from the objective technical capabilities of AI systems. Trust may therefore function as a key psychological mechanism through which users evaluate the reliability and acceptability of agentic AI before granting it operational authority (Akhil & Lyytinen, 2025; Dang & Li, 2026; Gor et al., 2026; Huynh & Aichner, 2025; Liu et al., 2025).

To address these theoretical and empirical gaps, this study extends the behavioral scope of privacy calculus theory from information disclosure to delegated agency, encompassing access permissions, operational discretion, and action authority (Dey & Kumar, 2026; Gor et al., 2026; Liu et al., 2025; Meier & Krämer, 2024; Mihale-Wilson, 2025a; Nguyen & Nallaperuma, 2026; Zhou & Tang, 2025). Specifically, the study integrates perceived benefits, perceived privacy risks, perceived human oversight, perceived explainability, and trust into a unified conceptual model explaining users' willingness to delegate tasks to agentic AI systems. It further examines the mediating role of trust in the relationships between perceived human oversight, perceived explainability, and willingness to delegate tasks. Rather than modifying the fundamental benefit–risk mechanism underlying privacy calculus theory, this study broadens its application to decisions involving the delegation of operational authority to autonomous AI systems. Accordingly, the study addresses two research questions: (1) How are perceived benefits, perceived privacy risks, perceived human oversight, perceived explainability, and trust associated with users' willingness to delegate tasks to agentic AI systems? (2) How are perceived human oversight and perceived explainability associated with

trust, and to what extent does trust mediate their respective relationships with users' willingness to delegate tasks to agentic AI systems?

METHOD

Sampling Method

This study employed purposive eligibility criteria combined with convenience sampling and voluntary self-selection through social media platforms, messaging applications, and community networks (Memon et al., 2025). Eligible participants were at least 18 years old, had used artificial intelligence (AI) within the preceding three months, had prior experience delegating tasks to AI systems, and demonstrated an understanding of the standardized agentic AI scenario. An a priori statistical power analysis, assuming five predictors of willingness to delegate tasks, a medium effect size ($f^2 = 0.15$), a significance level of 0.05, and statistical power of 0.80, indicated a minimum sample size of approximately 92 participants. A larger sample was targeted to accommodate mediation analysis and data quality screening, consistent with statistical power-based sample size recommendations rather than the conventional 10-times rule (Guenther et al., 2023).

Data Collection

Data were collected through an online questionnaire administered in Bahasa Indonesia using a five-point Likert scale (1 = strongly disagree to 5 = strongly agree). The questionnaire measured six constructs: Perceived Benefits, Perceived Privacy Risks, Perceived Human Oversight, Perceived Explainability, Trust, and Willingness to Delegate Tasks.

Before answering the measurement items, all respondents were presented with the same standardized hypothetical Agentic AI scenario. The scenario described an AI assistant capable of receiving user-defined goals, planning multiple steps, accessing authorized data and applications, and performing certain actions with delegated authority. It included examples across different professional contexts and illustrated how an assistant could prepare a monthly work report, retrieve and synthesize relevant information, prepare follow-up communications, and execute authorized actions. The scenario also described users' ability to monitor AI activities, review results, approve or reject actions, modify instructions, stop processes, or take over tasks. In addition, the assistant could explain its activities, data usage, reasons for actions, and potential consequences. This standardized scenario provided a common reference for respondents across different professional backgrounds and enabled them to evaluate Agentic AI based on the same hypothetical capabilities.

The questionnaire was administered in Bahasa Indonesia. To ensure linguistic and contextual equivalence with the original measurement items, the adapted items underwent translation and back-translation, expert review for content and contextual validity, and pilot testing with eligible respondents before the main data collection (Cruchinho et al., 2024; Elangovan & Sundaravel, 2021). The questionnaire was distributed through social media, messaging applications, and community networks from September to October 2026.

Measurement items were adapted from prior studies on Privacy Calculus (Dey & Kumar, 2026; Meier & Krämer, 2024), human oversight and perceived control (Adam et al., 2024; Langer et al., 2025; Nagpal & Cotte, 2026; Siebert et al., 2023), perceived explainability (Cheung & Ho, 2025; Zhang et al., 2024), trust in AI (Huynh & Aichner, 2025; Song & Lin, 2024), and human–AI delegation (Adam et al., 2024; Liu et al., 2025; Mihale-Wilson, 2025a; Song & Lin, 2024), as summarized in Table 1.

Table 1. Measurement Instrument

Construct	Items	Source	Adaptation
Perceived Benefits (PB)	4	(Dey & Kumar, 2026; Meier & Krämer, 2024)	Benefits of AI reframed toward efficiency, effort reduction, and task performance in Agentic AI delegation.
Perceived Privacy Risks (PPR)	4	(Dey & Kumar, 2026; Meier & Krämer, 2024)	Privacy risks extended to autonomous data access, unintended use, disclosure, and loss of data control.
Perceived Human Oversight (PHO)	4	(Adam et al., 2024; Langer et al., 2025; Nagpal & Cotte, 2026; Siebert et al., 2023)	Adapted to users' perceived ability to monitor, intervene, restrict, and terminate Agentic AI actions.
Perceived Explainability (PE)	4	(Cheung & Ho, 2025; Zhang et al., 2024)	Adapted to users' perceived understanding of Agentic AI rationale, processes, information use, and consequences.
Trust (TR)	4	(Huynh & Aichner, 2025; Song & Lin, 2024)	Cognitive trust adapted to Agentic AI competence, reliability, integrity, and delegated actions.
Willingness to Delegate Tasks (WTD)	5	(Adam et al., 2024; Liu et al., 2025; Mihale-Wilson, 2025a; Song & Lin, 2024)	Extended to task delegation, access, discretion, and action authority in Agentic AI.

Data Quality and Common Method Bias

Data quality was assessed by examining respondent eligibility, response completeness, potential duplicate submissions, and response patterns indicative of careless responding (Ward & Meade, 2023). To mitigate potential common method bias (CMB), procedural safeguards included respondent anonymity, neutral item wording, and the use of a standardized Agentic AI scenario to provide a consistent reference across respondents (Podsakoff et al., 2024). Although these procedures were intended to reduce potential method-related bias, they do not establish that CMB was absent. Screening variables were used to confirm respondent eligibility, while demographic characteristics and AI usage profiles were used to describe the sample.

Data Analysis Method

Data were analyzed using Partial Least Squares Structural Equation Modeling (PLS-SEM) to examine multiple direct and indirect relationships (Guenther et al., 2023). The reflective measurement model was assessed using outer loadings, Cronbach's alpha, rho_A, composite reliability, AVE, HTMT, and VIF, while the structural model was evaluated using standardized path coefficients (β), R^2, and f^2 (Guenther et al., 2023). Hypothesis testing and mediation analysis were performed using 5,000 bootstrap subsamples, a two-tailed test at $\alpha = .05$, and 95% BCa confidence intervals, examining direct, indirect, and total effects (Sarstedt & Moisescu, 2024).

Ethics and Informed Consent

Before participation, respondents are provided with information regarding the purpose of the study, voluntary participation, expected completion time, data usage, and their right to discontinue participation before submission. Informed consent is obtained before respondents proceed to the questionnaire. Responses are collected and analyzed confidentially, and directly identifying information is not included in the analytical dataset.

RESULTS AND DISCUSSION

Respondent Characteristics

The two-month data collection period resulted in 500 eligible and valid responses retained for analysis. Respondents were recruited through voluntary participation in an online survey distributed via social media, messaging applications, and community networks, following the eligibility criteria described in the Sampling Method section. Respondent information was categorized into screening variables, demographic characteristics, and AI usage profiles. Screening results confirmed respondent eligibility, while demographic and AI usage data described the composition and experience of the sample. Respondent characteristics are summarized in Tables II–IV.

Table 2. Respondent screening

Profile	Item	Frequency	Percentage
Age	≥ 18 years old	500	100%
AI Usage Experience	Have used AI	500	100%
	Have not used AI	0	0%
Experience Delegating Tasks to AI	Have delegated tasks to AI	500	100%
	Have not delegated tasks to AI	0	0%
Understanding of the Agentic AI Scenario	Understands the scenario	500	100%
	Does not understand the scenario	0	0%

Table 3. Respondent demographics

Profile	Item	Frequency	Percentage
Gender	Male	261	52.2%
	Female	239	47.8%
Age	18–24 years old	140	28.0%
	25–34 years old	180	36.0%
	35–44 years old	110	22.0%
	≥45 years old	70	14.0%
Educational Level	Senior High School or equivalent	32	6.4%
	Diploma	43	8.6%
	Bachelor's degree	353	70.6%
	Master's or doctoral degree	72	14.4%
Employment Status	Full-time	375	75.0%
	Part-time	22	4.4%
	Self-employed	33	6.6%
	Freelancer	21	4.2%
	Student	49	9.8%
Job Function	Information Technology (IT)	145	29.0%
	Data / Analytics / Artificial Intelligence	64	12.8%
	Student/Academic Activities	49	9.8%
	Marketing / Sales / Business Development	40	8.0%
	Creative / Media / Design	35	7.0%
	Operations / Supply Chain / Administration	28	5.6%
	Finance/Accounting	27	5.4%

Profile	Item	Frequency	Percentage
Industry	Business / Management / Consulting	27	5.4%
	Others	85	17.0%
	Technology / Telecommunications	114	22.8%
	Banking / Financial Services / Insurance / Fintech	104	20.8%
	Retail / E-commerce / FMCG	69	13.8%
	Student / Not currently working	49	9.8%
	Professional Services / Consulting	45	9.0%
	Media / Entertainment / Creative Industry	29	5.8%
	Others	90	18.0%

Table 4. AI Usage Profile of Respondents

Profile	Item	Frequency	Percentage
Duration of AI Use	<1 year	90	18,0%
	1–3 years	297	59,4%
	>3 years	113	22,6%
Frequency of AI Use	Rarely (≤1 time per month)	41	8,2%
	Occasionally (2–3 times per month)	87	17,4%
	Frequently (1–3 times per week)	197	39,4%
	Very Frequently (≥4 times per week)	175	35,0%
Frequency of Task Delegation to AI	Rarely (≤1 time/month)	71	14,2%
	Occasionally (2–3 times per month)	123	24,6%
	Frequently (1–3 times per week)	187	37,4%
	Very Frequently (≥4 times per week)	119	23,8%

Measurement Model Assessment

The measurement model was evaluated to assess indicator reliability, internal consistency reliability, and convergent validity (Guenther et al., 2023). As presented in Table 5, all indicator outer loadings exceeded the recommended threshold of approximately 0.70 (Guenther et al., 2023), ranging from 0.738 to 0.850, indicating satisfactory indicator reliability. Internal consistency reliability was also established, with Cronbach's alpha values ranging from 0.759 to 0.852, rho_A values from 0.771 to 0.852, and composite reliability (rho_C) values from 0.846 to 0.900. All reliability coefficients exceeded the recommended threshold of 0.70 (Guenther et al., 2023). Furthermore, the Average Variance Extracted (AVE) values ranged from 0.575 to 0.692, exceeding the recommended threshold of 0.50 (Guenther et al., 2023) and therefore supporting convergent validity. Overall, these results indicate that all constructs demonstrated satisfactory indicator reliability, internal consistency reliability, and convergent validity; therefore, no measurement items were removed from the model.

Table 5. Measurement Model Reliability and Convergent Validity

Construct	Outer Loading Range	Cronbach's Alpha	rho_A	rho_C	AVE
PB	0.761–0.823	0.808	0.821	0.873	0.633
PE	0.749–0.787	0.773	0.778	0.854	0.594
PHO	0.738–0.788	0.759	0.771	0.846	0.578
PPR	0.818–0.850	0.852	0.852	0.900	0.692
TR	0.744–0.792	0.784	0.786	0.860	0.606
WTD	0.741–0.785	0.816	0.818	0.871	0.575

Discriminant validity was assessed using the Heterotrait-Monotrait ratio (HTMT) (Henseler et al., 2015). As shown in Table 6, all HTMT values ranged from 0.033 to 0.589, remaining below the conservative threshold of 0.85 (Henseler et al., 2015). The highest HTMT value was observed between Trust (TR) and Willingness to Delegate Tasks (WTD) at 0.589. These results indicate satisfactory discriminant validity, demonstrating that all constructs were empirically distinct from one another.

Table 6. Discriminant Validity (HTMT)

Construct	PB	PE	PHO	PPR	TR	WTD
PB	—					
PE	0.102	—				
PHO	0.087	0.136	—			
PPR	0.058	0.033	0.063	—		
TR	0.232	0.336	0.226	0.404	—	
WTD	0.241	0.232	0.257	0.415	0.589	—

Indicator-level collinearity was assessed using the outer-model Variance Inflation Factor (VIF) (Guenther et al., 2023). As shown in Table 7, the outer-model VIF values ranged from 1.401 to 2.050. All values were below the recommended threshold, indicating that no problematic collinearity was identified among the measurement indicators.

Table 7. Outer Model VIF

Construct	VIF Range
PB	1.576–1.753
PE	1.489–1.518
PHO	1.401–1.493
PPR	1.811–2.050
TR	1.458–1.594
WTD	1.522–1.660

Structural Model Assessment

Prior to assessing the structural relationships, collinearity among predictor constructs was examined using the inner-model Variance Inflation Factor (VIF) (Guenther et al., 2023). As shown in Table 8, all VIF values ranged from 1.011 to 1.343. The highest VIF was observed for Trust → Willingness to Delegate Tasks (1.343), while the lowest values were found for Perceived Explainability → Trust and Perceived Human Oversight → Trust (1.011). These results indicate that no problematic multicollinearity was present in the structural model.

Table 8. Inner-Model Collinearity Assessment

Structural Relationship	VIF
Perceived Benefits → Willingness to Delegate Tasks	1.073
Perceived Explainability → Trust	1.011
Perceived Explainability → Willingness to Delegate Tasks	1.114
Perceived Human Oversight → Trust	1.011
Perceived Human Oversight → Willingness to Delegate Tasks	1.050
Perceived Privacy Risks → Willingness to Delegate Tasks	1.167
Trust → Willingness to Delegate Tasks	1.343

Following the evaluation of the measurement model, the structural model was assessed using the coefficient of determination (R^2), adjusted R^2, standardized path coefficients (β), and effect sizes (f^2), following established PLS-SEM guidelines (Guenther et al., 2023). The R^2 results are presented in Table 9.

Table 9. Coefficient of Determination

Endogenous Construct	R^2	Adjusted R^2
Trust (TR)	0.093	0.090
Willingness to Delegate Tasks (WTD)	0.315	0.308

The results indicate that Perceived Human Oversight (PHO) and Perceived Explainability (PE) jointly explained 9.3% of the variance in Trust (TR), with an adjusted R^2 of 0.090. Meanwhile, Perceived Benefits (PB), Perceived Privacy Risks (PPR), Perceived Human Oversight (PHO), Perceived Explainability (PE), and Trust (TR) collectively explained 31.5% of the variance in Willingness to Delegate Tasks (WTD), with an adjusted R^2 of 0.308. Thus, the model demonstrated greater explanatory power for users' willingness to delegate tasks than for trust.

The standardized path coefficients and effect sizes were subsequently examined to assess the direction and relative contribution of each predictor to its corresponding endogenous construct. The results are summarized in Table 10.

Table 10. Structural Path Coefficients and Effect Sizes

Structural Relationship	Path Coefficient (β)	f^2	Effect Size
PB → WTD	0.156	0.033	Small
PPR → WTD	−0.261	0.085	Small
PHO → TR	0.151	0.025	Small
PE → TR	0.250	0.068	Small
TR → WTD	0.304	0.101	Small
PHO → WTD	0.148	0.031	Small
PE → WTD	0.101	0.013	Negligible

Among the predictors of Trust, Perceived Explainability showed a larger positive path coefficient ($\beta = 0.250$) than Perceived Human Oversight ($\beta = 0.151$). For Willingness to Delegate Tasks, Trust showed the largest positive path coefficient ($\beta = 0.304$), whereas Perceived Privacy Risks showed a negative association ($\beta = -0.261$). Perceived Benefits ($\beta = 0.156$), Perceived Human Oversight ($\beta = 0.148$), and Perceived Explainability ($\beta = 0.101$) were also positively associated with Willingness to Delegate Tasks.

The f^2 values ranged from 0.013 to 0.101. Using conventional benchmarks of approximately 0.02, 0.15, and 0.35 for small, medium, and large effects, respectively (Guenther et al., 2023), most relationships demonstrated small effect sizes. Trust had the largest effect on Willingness to Delegate Tasks ($f^2 = 0.101$), followed by Perceived Privacy Risks on Willingness to Delegate Tasks ($f^2 = 0.085$) and Perceived Explainability on Trust ($f^2 = 0.068$). The direct effect of Perceived Explainability on Willingness to Delegate Tasks was negligible ($f^2 = 0.013$), suggesting that its unique contribution to explaining WTD was limited after accounting for the other predictors.

Hypothesis Testing

The hypothesized direct relationships were assessed using a bootstrapping procedure with 5,000 subsamples, a two-tailed test, and a significance level of 5% ($\alpha = .05$). Statistical significance was evaluated based on the standardized path coefficients (β), t-values, p-values, and 95% bias-corrected and accelerated (BCa) bootstrap confidence intervals (Guenther et al., 2023). A hypothesis was considered supported when the path coefficient was consistent with the hypothesized direction, the p-value was below .05, and the 95% BCa confidence interval did not include zero. The results are presented in Table 11.

Table 11. Hypothesis Testing Result

Hypothesis	β	t-value	p-value	95% BCa CI	Decision
H1: PB → WTD	0.156	4.391	< .001	[0.084, 0.221]	Supported
H2: PPR → WTD	−0.261	6.117	< .001	[−0.341, −0.174]	Supported
H3: PHO → TR	0.151	3.455	.001	[0.060, 0.231]	Supported
H4: PE → TR	0.250	5.777	< .001	[0.158, 0.331]	Supported
H5: TR → WTD	0.304	7.135	< .001	[0.226, 0.390]	Supported
H6: PHO → WTD	0.148	4.081	< .001	[0.073, 0.215]	Supported
H7: PE → WTD	0.101	2.606	.009	[0.022, 0.172]	Supported

All seven direct-effect hypotheses were statistically supported. Perceived Benefits showed a positive association with Willingness to Delegate Tasks ($\beta = 0.156$, $p < .001$), supporting H1. In contrast, Perceived Privacy Risks were negatively associated with Willingness to Delegate Tasks ($\beta = -0.261$, $p < .001$), supporting H2. The corresponding 95% confidence intervals, [0.084, 0.221] for H1 and [−0.341, −0.174] for H2, did not include zero, further supporting the statistical significance of these relationships.

Perceived Human Oversight was positively associated with Trust ($\beta = 0.151$, $p = .001$), while Perceived Explainability also showed a positive association with Trust ($\beta = 0.250$, $p < .001$), supporting H3 and H4, respectively. Among the two antecedents of Trust, Perceived Explainability exhibited the larger standardized path coefficient. The 95% confidence intervals for both relationships excluded zero, indicating statistically significant associations.

Trust demonstrated the largest positive direct association with Willingness to Delegate Tasks among the positive predictors ($\beta = 0.304$, $p < .001$), supporting H5. Perceived Human Oversight was also positively associated with Willingness to Delegate Tasks ($\beta = 0.148$, $p < .001$), supporting H6. Finally, Perceived Explainability showed a smaller but statistically significant positive association with Willingness to Delegate Tasks ($\beta = 0.101$, $p = .009$), supporting H7. In all three cases, the corresponding confidence intervals did not include zero.

Overall, the results provide empirical support for all proposed direct relationships. Trust showed the strongest positive association with Willingness to Delegate Tasks, whereas Perceived Privacy Risks exhibited a comparatively substantial negative association. These findings are consistent with the proposed research framework, in which perceived benefits, privacy risks, human oversight, explainability, and trust are associated with users' willingness to delegate access, discretion, and action authority to Agentic AI.

Mediation Analysis

The mediating role of Trust (TR) was examined by assessing the specific indirect effects, direct effects, and total effects using the bootstrapping procedure with 5,000 subsamples and 95% bias-corrected and accelerated (BCa) confidence intervals (Guenther et al., 2023; Sarstedt & Moisescu, 2024). An indirect effect was considered statistically significant when its p-value was below .05 and the corresponding 95% BCa confidence interval did not include zero (Guenther et al., 2023; Sarstedt & Moisescu, 2024). The mediation results for H8 and H9 are presented in Table 12.

Table 12. Mediation Analysis Result

Hypothesis	Indirect β	t-value	p-value	95% BCa CI	Direct Effect β (p)	Total Effect β (p)
H8: PHO → TR → WTD	0.046	3.063	.002	[0.020, 0.078]	0.148 (< .001)	0.194 (< .001)
H9: PE → TR → WTD	0.076	4.566	< .001	[0.047, 0.112]	0.101 (.009)	0.177 (< .001)

For H8, Perceived Human Oversight exhibited a significant positive indirect association with Willingness to Delegate Tasks through Trust ($\beta = 0.046$, $t = 3.063$, $p = .002$). The 95% BCa confidence interval ranged from 0.020 to 0.078 and did not include zero, confirming the significance of the indirect effect. The direct relationship between Perceived Human Oversight and Willingness to Delegate Tasks also remained statistically significant ($\beta = 0.148$, $p < .001$), while the total effect was $\beta = 0.194$ ($p < .001$). Therefore, Trust partially mediates the relationship between Perceived Human Oversight and Willingness to Delegate Tasks, supporting H8.

For H9, Perceived Explainability also demonstrated a significant positive indirect association with Willingness to Delegate Tasks through Trust ($\beta = 0.076$, $t = 4.566$, $p < .001$). Its 95% BCa confidence interval of [0.047, 0.112] excluded zero. At the same time, the direct relationship between Perceived Explainability and Willingness to Delegate Tasks remained significant ($\beta = 0.101$, $p = .009$), and the total effect was also significant ($\beta = 0.177$, $p < .001$). These results indicate that Trust partially mediates the relationship between Perceived Explainability and Willingness to Delegate Tasks, thereby supporting H9.

Because both the direct and indirect effects were statistically significant and operated in the same positive direction, the two mediation relationships can more specifically be characterized as complementary partial mediation (Guenther et al., 2023; Sarstedt & Moisescu, 2024). In other words, Perceived Human Oversight and Perceived Explainability are associated with users' willingness to delegate tasks both directly and indirectly through Trust. The indirect effect of Perceived Explainability through Trust ($\beta = 0.076$) was numerically larger than that of Perceived Human Oversight ($\beta = 0.046$), suggesting that Trust represents a particularly relevant pathway linking users' perceived understanding of Agentic AI with their willingness to delegate tasks.

This finding is consistent with the theoretical positioning of Trust in the research framework as a psychological mechanism connecting users' perceptions of oversight and explainability with their willingness to grant Agentic AI greater access, discretion, and action authority.

The findings answer RQ1 by showing that users' Willingness to Delegate Tasks to Agentic AI is associated with a broader set of considerations than the conventional benefit–risk trade-off. Perceived Benefits were positively associated with Willingness to Delegate Tasks, whereas Perceived Privacy Risks were negatively associated with it. This pattern can be explained by the nature of Agentic AI delegation: users may value efficiency, reduced effort, and improved task performance, but these benefits are evaluated against the possibility that an autonomous system may obtain broader data access or perform actions involving personal or work-related information. This finding is consistent with Privacy Calculus research showing that technology-related decisions involve weighing expected benefits against perceived privacy-related costs (Dey & Kumar, 2026; Meier & Krämer, 2024; Wang et al., 2024; Zhou & Tang, 2025). However, the contribution of this study is to show that this calculus remains relevant when the behavioral decision extends beyond information disclosure. In Agentic AI, users evaluate not only whether information should be shared, but also whether the system

should be granted access, discretion, and action authority (Gor et al., 2026; Liu et al., 2025; Mihale-Wilson, 2025a; Nguyen & Nallaperuma, 2026). Theoretically, this extends the behavioral domain of Privacy Calculus from disclosure decisions to delegated agency without changing its underlying benefit–risk mechanism. Practically, Agentic AI should therefore preserve visible task benefits while limiting unnecessary privacy exposure through data minimization, least-privilege access, purpose-specific permissions, and time-bounded access to sensitive resources.

The findings further show that Perceived Human Oversight is positively associated with both Trust and Willingness to Delegate Tasks. A plausible explanation is that delegation requires users to transfer some decision or action authority to the AI, making the perceived ability to intervene particularly relevant when the system operates beyond simple recommendation. When users believe that they can monitor, restrict, override, or terminate Agentic AI actions, delegation may involve less perceived vulnerability. This interpretation is consistent with meaningful-human-control research emphasizing users' ability to intervene in autonomous systems (Langer et al., 2025; Siebert et al., 2023) and with evidence suggesting that perceived adequacy of oversight can be more consequential than the mere technical presence of a human-in-the-loop (Nagpal & Cotte, 2026). At the same time, previous studies caution that excessive human oversight may create cognitive burden (Mihale-Wilson, 2025b; Nguyen & Nallaperuma, 2026; Tang et al., 2026). The theoretical implication is therefore not that more oversight will always produce greater delegation, but that Perceived Human Oversight represents a user-level condition through which delegated autonomy becomes more acceptable. Practically, oversight mechanisms should be selective and risk-based rather than requiring constant manual supervision. High-risk actions can require approval checkpoints, while lower-risk actions may proceed automatically within predefined permission boundaries. Users should also be provided with pause, revoke, cancel, and override capabilities, together with accessible audit logs showing what the Agentic AI has done and which resources it accessed.

Perceived Explainability was positively associated with both Trust and Willingness to Delegate Tasks, addressing RQ1. However, its direct contribution to Willingness to Delegate Tasks was negligible ($f^2 = 0.013$), despite the statistically significant relationship ($\beta = 0.101$, $p = .009$). This suggests that understanding Agentic AI actions may reduce uncertainty but is insufficient on its own to justify granting access and action authority when other considerations are taken into account. Delegation also involves evaluating privacy risks, task consequences, and the ability to intervene. Notably, Perceived Explainability showed a larger contribution to Trust ($f^2 = 0.068$), while its indirect relationship with Willingness to Delegate Tasks through Trust was significant ($\beta = 0.076$). Thus, explainability may support delegation partly by strengthening Trust, consistent with research showing that explanations do not always translate directly into appropriate reliance (Leichtmann et al., 2023; Naiseh et al., 2023; Scharowski et al., 2023). Theoretically, Perceived Explainability appears more relevant as a contributor to Trust than as a standalone determinant of delegation. Practically, explanations should be provided at meaningful decision points, particularly before sensitive or irreversible actions, alongside permission boundaries and user approval mechanisms.

The findings answering RQ2 show that Trust functions as an important mechanism connecting both Perceived Human Oversight and Perceived Explainability with Willingness to Delegate Tasks. Both relationships remained direct while also operating indirectly through Trust, indicating complementary partial mediation. This means that Perceived Human Oversight and Perceived Explainability are relevant not only because they are associated with stronger Trust, but also because they have relationships with delegation beyond Trust alone. This is consistent with prior work identifying Trust as central to human–AI delegation and reliance (Akhil & Lyytinen, 2025; Gor et al., 2026; Liu et al., 2025), while extending that

literature by empirically positioning Trust between users' subjective evaluations of Agentic AI governance mechanisms and their willingness to grant the system greater authority. Theoretically, Trust should therefore be understood as one psychological pathway rather than the sole explanation for delegation. Perceived Human Oversight and Perceived Explainability can influence delegation both through users' willingness to accept vulnerability and through their more immediate evaluation of whether Agentic AI can be safely authorized to act. Practically, organizations should not attempt merely to "increase trust." The objective should be appropriately calibrated Trust, supported by observable safeguards such as permission limits, action previews, approval requirements for sensitive operations, intervention capabilities, and traceable execution records.

The relatively small effect sizes across the structural relationships suggest that no single construct provides a dominant explanation of users' Willingness to Delegate Tasks. Among the predictors, Trust ($f^2 = 0.101$) and Perceived Privacy Risks ($f^2 = 0.085$) showed the largest relative contributions, followed by Perceived Benefits ($f^2 = 0.033$) and Perceived Human Oversight ($f^2 = 0.031$). Meanwhile, Perceived Explainability demonstrated a negligible direct contribution ($f^2 = 0.013$), despite its statistically significant positive association with Willingness to Delegate Tasks. This suggests that its unique explanatory contribution is limited after accounting for the other predictors. One possible explanation is that expected benefits, oversight mechanisms, and explanations may not be sufficient to justify delegating actions involving sensitive data or consequential decisions. Notably, Perceived Explainability ($f^2 = 0.068$) contributed more to Trust than Perceived Human Oversight ($f^2 = 0.025$), while its significant indirect association with delegation through Trust suggests an additional pathway linking explainability to delegation willingness. Nevertheless, both constructs explained only 9.3% of the variance in Trust, indicating that other factors may be relevant. Potential influences include perceived AI competence, reliability, previous AI experiences, accountability, task criticality, and organizational safeguards. These factors were not examined in the present model and therefore warrant further investigation. Overall, the findings suggest that Trust and delegation involve broader evaluations of Agentic AI capabilities and task-related risks beyond the constructs examined in this study (Bach et al., 2024; Kaplan et al., 2023; Lukyanenko et al., 2022).

This study provides three theoretical implications. First, Perceived Benefits were positively associated with Willingness to Delegate Tasks, while Perceived Privacy Risks were negatively associated, supporting the Privacy Calculus view that users balance expected benefits and privacy risks (Dey & Kumar, 2026; Meier & Krämer, 2024; Wang et al., 2024; Zhou & Tang, 2025). The study extends this perspective from information disclosure and technology acceptance to delegated agency, where users decide how much access, discretion, and action authority to grant Agentic AI (Gor et al., 2026; Liu et al., 2025; Mihale-Wilson, 2025a; Nguyen & Nallaperuma, 2026). Second, Perceived Human Oversight and Perceived Explainability were positively associated with both Trust and Willingness to Delegate Tasks. This is consistent with prior research on meaningful human control, perceived oversight, and explainability (Cheung & Ho, 2025; Langer et al., 2025; Nagpal & Cotte, 2026; Shin, 2021; Siebert et al., 2023; Zhang et al., 2024), while also clarifying mixed findings showing that oversight may increase cognitive burden and explanations do not always strengthen Trust (Leichtmann et al., 2023; Mihale-Wilson, 2025b; Naiseh et al., 2023; Scharowski et al., 2023; Tang et al., 2026). Thus, oversight and explainability complement, rather than replace, the benefit–risk mechanism of Privacy Calculus. Third, Trust partially mediated the relationships of Perceived Human Oversight and Perceived Explainability with Willingness to Delegate Tasks. While prior research emphasizes Trust in human–AI reliance and delegation (Akhil & Lyytinen, 2025; Bach et al., 2024; Gor et al., 2026; Kaplan et al., 2023; Liu et al., 2025;

Lukyanenko et al., 2022), these findings show that Trust is important but not the sole mechanism, as oversight and explainability are associated with delegation both directly and indirectly through Trust.

For practitioners, the findings suggest that Agentic AI governance should translate Perceived Human Oversight, Perceived Explainability, and Perceived Privacy Risks into concrete interaction and control mechanisms rather than broad principles alone. A risk-based delegation architecture can be implemented in which low-risk and reversible actions are allowed to proceed within predefined boundaries, while high-risk actions require explicit user confirmation. In practice, this can include approval checkpoints before irreversible or sensitive actions; action previews describing what the Agentic AI intends to do; granular permission boundaries specifying accessible data, systems, and actions; time-limited or task-specific authorization; revoke, pause, cancel, and override capabilities; and audit logs that record actions, accessed resources, and user approvals. For actions involving sensitive information or significant external consequences, explanations should be provided before execution, including the rationale, relevant data used, intended action, and expected consequences. These mechanisms also address the trade-off identified in prior oversight research. Because requiring approval for every action may introduce excessive cognitive burden (Mihale-Wilson, 2025b; Nguyen & Nallaperuma, 2026; Tang et al., 2026), organizations should apply stronger oversight to higher-risk or less reversible actions while allowing routine actions to operate under predefined limits. In this way, governance does not simply maximize human intervention or explanation volume; it aligns the degree of oversight and explanation with the risk, sensitivity, and reversibility of the delegated task. This approach can support users in obtaining the benefits of Agentic AI while retaining meaningful authority over how and when the system acts on their behalf.

CONCLUSION

This study concludes that perceived benefits, perceived privacy risks, perceived human oversight, perceived explainability, and trust are significantly associated with users' willingness to delegate tasks to agentic artificial intelligence (agentic AI) systems. However, their relative contributions vary, with trust and perceived privacy risks exhibiting the largest effects, although both effect sizes remain small. Perceived explainability demonstrated a negligible direct effect size despite its statistically significant association with willingness to delegate tasks, whereas its indirect association through trust remained significant. Trust also partially mediated the relationships between perceived human oversight, perceived explainability, and willingness to delegate tasks. The relatively small effect sizes suggest that users' delegation decisions may be influenced by additional factors beyond the constructs examined in this study. Furthermore, perceived human oversight and perceived explainability collectively explained only 9.3% of the variance in trust, indicating the potential relevance of other determinants, including perceived AI competence, system reliability, previous interaction experiences, and accountability. Overall, these findings extend the application of privacy calculus theory beyond information disclosure to delegated agency, highlighting the multidimensional nature of trust and users' decisions to delegate tasks to agentic AI systems.

Several limitations should be considered when interpreting these findings. First, the cross-sectional research design captured respondents' perceptions at a single point in time, preventing causal inference and limiting the ability to examine changes in trust and delegation willingness over time. Second, the study relied on self-reported behavioral intentions, which may be subject to response biases and may not accurately reflect actual behavior when interacting with operational agentic AI systems. Third, participants were recruited through purposive eligibility criteria combined with convenience sampling and voluntary self-selection, resulting in a nonprobability sample that may limit the generalizability of the findings.

Additionally, respondents exhibited heterogeneous AI experience, including differences in the duration and frequency of AI use and previous task delegation practices, which may have influenced their perceptions and willingness to delegate. Finally, the findings were context-dependent because respondents evaluated a standardized agentic AI scenario rather than multiple task categories or operational AI systems. Willingness to delegate may vary according to task sensitivity, perceived risk, reversibility, required autonomy, and the potential consequences of AI-generated actions. Therefore, the measured willingness to delegate tasks should not be interpreted as equivalent to actual delegation behavior

FUNDING

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

CONFLICT OF INTEREST

The author declares no financial or non-financial conflicts of interest related to the research, authorship, or publication of this article.

AUTHOR CONTRIBUTIONS

The author[1]: Conceptualization, Methodology, Investigation, Formal analysis, Writing – original draft, Writing – review and editing, Visualization. Author[2]: Conceptualization, Validation, Formal analysis, Writing – review and editing, Supervision. Author[3]: Writing – review and editing. All authors have read and agreed to the published version of the manuscript.

GENERATIVE AI STATEMENT

Generative AI assistance was used to adapt manuscript formatting and convert citations into APA style.

REFERENCES

Adam, M., Diebel, C., Goutier, M., & Benlian, A. (2024). Navigating autonomy and control in human-AI delegation: User responses to technology- versus user-invoked task allocation. Decision Support Systems, 180, Article 114193. https://doi.org/10.1016/j.dss.2024.114193

Akhil, A. S. G., & Lyytinen, K. (2025). Distributed cognition and human-AI delegation in knowledge work. In Proceedings of International Conference on Information Systems (ICIS) (Paper ICIS2025-2428). https://aisel.aisnet.org/icis2025/conf_theme/conf_theme/5/

Bach, T. A., Khan, A., Hallock, H., Beltrão, G., & Sousa, S. (2024). A systematic literature review of user trust in AI-enabled systems: An HCI perspective. International Journal of Human–Computer Interaction, 40(5), 1251–1266. https://doi.org/10.1080/10447318.2022.2138826

Cheung, J. C., & Ho, S. S. (2025). Explainable AI and trust: How news media shapes public support for AI-powered autonomous passenger drones. Public Understanding of Science, 34(3), 344–362. https://doi.org/10.1177/09636625241291192

Cruchinho, P., López-Franco, M. D., Capelas, M. L., Almeida, S., Bennett, P. M., da Silva, M. M., Teixeira, G., Nunes, E., Lucas, P., & Gaspar, F. (2024). Translation, cross-cultural adaptation, and validation of measurement instruments: A practical guideline for novice researchers. Journal of Multidisciplinary Healthcare, 17, 2701–2728. https://doi.org/10.2147/JMDH.S419714

Dang, Q., & Li, G. (2026). Unveiling trust in AI: The interplay of antecedents, consequences, and cultural dynamics. AI & Society, 41, 669–692. https://doi.org/10.1007/s00146-025-02477-6

Dey, S., & Kumar, S. (2026). Privacy calculus and disclosure in LLMs: The role of privacy controls, nudges, and perceived anthropomorphism. International Journal of Information Management, 89, Article 103071. https://doi.org/10.1016/j.ijinfomgt.2026.103071

Elangovan, N., & Sundaravel, E. (2021). Method of preparing a document for survey instrument validation by experts. MethodsX, 8, Article 101326. https://doi.org/10.1016/j.mex.2021.101326

Fernandes, T., & Pereira, N. (2021). Revisiting the privacy calculus: Why are consumers (really) willing to disclose personal data online? Telematics and Informatics, 65, Article 101717. https://doi.org/10.1016/j.tele.2021.101717

Gartner. (2025, September 30). Gartner survey finds just 15% of IT application leaders are considering, piloting, or deploying fully autonomous AI agents. Gartner Newsroom. https://www.gartner.com/en/newsroom/press-releases/2025-09-30-gartner-survey-finds-just-15-percent-of-it-application-leaders-are-considering-piloting-or-deploying-fully-autonomous-ai-agents

Google Cloud. (2025, September 4). Google Cloud study reveals 52% of executives say their organizations have deployed AI agents, unlocking a new wave of business value. Google Cloud Press Corner. https://www.googlecloudpresscorner.com/2025-09-04-Google-Cloud-Study-Reveals-52-of-Executives-Say-Their-Organizations-Have-Deployed-AI-Agents%2C-Unlocking-a-New-Wave-of-Business-Value%2C1

Gor, M., Sung, Y. Y., Hou, Y., Fleisig, E., Ying, Z. I., Zhou, T., & Boyd-Graber, J. L. (2026). AI, take the wheel: What drives delegation and trust in human–computer cooperative question answering? In Findings of the Association for Computational Linguistics: ACL 2026 (pp. 8664–8689). https://doi.org/10.18653/v1/2026.findings-acl.422

Guenther, P., Guenther, M., Ringle, C. M., Zaefarian, G., & Cartwright, S. (2023). Improving PLS-SEM use for business marketing research. Industrial Marketing Management, 111, 127–142. https://doi.org/10.1016/j.indmarman.2023.03.010

Henseler, J., Ringle, C. M., & Sarstedt, M. (2015). A new criterion for assessing discriminant validity in variance-based structural equation modeling. Journal of the Academy of Marketing Science, 43(1), 115–135. https://doi.org/10.1007/s11747-014-0403-8

Huynh, M.-T., & Aichner, T. (2025). In generative artificial intelligence we trust: Unpacking determinants and outcomes for cognitive trust. AI & Society, 40, 5849–5869. https://doi.org/10.1007/s00146-025-02378-8

Kaplan, A. D., Kessler, T. T., Brill, J. C., & Hancock, P. A. (2023). Trust in artificial intelligence: Meta-analytic findings. Human Factors, 65(2), 337–359. https://doi.org/10.1177/00187208211013988

Kim, Y., Kim, S. H., Peterson, R. A., & Choi, J. (2023). Privacy concern and its consequences: A meta-analysis. Technological Forecasting and Social Change, 196, Article 122789. https://doi.org/10.1016/j.techfore.2023.122789

Langer, M., Baum, K., & Schlicker, N. (2025). Effective human oversight of AI-based systems: A signal detection perspective on the detection of inaccurate and unfair outputs. Minds and Machines, 35, Article 1. https://doi.org/10.1007/s11023-024-09701-0

Leichtmann, B., Humer, C., Hinterreiter, A., Streit, M., & Mara, M. (2023). Effects of Explainable Artificial Intelligence on trust and human behavior in a high-risk decision task. Computers in Human Behavior, 139, Article 107539. https://doi.org/10.1016/j.chb.2022.107539

Liu, J., Yue, W., Leung, A., & Zhang, X. (2025). Find the good. Seek the unity: A hidden Markov model of human-AI delegation dynamics. MIS Quarterly, 49(3), 1185–1204. https://doi.org/10.25300/MISQ/2024/18232

Lukyanenko, R., Maass, W., & Storey, V. C. (2022). Trust in artificial intelligence: From a Foundational Trust Framework to emerging research opportunities. Electronic Markets, 32, 1993–2020. https://doi.org/10.1007/s12525-022-00605-4

Meier, Y., & Krämer, N. C. (2024). The privacy calculus revisited: An empirical investigation of online privacy decisions on between- and within-person levels. Communication Research, 51(2), 178–202. https://doi.org/10.1177/00936502221102101

Memon, M. A., Ramayah, T., Ting, H., & Cheah, J.-H. (2025). Purposive sampling: A review and guidelines for quantitative research. Journal of Applied Structural Equation Modeling, 9(1), 1–23. https://doi.org/10.47263/JASEM.9(1)01

Mihale-Wilson, C. A. (2025a). From whether to how much: A multi-tiered perspective on delegation to agentic IS. In Proceedings of International Conference on Information Systems (ICIS) (Paper 32). https://aisel.aisnet.org/icis2025/hti/hti/32/

Mihale-Wilson, C. A. (2025b). Invisible paternalism in agentic IS: Rethinking autonomy in the design of AI-based support. In Proceedings of International Conference on Information Systems (ICIS) (Paper ICIS2025-2504). https://aisel.aisnet.org/icis2025/ethical_is/ethical_is/11/

Nagpal, G. K., & Cotte, J. (2026). Human in the loop, or perceived oversight? The psychological inference that drives AI credibility. Psychology & Marketing, 43(9), 2305–2316. https://doi.org/10.1002/mar.70174

Naiseh, M., Al-Thani, D., Jiang, N., & Ali, R. (2023). How the different explanation classes impact trust calibration: The case of clinical decision support systems. International Journal of Human-Computer Studies, 169, Article 102941. https://doi.org/10.1016/j.ijhcs.2022.102941

Nguyen, L., & Nallaperuma, K. (2026). Power calibration and delegation in human–Agentic AI interaction: Psychological pathways. In Proceedings of Pacific Asia Conference on Information Systems (PACIS) (Paper PACIS2026-1319). https://aisel.aisnet.org/pacis2026/ai_ethic/ai_ethic/4/

Podsakoff, P. M., Podsakoff, N. P., Williams, L. J., Huang, C., & Yang, J. (2024). Common method bias: It's bad, it's complex, it's widespread, and it's not easy to fix. Annual Review of Organizational Psychology and Organizational Behavior, 11, 17–61. https://doi.org/10.1146/annurev-orgpsych-110721-040030

Rodea, A.-B. G. (2025). The Impact of Artificial Intelligence on Corporate Decision-Making: The Role of Human Directors and Delegation to AI Systems under EU Law. *UC Dublin L. Rev.*, *25*, 1.

Ruan, Y., Dong, H., Wang, A., Pitis, S., Zhou, Y., Ba, J., Dubois, Y., Maddison, C. J., & Hashimoto, T. (2024). Identifying the risks of LM agents with an LM-emulated sandbox. In Proceedings of Int. Conf. Learn. Representations (ICLR). https://proceedings.iclr.cc/paper_files/paper/2024/hash/7274ed909a312d4d869cc328ad1c5f04-Abstract-Conference.html

Sarstedt, M., & Moisescu, O.-I. (2024). Quantifying uncertainty in PLS-SEM-based mediation analyses. Journal of Marketing Analytics, 12, 87–96. https://doi.org/10.1057/s41270-023-00231-9

Scharowski, N., Perrig, S. A. C., Svab, M., Opwis, K., & Brühlmann, F. (2023). Exploring the effects of human-centered AI explanations on trust and reliance. Frontiers in Computer Science, 5, Article 1151150. https://doi.org/10.3389/fcomp.2023.1151150

Shao, Y., Li, T., Shi, W., Liu, Y., & Yang, D. (2024). PrivacyLens: Evaluating privacy norm awareness of language models in action. In Advances in Neural Information Processing Systems, 37.

https://proceedings.neurips.cc/paper_files/paper/2024/hash/a2a7e58309d5190082390ff10ff3b2b8-Abstract-Datasets_and_Benchmarks_Track.html

Shin, D. (2021). The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI. International Journal of Human-Computer Studies, 146, Article 102551. https://doi.org/10.1016/j.ijhcs.2020.102551

Siebert, L. C., Lupetti, M. L., Aizenberg, E., Beckers, N., Zgonnikov, A., Veluwenkamp, H., Abbink, D., Giaccardi, E., Houben, G.-J., & Jonker, C. M. (2023). Meaningful human control: Actionable properties for AI system development. AI and Ethics, 3, 241–255. https://doi.org/10.1007/s43681-022-00167-3

Song, J., & Lin, H. (2024). Exploring the effect of artificial intelligence intellect on consumer decision delegation: The role of trust, task objectivity, and anthropomorphism. Journal of Consumer Behaviour, 23(2), 727–747. https://doi.org/10.1002/cb.2234

Tang, J., Chen, C., Li, J., Zhang, Z., Guo, B., Khalilov, I., Gebreegziabher, S. A., Yao, B., Wang, D., Ye, Y., Li, T., Xiao, Z., Yao, Y., & Li, T. J.-J. (2026). Dark patterns meet GUI agents: LLM agent susceptibility to manipulative interfaces and the role of human oversight. In Proceedings of CHI Conference on Human Factors in Computing Systems (CHI '26) (pp. 1–26). https://doi.org/10.1145/3772318.3791568

Wagner, R. L. (2026). *Consumer Agency and Delegation in AI-Mediated Environments*.

Wang, Y., Zhu, J., Liu, R., & Jiang, Y. (2024). Enhancing recommendation acceptance: Resolving the personalization–privacy paradox in recommender systems: A privacy calculus perspective. International Journal of Information Management, 76, Article 102755. https://doi.org/10.1016/j.ijinfomgt.2024.102755

Ward, M. K., & Meade, A. W. (2023). Dealing with careless responding in survey data: Prevention, identification, and recommended best practices. Annual Review of Psychology, 74, 577–596. https://doi.org/10.1146/annurev-psych-040422-045007

Yankouskaya, A., Almourad, M. B., Liebherr, M., Beyahi, F., Xu, G., & Ali, R. (2026). Who lets AI take over? Cross-national variation in willingness to delegate socially important roles to artificial intelligence. AI & Society, 41, 6151–6169. https://doi.org/10.1007/s00146-026-02858-5

Zhang, T., Li, W., Huang, W., & Ma, L. (2024). Critical roles of explainability in shaping perception, trust, and acceptance of autonomous vehicles. International Journal of Industrial Ergonomics, 100, Article 103568. https://doi.org/10.1016/j.ergon.2024.103568

Zhou, T., & Tang, B. (2025). Examining user information disclosure on artificial intelligence-generated content platforms: A privacy calculus perspective. Information Development. https://doi.org/10.1177/02666669251384288