Islamic Digital Religion and Intelligent Media

Vol. 1, No. 1 (October 2026), 63-86
P-ISSN: XXXX-XXXX | E-ISSN: XXXX-XXXX
DOI: https://doi.org/10.xxxxx/drim.vXiX.XX

Legitimacy, Trust, and Perceived Authority: A Serial Mediation Model of Muslim Undergraduates' Intention to Consult Generative Artificial Intelligence on Religious Questions

OPEN ACCESS

Umi Musyarrofah[1*], Kalsum Minangsih[1], Athia Nur Kamilah[2]

[1] Universitas Islam Negeri Syarif Hidayatullah Jakarta, Indonesia
[2] Universiti Utara Malaysia, Malaysia

[*]Correspondence: umi.musyarrofah@uinjkt.ac.id

Received: 21-10-2026	Revised: 07-10-2026	Accepted: 29-10-2026

Abstract. This study examines how Muslim undergraduate students come to treat generative artificial intelligence as an acceptable source of religious guidance, and it asks whether judging such guidance legitimate is sufficient to shape religious information behaviour. A cross-sectional survey was administered to 320 Muslim undergraduates aged 18 to 24 across five university categories in Greater Jakarta who had consulted a generative system on an Islam-related question during the previous three months. Six reflective constructs were measured on five-point Likert scales, namely information quality, source transparency, religious legitimacy, religious artificial intelligence trust, perceived artificial intelligence religious authority, and behavioural intention. Data were analysed using partial least squares structural equation modelling with 5,000 bootstrap resamples. The measurement model satisfied all conventional thresholds, with outer loadings between 0.824 and 0.885, composite reliability between 0.910 and 0.920, average variance extracted between 0.717 and 0.743, and a maximum heterotrait-monotrait ratio of 0.627. Seven of the eight direct hypotheses were statistically supported. Information quality and source transparency both predicted religious legitimacy, legitimacy strongly predicted trust, and perceived authority carried the largest direct effect on behavioural intention. The direct path from legitimacy to behavioural intention was not significant, whereas the serial indirect path through trust and authority was significant. The model explained 35.8 per cent of the variance in behavioural intention. Legitimacy therefore operates as a distal judgement that reaches conduct only after it has become reliance and then deference, a theoretical ordering supported by model fit in a cross-sectional design, rather than a demonstrated temporal sequence; see Limitations which locates religious literacy intervention at the trust-to-authority transition.

Keywords: generative artificial intelligence; Muslim students; PLS-SEM; religious authority; religious legitimacy

Introduction

Generative artificial intelligence has entered Muslim religious life at a speed that institutional scholarship has not matched. Students now pose questions about worship, permissible conduct, and scriptural meaning to systems that answer instantly, fluently, and without the credentials that Islamic tradition has long required of anyone who issues a ruling (Kurniawan et al., 2025). Empirical audits of how these systems answer religious questions have documented uneven accuracy, unverifiable citations, and inconsistent jurisprudential positions across models and languages (Tsuria & Tsuria, 2024; Wahid, 2025). The practical question for religious education

Copyright © 2026 by Author(s)
This work is licensed under a **Creative Commons Attribution-NonCommercial 4.0 International License**.

is therefore no longer whether students encounter machine-generated religious content, because they plainly do, but how they come to grant that content standing in their own religious reasoning. Qualitative work has begun to describe what that granting looks like in practice. A study of Muslim users of generative tools found that respondents cast these systems in three distinct roles, treating them as religious assistants, as perceived religious scholars, and as faith counsellors, while simultaneously voicing fears of misguidance (Jahan et al., 2026). A phenomenological study of madrasah students identified several engagement types and found that institutional autonomy becomes superficial without explicit critical literacy (Afzainizam & Rahman, 2026). Research on technology acceptance offers one established way to approach this question. Studies applying the unified theory of acceptance and use of technology to generative systems have shown robust relationships between perceived usefulness, social influence, and behavioural intention among university students, including large surveys conducted in Indonesian universities (Habibi et al., 2024). Parallel work has modelled continuance intention for generative systems using stimulus-organism-response logic and has confirmed that information quality operates through intermediate evaluative states rather than directly on intention (Mun & Hwang, 2025; Zhou & Ma, 2025). An extended technology acceptance model applied to Muslim students in Islamic religious education has likewise demonstrated that general acceptance constructs transfer to religious learning settings (Faizin, Alfan, et al., 2025). That model's acceptance constructs describe whether machine-generated religious content is perceived as useful or easy to engage with, not whether it is recognised as consistent with Islamic doctrinal principles and admissible as a preliminary basis for religious understanding; the religious legitimacy construct proposed here targets this narrower doctrinal-admissibility judgement, which perceived-usefulness items were not designed to capture. Comparable continuance-intention studies conducted with student samples in other Muslim-majority settings report the same broad structure of relationships (Bekturova et al., 2025), and recent comparative work in higher education confirms that these constructs behave consistently across student and instructor populations (Ayalon & Ido, 2026; Fatmawati, Khoiriyah, et al., 2026).

This acceptance literature nonetheless leaves something important unmeasured. Religious information is not an ordinary informational good, because its acceptability depends on whether a community recognises the source as entitled to speak (Noh et al., 2026). Scholarship on digital religion has argued for two decades that authority shifts rather than disappears when religious life moves into networked media, and that legitimacy is contested through visibility, platform affordance, and institutional recognition (Campbell & Tsuria, 2022; Ramadhan et al., 2026). Recent work extends this reasoning to algorithmically mediated environments, proposing that platform infrastructure and institutional gatekeeping jointly determine which religious expressions count as authoritative (Akhmetova et al., 2026). Studies of Indonesian religious organisations have shown that digital platforms make internal disagreement visible without dissolving institutional control (Akmaliah, 2022). Related analysis proposes that digital mediation shifts the basis of juristic authority away from evidentiary reasoning and towards infrastructural prominence, so that visibility performs work that proof once performed (Atallah, 2026). Platform-mapping research describes contemporary clerical standing as produced through technocultural mediation, in which a scholar exists both as a person and as an algorithmically computed representation (Lohlker & Wahid, 2026). Interview research with younger Indonesian scholars found that they routinely source and trust online material, which has weakened the mediating role of senior figures and formal institutions (Rusli & Nurdin, 2022).

A third body of work concerns trust in artificial intelligence and its calibration. Meta-analytic evidence establishes that trust in artificial systems responds to perceived performance and transparency, and that trust in turn drives acceptance and reliance (Choung et al., 2023;

Kaplan et al., 2023). Experimental work has repeatedly demonstrated overreliance, in which users accept machine advice beyond what its accuracy warrants, and has shown that cognitive friction reduces this tendency (Bučinca et al., 2021; Schemmer et al., 2023). Source disclosure and citation have been found to raise trust even when users do not verify the sources supplied (Ding et al., 2025; Schiedermair et al., 2026). One study of Muslim American participants found that stated distrust of artificial intelligence coexisted with a strong preference for machine answers over scholar-authored ones in blind comparison (Alam et al., 2025). Methodological reviews of this field caution that trust is operationalised inconsistently across studies, which complicates comparison between them (Vereschak et al., 2021). Systematic review evidence on algorithm aversion further shows that rejection of machine advice varies by task domain rather than being a stable disposition (Mahmud et al., 2022), and integrative work in information systems has sought to reconcile aversion and appreciation within a single account (Jussupow et al., 2024). How presentation format shapes trust in machine recommendations has also been demonstrated experimentally (Kim et al., 2021).

These three literatures have not been brought together in a single measurement model, and their separation conceals a specific theoretical problem. Specifically, none of the three literatures alone can generate a prediction about whether a judgement that machine-generated religious content is doctrinally acceptable will itself predict reliance on, and deference to, that content: acceptance research lacks a doctrinal-admissibility construct, digital religion research theorises legitimacy and authority without a measurement model linking them to trust, and trust research lacks the normative criterion that would make a source doctrinally eligible for trust in the first place. Acceptance research treats trust as an antecedent of intention without asking what makes a source eligible for trust in a religious domain. Digital religion research theorises legitimacy and authority richly but rarely operationalises them as measurable constructs in a predictive model. Trust research measures calibration without reference to the normative traditions that decide who may speak on religious matters. No published study, to the author's knowledge, has modelled religious legitimacy as a measured antecedent of trust in artificial intelligence, or distinguished legitimacy from perceived authority as separate constructs with different consequences for behaviour. Legitimacy has been operationalised for algorithmic systems in secular governance settings, where perceived legitimacy of automated decisions has been measured experimentally among citizens and administrators (Hillo et al., 2026), and legitimacy theory has been extended to artificial systems in political analysis (Erman & Furendal, 2024). These treatments supply a measurement precedent without addressing the religious case, where the criteria for legitimate speech are set by a tradition rather than by a state.

The distinction matters because legitimacy and authority are not the same judgement. A student may consider machine-generated guidance broadly consistent with recognised Islamic principles, which is a judgement about content, while still declining to let that guidance shape which interpretations they actually weigh, which is a judgement about standing. Islamic legal epistemology has always held these apart, since a ruling is valid not merely because its content is defensible but because a qualified person who bears responsibility has issued it (Ahmed et al., 2026; Mahendra et al., 2026). Work on digital epistemic authority in religious learning makes a comparable distinction between accepting content and deferring to a source (Berger & Golan, 2024). This existing distinction means the present study's contribution is best framed as operationalising and empirically testing legitimacy and authority as two separable, measured constructs within a predictive PLS-SEM model, rather than as introducing the conceptual distinction itself. Whether the two constructs behave differently in practice is an empirical question that has not been tested. Survey evidence from Muslim communities elsewhere indicates that respondents value digital access while continuing to distrust unqualified online voices, and that local religious figures remain the most trusted single source

(Whyte, 2022). If legitimacy and authority were simply two expressions of one favourable evaluation, that coexistence of openness and reservation would be difficult to account for. This study addresses that question using survey data from 320 Muslim undergraduates in Greater Jakarta who had recently consulted generative systems on Islam-related matters. It models information quality and source transparency as antecedents of religious legitimacy, legitimacy as an antecedent of religious artificial intelligence trust and perceived religious authority, and all three as predictors of behavioural intention. Three research questions guide the analysis reported below. The first asks which perceived attributes of machine-generated religious content predict judgements of religious legitimacy. The second asks whether legitimacy influences behavioural intention directly or through trust and perceived authority. The third asks which construct carries the largest direct effect on intention, and what that ordering implies for religious literacy education.

The remainder of the article follows the conventional sequence. The method section describes the sampling frame, instrument development, and the partial least squares procedure, and it presents the hypothesised model. The results section reports the measurement model, the structural estimates, and the mediation analysis. The discussion interprets the pattern through a four-stage account running from evaluation to conduct, and it compares the findings with published work on trust calibration and religious authority. The conclusion states the answers to the three research questions and identifies directions for subsequent research.

Method

This study employed a cross-sectional explanatory survey design analysed with variance-based structural equation modelling. The design suits the research questions because the model contains several mediating constructs, because the constructs are latent and measured reflectively, and because the aim is to explain variance in a target construct rather than to confirm a covariance structure derived from established theory. Partial least squares structural equation modelling is appropriate for models of this kind, particularly where the theoretical framework is still being assembled from separate literatures and where prediction of a key endogenous construct is the analytic goal (Hair & Alamer, 2022; Sarstedt et al., 2023). The method has been defended against recurring objections and its current reporting conventions have been set out in detail (Hair et al., 2024; Sabol et al., 2023).

Two further considerations supported the choice of a variance-based rather than a covariance-based estimator. The model contains a serial mediation chain of four constructs in sequence, and the study aims to identify which link in that chain carries the largest effect on the target construct, which is an explanatory-predictive objective rather than a test of global model fit (Hair et al., 2022). Two of the six constructs were measured with instruments developed for this study, and variance-based estimation is more tolerant of measurement instruments at an early stage of validation than covariance-based alternatives are (Becker et al., 2023). Reviews of applied practice note that the method is often chosen without justification, so the reasoning is stated explicitly here (Guenther et al., 2023; Sarstedt et al., 2022). The target population comprised Muslim undergraduate students aged 18 to 24 enrolled in universities in Greater Jakarta who had used a generative artificial intelligence tool for an Islam-related question during the preceding three months. Greater Jakarta was selected because it contains a dense concentration of public and private universities of both general and Islamic orientation, which permits variation in institutional religious culture. Five institutional categories were distinguished, namely public Islamic, public general, private Islamic, and two private general categories, so that respondents would not be drawn from a single type of religious educational environment.

Purposive quota sampling was applied with quotas set across the five institutional categories. Eligibility required that the respondent identify as Muslim, hold active

undergraduate status, fall within the specified age range, have used a generative tool at least twice in the previous three months for a religious or Islamic question, and provide informed consent. A screening stage identified 1,180 eligible students, from whom 360 responses were obtained. Forty cases were removed during screening and data-quality checks for failed attention checks, ineligibility on the screening criteria, or patterned responding, which produced a final analytical sample of 320 observations. The retained sample exceeds the reference minimum of approximately 244 cases for a model with four predictors of a single construct at an alpha of .05, statistical power of .80, and a small effect size of .05. None of the four structural equations implied by Figure 1 has four direct predictors of a single construct: legitimacy has two antecedents (information quality, source transparency), trust has one (legitimacy), authority has two (legitimacy, trust), and intention has three (legitimacy, trust, authority); which specific equation this power calculation was based on is accordingly unclear from the information given and would benefit from clarification.

Six latent constructs were measured, each with four reflective indicators scored on a five-point Likert scale anchored at strongly disagree and strongly agree. Information quality captured perceived accuracy, relevance, internal consistency, and clarity of machine-generated Islamic information. Source transparency captured whether the system makes its reasoning traceable, flags uncertainty, permits attribution to identifiable Islamic sources, and distinguishes fact from interpretation. Religious legitimacy captured whether machine-generated guidance can be consistent with recognised Islamic principles and can serve as an acceptable preliminary reference when appropriately verified. Religious artificial intelligence trust captured willingness to rely on such guidance under conditions of transparency and verification. Perceived artificial intelligence religious authority captured whether machine-generated explanations actually shape the respondent's religious understanding and the interpretations they consider. Behavioural intention captured stated intention to continue consulting generative systems for future religious questions.

Instrument items were developed through a three-step procedure. Existing scales for information quality, transparency, trust, and behavioural intention were adapted from published instruments in information systems and human-computer interaction research, then reworded for the religious information context (Choung et al., 2023; Shin, 2021). Legitimacy and authority items were newly written, because no validated instrument measures these constructs for artificial systems in Islamic settings, and were drafted from conceptual accounts of religious legitimacy and digital epistemic authority (Abdelnour, 2025; Berger & Golan, 2024; Grimmelikhuijsen & Meijer, 2022). A panel review then examined the draft instrument for doctrinal appropriateness, clarity, and the avoidance of wording that might presuppose a ruling on permissibility, after which items were revised and piloted before full administration. The panel comprised scholars of Islamic studies and of educational measurement, and it screened for items that assumed a jurisprudential position or invited a socially desirable response. The number of panel members, their specific qualifications beyond these two disciplinary categories, and any quantitative content-validity agreement statistic are not reported, so the panel's precise effect on the final item set, including which items were revised or dropped as a result of this review, cannot be fully assessed from the information given. Item wording was deliberately kept conditional, so that agreement with a legitimacy item indicates that guidance can be acceptable under stated conditions rather than that it is acceptable without qualification. This phrasing decision matters for interpretation, because credibility assessment in mediated settings is known to depend heavily on how a claim is framed to the respondent (Malhotra & Shin, 2025). This conditional phrasing is described here specifically for the legitimacy items; whether the trust and authority items were worded with comparable conditionality is not stated, and a systematic wording difference across constructs, rather than a purely psychological

mediation process, could itself partly contribute to the weaker direct path observed for legitimacy relative to trust and authority.

Figure 1 presents the hypothesised model in full. Information quality and source transparency are modelled as antecedents of religious legitimacy, reflecting the argument that legitimacy judgements in religious information settings rest on the perceived soundness of content and on the traceability of its basis. Legitimacy is modelled as an antecedent of both trust and perceived authority. Trust is modelled as an antecedent of perceived authority, on the reasoning that deference typically follows reliance rather than preceding it. All three of legitimacy, trust, and authority are modelled as predictors of behavioural intention, with the direct legitimacy path retained so that it can be tested against the mediated pathways rather than assumed away.

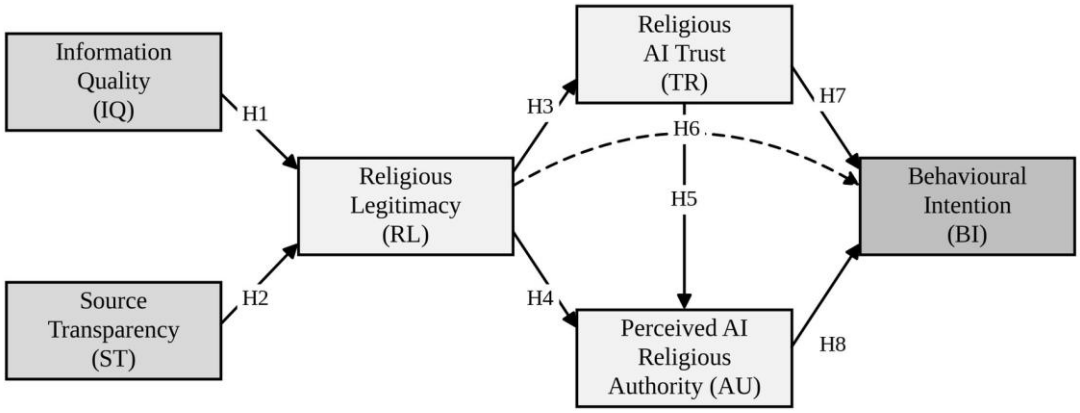


Figure 1. Hypothesised research model with direct and mediated pathways

The figure specifies eight direct hypotheses and two mediation hypotheses. Retaining the direct path from legitimacy to intention allows the analysis to distinguish an indirect-only pattern from partial mediation, which is a distinction that carries substantive meaning here rather than being a technical nicety. Should the direct path prove non-significant while the indirect paths are significant, the interpretation is that judging machine guidance acceptable is not by itself a behavioural disposition. Data were collected through a self-administered online questionnaire distributed via campus networks and student organisations across the five institutional categories. Participation was voluntary and no incentive was offered. Recruitment through campus networks and student organisations for a voluntary online survey is a convenience and self-selection mechanism layered on top of the purposive quota sampling described above; this may over-represent digitally engaged or organisationally connected students relative to the broader population of eligible Muslim undergraduates, which is a distinct source of potential bias from the single-region generalisability limitation noted later. Respondent identifiers were anonymised at the point of collection, and no directly identifying information was retained in the analytical file. Ethical approval, consent documentation, and institutional permissions were obtained in accordance with the procedures of the host institution and are recorded in the study documentation. Procedural remedies against common method

bias were applied during design, including the separation of predictor and criterion blocks in the questionnaire, assurance of anonymity, and the avoidance of leading item wording (Podsakoff et al., 2024). All six constructs, including both predictors and outcomes, were nonetheless collected from the same respondents in a single self-administered sitting, which keeps the study exposed to common method variance despite these procedural remedies and the supplementary, non-conclusive Harman's single-factor check reported below; no marker-variable test or other design-based remedy beyond block separation and anonymity was implemented, which is noted here as a limitation given how central several same-source structural paths are to the paper's main conclusions.

Analysis proceeded in the sequence recommended for reflective models. The measurement model was assessed first for indicator reliability through outer loadings, for internal consistency through Cronbach's alpha and composite reliability, and for convergent validity through average variance extracted. Discriminant validity was assessed through the Fornell-Larcker criterion, through cross-loadings, and through the heterotrait-monotrait ratio of correlations, the most sensitive available criterion (Roemer et al., 2021). The structural model was then assessed for collinearity through variance inflation factors, for explanatory power through the coefficient of determination, for effect size through f squared, and for predictive relevance through a cross-validated redundancy proxy. Path coefficients and indirect effects were tested through bootstrapping with 5,000 resamples and percentile confidence intervals, following current guidance for mediation analysis in this framework (Cheah et al., 2021; Rasoolimanesh et al., 2021). Common method bias was examined post hoc through Harman's single-factor diagnostic, which is treated here as supplementary rather than conclusive (Kock, 2021). Estimation was carried out in SmartPLS 4 using the path weighting scheme (Ringle et al., 2024). Predictive relevance was evaluated through a cross-validated redundancy proxy computed for each endogenous construct, which follows the logic of out-of-sample assessment developed for this estimator (Liengaard et al., 2021; Sharma et al., 2023). The phrase 'cross-validated redundancy proxy' mixes terminology from two distinct approaches, the older blindfolding-based Q^2 redundancy statistic and the newer cross-validated predictive ability test (CVPAT) introduced in the cited sources, which does not report results as a simple positive or negative redundancy value in the way described here; the specific statistic actually computed should be named and reported using the terminology and benchmarks appropriate to that method. Robustness was considered along the dimensions that current guidance identifies as standard, namely nonlinearity, endogeneity, and unobserved heterogeneity, and no indication of a specification problem emerged that would alter the substantive conclusions (Vaithilingam et al., 2024).

Results

Respondent characteristics

Table 1 summarises the composition of the analytical sample.

Table 1. Profile of the analytical sample (n = 320)

Characteristic	Category	n	%
Institution type	Public Islamic	74	23.1
	Public general A	72	22.5
	Private Islamic	66	20.6
	Private general A	58	18.1
	Private general B	50	15.6
Gender	Female	176	55.0

Characteristic	Category	n	%
	Male	144	45.0
Semester	1–2	38	11.9
	3–4	107	33.4
	5–6	85	26.6
	7–8	90	28.1
Primary platform	ChatGPT	231	72.2
	Gemini	49	15.3
	Perplexity	21	6.6
	Microsoft Copilot	16	5.0
	Other	3	0.9
Religious use frequency	1–2 times per month	121	37.8
	3–5 times per month	108	33.8
	6–10 times per month	52	16.3
	More than 10 times per month	39	12.2
Primary query type	Fiqh and permissible conduct	108	33.8
	Worship practice	73	22.8
	Qur'an and hadith interpretation	53	16.6
	Islamic ethics	50	15.6
	Personal religious decision	36	11.3

The sample is distributed across the five institutional categories without heavy concentration in any one, and the gender split is close to even. Respondents cluster in the middle undergraduate semesters, which is consistent with an age range of 18 to 24. Platform use is heavily concentrated, since almost three quarters of respondents named one system as their primary tool, a concentration that matches patterns reported in Indonesian university samples (Habibi et al., 2024). Most respondents consulted generative systems for religious purposes between one and five times per month, and questions of permissible conduct and worship practice together accounted for more than half of primary query types. The predominance of practical over interpretive questions matters, since permissible conduct is precisely the domain in which a ruling traditionally requires a qualified issuer. Only a minority of respondents reported consulting these systems more than ten times per month, so the sample describes occasional rather than habitual religious consultation.

Measurement model

Table 2 reports item descriptive statistics together with the measurement model assessment.

Table 2. Item descriptive statistics and reflective measurement model assessment

Construct	Item	M	SD	Loading	α	CR	AVE
Information quality	IQ1	3.51	0.86	0.873	0.874	0.914	0.726
	IQ2	3.45	0.90	0.867			
	IQ3	3.50	0.88	0.831			
	IQ4	3.48	0.91	0.835			
Source transparency	ST1	3.30	0.92	0.866	0.878	0.916	0.733
	ST2	3.32	0.90	0.862			
	ST3	3.33	0.89	0.849			

Construct	Item	M	SD	Loading	α	CR	AVE
	ST4	3.39	0.92	0.846			
Religious legitimacy	RL1	3.45	0.86	0.871	0.873	0.913	0.725
	RL2	3.43	0.85	0.867			
	RL3	3.40	0.87	0.835			
	RL4	3.44	0.82	0.833			
Religious AI trust	TR1	3.33	0.90	0.884	0.884	0.920	0.741
	TR2	3.34	0.84	0.848			
	TR3	3.37	0.88	0.852			
	TR4	3.41	0.88	0.860			
Perceived AI religious authority	AU1	3.22	0.91	0.872	0.868	0.910	0.717
	AU2	3.19	0.88	0.863			
	AU3	3.21	0.90	0.824			
	AU4	3.17	0.90	0.827			
Behavioural intention	BI1	3.53	0.87	0.879	0.885	0.920	0.743
	BI2	3.51	0.90	0.885			
	BI3	3.50	0.91	0.847			
	BI4	3.48	0.91	0.836			

Note. α = Cronbach's alpha; CR = composite reliability; AVE = average variance extracted. Items used a five-point Likert scale. Construct-level statistics appear once, on each construct's first indicator row. Absolute skewness did not exceed 0.31 and absolute excess kurtosis did not exceed 0.64.

All indicator means fall between 3.17 and 3.53 on a five-point scale, which indicates moderate rather than enthusiastic agreement across every construct. Perceived authority records the lowest construct mean at 3.20, and behavioural intention the highest at 3.51, a gap that is itself informative and is taken up in the discussion. Absolute skewness does not exceed 0.31 and absolute excess kurtosis does not exceed 0.64, so the distributions suit the estimation procedure used.

All outer loadings fall between 0.824 and 0.885, comfortably above the 0.708 benchmark, so every indicator is retained. Cronbach's alpha ranges from 0.868 to 0.885 and composite reliability from 0.910 to 0.920, both within the range that indicates satisfactory internal consistency without the redundancy that values above 0.95 would suggest. Average variance extracted ranges from 0.717 to 0.743, so each construct explains at least seventy per cent of the variance in its indicators. The measurement model therefore satisfies the conventional criteria for reflective constructs. No indicator was removed, so each construct retains its full complement of four items. Reliability estimates are uniform across all six constructs, spanning only 0.017 in alpha, which indicates that the new legitimacy and authority scales perform comparably to the adapted ones.

Table 3 reports the two discriminant validity assessments.

Table 3. Discriminant validity assessment

Panel and construct	IQ	ST	RL	TR	AU	BI
Panel A. Fornell–Larcker criterion						
Information quality (IQ)	**0.852**					
Source transparency (ST)	0.311	**0.856**				
Religious legitimacy (RL)	0.475	0.404	**0.851**			
Religious AI trust (TR)	0.227	0.224	0.546	**0.861**		

Panel and construct	IQ	ST	RL	TR	AU	BI
Perceived authority (AU)	0.191	0.166	0.450	0.549	**0.847**	
Behavioural intention (BI)	0.214	0.155	0.376	0.493	0.550	**0.862**
Panel B. Heterotrait–monotrait ratio						
Information quality (IQ)	–					
Source transparency (ST)	0.355	–				
Religious legitimacy (RL)	0.543	0.461	–			
Religious AI trust (TR)	0.259	0.254	0.622	–		
Perceived authority (AU)	0.219	0.190	0.517	0.627	–	
Behavioural intention (BI)	0.244	0.175	0.428	0.559	0.627	–

Note. Panel A diagonal values in bold are square roots of average variance extracted; off-diagonal values are construct correlations. Panel B reports heterotrait–monotrait ratios against a conservative threshold of 0.85.

In the upper panel the square root of average variance extracted exceeds every corresponding inter-construct correlation for all six constructs, which satisfies the Fornell-Larcker criterion. In the lower panel the largest heterotrait-monotrait ratio is 0.627, for the trust and authority pair and again for the authority and intention pair, well below both the conservative 0.85 threshold and the more permissive 0.90 benchmark. Examination of cross-loadings confirmed that every indicator loaded most strongly on its assigned construct. Discriminant validity is therefore established, and the constructs of legitimacy, trust, and authority are empirically distinguishable rather than a single underlying evaluation expressed three ways.

Harman's single-factor diagnostic returned a first unrotated factor accounting for 34.6 per cent of variance, which is below the fifty per cent level often used as an informal warning sign. This diagnostic is supplementary only, since it cannot establish the absence of common method variance, and the procedural remedies described earlier carry the principal weight.

Structural model

Collinearity was examined before interpreting the path coefficients. Inner variance inflation factors range from 1.000 to 1.705, far below the threshold of 5.0 and below the more cautious threshold of 3.3, so multicollinearity does not distort the estimates. Table 4 reports the structural results and Figure 2 presents them in diagrammatic form.

Table 4. Structural model results for the direct hypotheses

H	Path	β	SE	t	p	95% CI	f^2	VIF	Decision
H1	IQ → RL	0.386	0.050	7.803	< .001	[0.286, 0.482]	0.192	1.107	Supported
H2	ST → RL	0.284	0.049	5.767	< .001	[0.187, 0.380]	0.104	1.107	Supported
H3	RL → TR	0.546	0.041	13.312	< .001	[0.460, 0.620]	0.424	1.000	Supported
H4	RL → AU	0.214	0.053	4.066	< .001	[0.109, 0.316]	0.048	1.424	Supported
H5	TR → AU	0.432	0.048	8.982	< .001	[0.333, 0.523]	0.197	1.424	Supported
H6	RL → BI	0.070	0.053	1.323	.186	[–0.031, 0.175]	0.005	1.493	Not supported
H7	TR → BI	0.244	0.057	4.296	< .001	[0.131, 0.353]	0.054	1.705	Supported
H8	AU → BI	0.385	0.055	6.973	< .001	[0.273, 0.492]	0.154	1.500	Supported

Note. Standardised coefficients from 5,000 bootstrap resamples with percentile confidence intervals. R^2 = .298 (legitimacy), .298 (trust), .333 (authority), and .358 (intention). Approximate SRMR = .050.

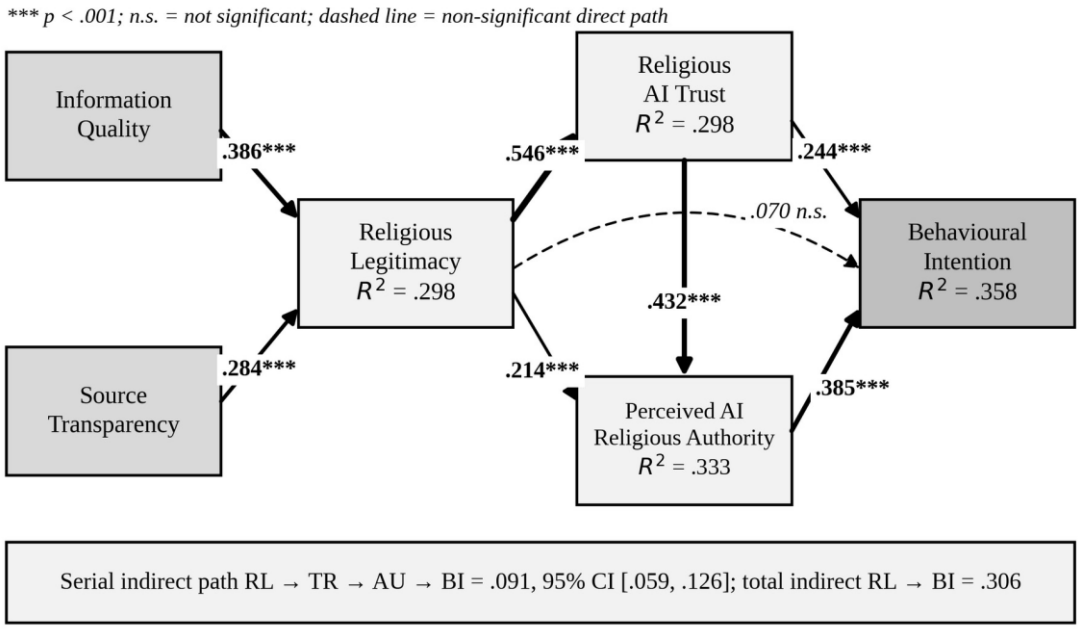


Figure 2. Estimated structural model with standardised path coefficients and explained variance

Seven of the eight direct hypotheses are supported. Information quality predicts religious legitimacy most strongly among the two antecedents, with a standardised coefficient of 0.386 and a medium effect size, while source transparency contributes a smaller but clearly significant coefficient of 0.284. Together the two explain 29.8 per cent of the variance in legitimacy. Religious legitimacy predicts trust with a coefficient of 0.546, which is the largest single path in the model and corresponds to a large effect size of 0.424. Legitimacy and trust both predict perceived authority, with trust contributing roughly twice the coefficient of legitimacy at 0.432 against 0.214. This ordering is itself a finding, because it indicates that the willingness to depend on a system contributes more to conceding its standing than the prior judgement that its content is acceptable. Effect sizes follow the same ordering, with trust reaching 0.197 against 0.048 for legitimacy, so the difference is not an artefact of coefficient scaling alone.

The results for behavioural intention are the most consequential. Perceived authority carries the largest direct effect at 0.385, trust carries a smaller but significant effect at 0.244, and the direct effect of religious legitimacy is 0.070 with a confidence interval spanning zero. Given the strong upstream paths already estimated (legitimacy to trust = 0.546; trust to authority = 0.432), a small residual direct effect of legitimacy on intention is close to what the model structure would produce almost by construction once most of legitimacy's variance is channelled through trust and authority; this structural consideration should be read alongside, not instead of, the substantive interpretation offered below. Hypothesis six is therefore not supported by these data. The model explains 35.8 per cent of the variance in behavioural intention, with moderate explanatory power in the remaining endogenous constructs and positive predictive relevance proxies throughout. The approximate standardised root mean square residual of 0.050 falls below the 0.08 diagnostic benchmark. Explained variance rises monotonically along the causal chain, from 29.8 per cent for legitimacy to 35.8 per cent for intention, as expected where later constructs accumulate earlier influence. All four cross-

validated redundancy proxies are positive, ranging from 0.280 to 0.337, which indicates predictive relevance for every endogenous construct.

Mediation analysis

Because the direct legitimacy path proved non-significant, the mediation analysis carries the interpretive weight of the study. Table 5 reports the full set of indirect effects.

Table 5. Indirect, direct, and total effects

H	Indirect pathway	Estimate	SE	t	p	95% CI	Significant
H9	RL → TR → BI	0.133	0.034	3.948	< .001	[0.069, 0.202]	Yes
	RL → AU → BI	0.082	0.025	3.345	.001	[0.038, 0.134]	Yes
H10	RL → TR → AU → BI	0.091	0.017	5.225	< .001	[0.059, 0.126]	Yes
	TR → AU → BI	0.166	0.030	5.491	< .001	[0.110, 0.229]	Yes
	IQ → RL → TR	0.211	0.030	6.935	< .001	[0.152, 0.271]	Yes
	ST → RL → TR	0.155	0.030	5.222	< .001	[0.100, 0.215]	Yes
	RL → BI total indirect	0.306					Yes
	RL → BI direct	0.070					No
	RL → BI total effect	0.376					Yes

Note. Estimates derive from 5,000 bootstrap resamples with percentile confidence intervals. Significant indirect effects alongside a non-significant direct effect indicate an indirect-only mediation pattern.

Every indirect pathway tested is positive and statistically significant with a bootstrap confidence interval excluding zero. The single-mediator path through trust returns an indirect effect of 0.133, and the path through authority returns 0.082. The serial path running from legitimacy through trust and authority to intention returns 0.091 with a confidence interval from 0.059 to 0.126, which supports hypothesis ten. The total indirect effect of legitimacy on intention is 0.306 against a non-significant direct effect of 0.070, giving a total effect of 0.376.

This configuration corresponds to an indirect-only mediation pattern in the established typology, in which a significant indirect effect accompanies a non-significant direct effect (Rasoolimanesh et al., 2021). The substantive reading is that religious legitimacy is a genuine antecedent of behavioural intention, since its total effect is substantial, but that none of its influence reaches intention except through trust and perceived authority. Judging machine-generated religious guidance acceptable does not, on this evidence, translate into an intention to consult it again unless that judgement first becomes reliance and then deference. The two antecedent constructs behave consistently with this mediated account. The indirect effect of information quality on trust through legitimacy reaches 0.211, and the corresponding effect of source transparency reaches 0.155, so both perceived attributes transmit their influence forward through the legitimacy judgement rather than bypassing it. Because these structural results are reported for the pooled sample across all five institutional categories without a preceding measurement-invariance test, the interpretation that follows assumes the six constructs carry the same meaning in Islamic-affiliated and general university settings; this pooling assumption is examined further in the Limitations section below.

Discussion

What predicts legitimacy

The first research question asked which perceived attributes of machine-generated religious content predict judgements of religious legitimacy, and the answer is that both tested attributes matter while content quality matters more. The coefficient for information quality exceeds that for source transparency by a third, and its effect size is nearly double. This ordering is consistent

with evidence from general information systems research, where information quality has been shown to drive source trust and continuance intention for generative systems (Mun & Hwang, 2025). It also carries a warning that the transparency literature has already identified. Studies of citation in machine-generated answers have found that the presence of citations raises trust while actually checking those citations lowers it, which indicates that transparency features can function as credibility signals rather than as verification aids (Ding et al., 2025; Schiedermair et al., 2026). Related experimental work shows that the way numerical and source information is presented shifts trust independently of the underlying accuracy (Kim et al., 2021), and that disclosing system accuracy changes reliance in ways users do not always anticipate (de Jong et al., 2025). Transparency, in other words, is not self-executing, and its effect depends on whether the user treats a disclosure as an invitation to check or as a reassurance that checking is unnecessary.

The implication for religious settings is an uncomfortable one. If students weigh perceived accuracy and clarity more heavily than traceability when deciding whether machine guidance could be Islamically acceptable, then a fluent and confident answer with weak sourcing will attract a favourable legitimacy judgement more readily than a hedged answer with careful attribution. This inverts the ordering that Islamic legal epistemology applies, in which the chain of transmission and the qualification of the transmitter are prior to the persuasiveness of the content (Ahmed et al., 2026). The gap between the two orderings is a precise target for religious literacy instruction rather than a general concern about technology.

How legitimacy reaches conduct

The second research question asked whether legitimacy influences behavioural intention directly or through trust and authority, and the evidence supports the mediated account without qualification. Figure 3 sets out the interpretation that follows.

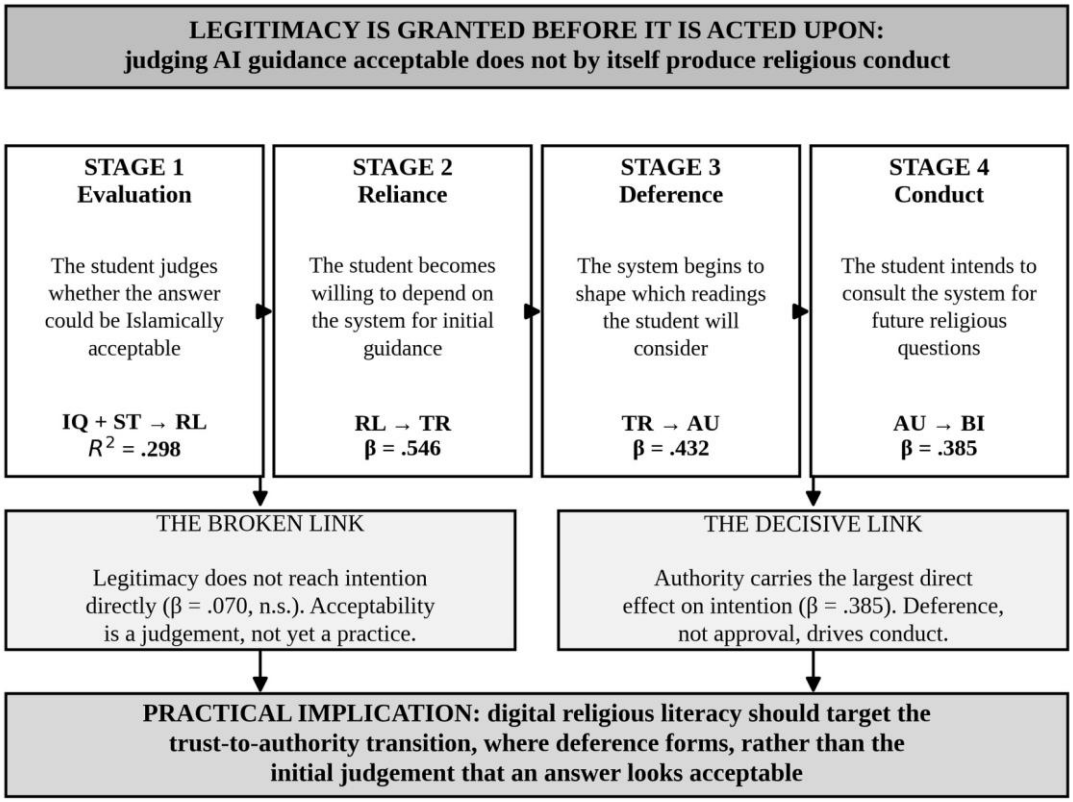


Figure 3. Four-stage interpretation of the pathway from legitimacy judgement to religious information conduct

The figure distinguishes four positions that a student may occupy. At the evaluation stage the student judges whether an answer could be Islamically acceptable, and that judgement is predicted jointly by content quality and traceability. At the reliance stage the student becomes willing to depend on the system for initial guidance, which is the strongest transition in the model. At the deference stage the system begins to shape which interpretations the student considers, which is where standing rather than mere usefulness is conceded. Only at the conduct stage does the student form an intention to consult the system again for religious questions. This four-stage account is an interpretive overlay constructed to describe the obtained pattern of significant and non-significant paths in the serial mediation chain that was specified and tested before data collection, rather than an independently derived and separately tested stage theory; the four positions track the pre-specified path sequence (legitimacy to trust to authority to intention) rather than constituting new evidence for discrete psychological stages.

The non-significant direct path is the theoretically interesting result rather than a null finding to be explained away. Legitimacy in this model behaves as a distal evaluative judgement rather than a behavioural disposition, and its substantial total effect of 0.376 travels entirely through the intermediate states. This distinguishes religious legitimacy from the perceived usefulness constructs of acceptance research, which typically show direct effects on intention (Habibi et al., 2024; Sun et al., 2025). It is also worth noting that the acceptance studies cited for comparison typically specify simpler two-step models (usefulness predicting intention directly) rather than the four-construct serial mediation chain estimated here, so the non-significant direct effect observed in the present model may partly reflect how much variance upstream constructs already absorb in a longer causal chain, rather than reflecting only what legitimacy means as a construct. The distinction supports treating legitimacy as a

construct in its own right rather than as a religiously flavoured relabelling of perceived usefulness, which addresses the theoretical concern raised in the introduction.

Why authority rather than trust drives intention

The third research question asked which construct carries the largest direct effect on intention. Perceived authority does so, at 0.385 against 0.244 for trust, and this ordering has implications that reach beyond the present sample. Trust research has established that trust predicts acceptance and reliance across domains (Choung et al., 2023; Kaplan et al., 2023). The present results suggest that in religious settings trust is necessary but not sufficient, because its influence on intention is roughly matched by a further transition in which the system acquires standing to shape what the student considers.

That transition is what scholarship on digital religion has been describing qualitatively. Accounts of authority in algorithmically mediated environments argue that platform prominence and institutional recognition jointly determine which expressions count as authoritative (Akhmetova et al., 2026), and studies of religious learning have introduced digital epistemic authority to name the standing that mediated sources acquire (Berger & Golan, 2024). Work on the attribution of epistemic authority to artificial systems in secular settings has found comparable patterns of deference (Milella & Cabitza, 2026). Experimental measurement of perceived legitimacy in algorithmic governance provides a further parallel, since citizens there also distinguish between accepting that an automated decision is defensible and accepting that the system is entitled to make it (Hillo et al., 2026). The religious case appears to follow the same two-step structure while raising the stakes, because the entitlement at issue concerns the interpretation of revelation rather than the administration of a public service. The present study contributes a measurement result to that conversation, namely that authority and trust are empirically separable, that authority sits downstream of trust, and that authority rather than trust is the proximate predictor of religious information conduct. The claim that authority sits downstream of trust reflects the a priori structural specification, in which the trust-to-authority path was modelled on the reasoning that deference follows reliance, rather than a direction independently established against a rival ordering; a cross-sectional design estimated under this specification cannot by itself distinguish it from the reverse direction or from a non-causal association, which is the same consideration noted in the Limitations regarding the legitimacy-trust-authority sequence.

A further observation from the descriptive data sharpens this point. Perceived authority recorded the lowest construct mean in the sample at 3.20, while behavioural intention recorded the highest at 3.51. Respondents were therefore more willing to keep consulting these systems than to concede that the systems shape their religious understanding. The gap is consistent with the pattern documented among Muslim American participants, who voiced clear reservations about artificial intelligence while preferring its answers to scholar-authored ones in blind comparison (Alam et al., 2025). Stated reservation and actual reliance appear to diverge in this domain, and self-report of authority may understate the influence that the structural model detects.

Implications for education and institutions

If deference rather than approval drives conduct, then religious literacy programmes aimed at the initial judgement will intervene at the wrong point. Telling students that machine-generated guidance requires verification addresses the evaluation stage, where the present model already shows students applying reasonable criteria. The consequential transition occurs later, between reliance and deference, and it is there that the ordinary practices of Islamic knowledge transmission have most to offer. Research on Indonesian religious organisations documents both the migration of religious learning online and the authority vacuum that migration creates

(Akmaliah, 2022; Syafaq et al., 2026). Studies of students in Islamic higher education report that learners themselves raise the absence of a valid chain of transmission when questioning machine guidance (Darmu'in et al., 2025), which suggests that the conceptual resources for such instruction are already present among students. What appears to be missing is not the concept but its application at the moment when reliance is forming. Evidence that curricular autonomy becomes superficial without explicit critical literacy points in the same direction (Afzainizam & Rahman, 2026; Haditia et al., 2026; Ahmad et al., 2025), as does the finding that users readily assign these systems the role of a perceived religious scholar while remaining uneasy about doing so (Jahan et al., 2026; Fatmawati, Musyarofah, et al., 2026).

For institutions the findings support a specific division of labour. Because information quality drives legitimacy most strongly, technical improvements in accuracy and retrieval grounding will raise legitimacy judgements whether or not students become better verifiers. Because authority is the proximate determinant of conduct, institutional response cannot rest on technical improvement alone and must address the social transition through which students concede standing. Work comparing machine-generated rulings with the collective deliberative processes of Indonesian religious organisations illustrates what that response might draw upon (Faizin, Maarif, et al., 2025; Tawakkal et al., 2026). Research on digitally mediated religious authority among Indonesian women scholars shows that platform dynamics can amplify some positions while confining others to narrow audiences (Rochmawati, 2026), which indicates that institutional presence online is necessary but does not guarantee reach. The broader lesson from digital religion scholarship applies directly, namely that authority is renegotiated rather than lost when religious life moves into mediated environments (Campbell & Bellar, 2023).

Limitations

Four limitations qualify the conclusions that are drawn above. The design is cross-sectional, so the sequence from legitimacy through trust and authority to intention is a theoretical ordering supported by model fit rather than a demonstrated temporal sequence, and longitudinal or experimental work is needed to establish direction. The sample is confined to undergraduates in one metropolitan region of Indonesia, so generalisation to older Muslims, to rural populations, or to other national contexts is not warranted without replication. The legitimacy and authority instruments were developed for this study, and although they performed well psychometrically they require validation in independent samples before they can be treated as established scales. Finally, behavioural intention is a stated disposition rather than observed behaviour, and the divergence between reported authority and reported intention noted above suggests that self-report may not capture the full influence at work. Two further analyses would strengthen future applications of this model. Measurement invariance testing across the five institutional categories would establish whether the constructs carry the same meaning in Islamic and general university settings (Liengaard, 2024). Because so much of the Discussion's interpretation depends on the specific legitimacy-to-trust-to-authority ordering being correct, the competing-model comparison recommended next is flagged here as a priority for a follow-up analysis of the present dataset rather than as an open-ended direction for future research in general. Treating the mediation result as a comparison between competing model specifications rather than as a single test would also place the indirect-only claim on firmer ground (Sarstedt & Moisescu, 2024).

Conclusion

This study modelled how 320 Muslim undergraduates in Greater Jakarta come to treat generative artificial intelligence as a source of religious guidance, and it reached four conclusions. First, judgements of religious legitimacy are predicted by both perceived information quality and perceived source transparency, with content quality carrying the

stronger influence, which means that persuasiveness currently outweighs traceability in how students evaluate machine-generated religious content. Second, religious legitimacy does not influence behavioural intention directly, since its direct coefficient was not significant while its total indirect effect through trust and perceived authority reached 0.306, which establishes an indirect-only mediation pattern. Third, perceived religious authority rather than trust is the proximate determinant of intention, carrying the largest direct effect in the model, which indicates that deference rather than approval governs religious information conduct. This statement should be read with the same caution given in the Limitations section, where the legitimacy-trust-authority sequence is described as a theoretical ordering supported by model fit rather than a demonstrated temporal sequence; firmer support for this conclusion awaits the competing-model comparison recommended there. Fourth, the constructs of legitimacy, trust, and authority are empirically distinguishable, which supports treating religious legitimacy as a construct in its own right rather than as a variant of perceived usefulness.

Taken together these findings locate the decisive moment later in the process than religious literacy programmes typically assume. Students apply defensible criteria when they first evaluate machine-generated guidance, and the risk does not lie principally in that evaluation. It lies in the subsequent transition through which reliance becomes deference, because that is where a system without qualification or accountability begins to shape which interpretations a student will weigh.

Three directions would advance this work most directly. Longitudinal and experimental designs should test whether the ordering proposed here holds over time, particularly whether trust precedes deference as the model assumes. Replication should extend beyond the metropolitan undergraduate population to older Muslims, to pesantren settings, and to other Muslim-majority and Muslim-minority contexts, since the legitimacy criteria that students apply may not be shared across these groups. Instrument development should continue, so that religious legitimacy and perceived religious authority acquire validated multi-sample measurement properties and can be carried into the comparative research that this field now requires.

Acknowledgement

The authors thank the students who participated in the survey, the campus organisations that assisted with distribution, and the panel members who reviewed the draft instrument for doctrinal appropriateness and clarity.

Declaration of Generative AI Use

The authors used Claude, an AI assistant developed by Anthropic, during the preparation of this manuscript. Claude was used to assist with statistical computation and verification of the reported estimates, to assist with literature searching and the verification of bibliographic metadata, to assist with preparing the tables and figures, and to assist with language editing of the draft, and to typeset the manuscript into the journal template. The research design, data collection, instrument development, interpretation of results, and all substantive argumentation are the responsibility of the authors. The authors verified every reported statistic and every cited source, reviewed and edited the full manuscript, and take full responsibility for its content.

Authors' Contribution Statement

Umi Musyarrofah: conceptualization, research design, development of the hypothesized model, formal analysis with PLS-SEM, and writing of the original draft.
Kalsum Minangsih: instrument development, data collection, data curation, and writing (review and editing).

Athia Nur Kamilah: literature review, validation of the measurement model, interpretation of results, and writing (review and editing).
All authors contributed to the discussion of findings, read, and approved the final manuscript.

References

Abdelnour, M. G. (2025). Artificial intelligence and the Islamic theology of technology: From "means" to "meanings" and from "minds" to "hearts". Religions, 16(6), Article 796. https://doi.org/10.3390/rel16060796

Afzainizam, M., & Rahman, M. F. (2026). Constructing digital religious agency: Generative AI literacy and Islamic learning within autonomous education frameworks. Al-Tadzkiyyah: Jurnal Pendidikan Islam, 17(1), 16–32. https://doi.org/10.24042/atjpi.v17i1.30279

Ahmad, Y. S., Abdulrasaq, Y., Ali, A. M., & Ali, F. (2025). Integrating Islamic ethical principles into AI education and character building in digital society. *Tunjuk Ajar: Journal of Education and Culture, 1*(2), 130–146. https://doi.org/10.64929/ta.v1i2.38

Ahmed, A., Masuwd, M., Sabir, M. R., & Mahmudulhassan, M. (2026). Can artificial intelligence perform ijtihad? Authority, epistemology, and legitimacy in digital Islamic jurisprudence. Suhuf: International Journal of Islamic Studies, 38(1), 154–169. https://doi.org/10.23917/suhuf.v38i1.16437

Akhmetova, A., Rakhimova, G., Onngar, Y., Issakhan, M., & Baitenova, N. (2026). Digital religion and hybrid religious authority in algorithmically mediated environments. Frontiers in Sociology, 11, Article 1892383. https://doi.org/10.3389/fsoc.2026.1892383

Akmaliah, W. (2022). Reclaiming moderate Islam in Nahdlatul Ulama: Challenging the dominant religious authority in digital platform. Journal of Indonesian Islam, 16(1), 223–248. https://doi.org/10.15642/jiis.2022.16.1.223-248

Alam, S. M., Abdulhai, M., & Salehi, N. (2025). Blind faith? User preference and expert assessment of AI-generated religious content. In Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (pp. 2410–2424). Association for Computing Machinery. https://doi.org/10.1145/3715275.3732162

Atallah, F. A. (2026). Digital mediation and fatwa authority in contemporary Islam: A critical Islamic legal and media-theoretical framework. Religions, 17(3), Article 350. https://doi.org/10.3390/rel17030350

Ayalon, A., & Ido, O. (2026). Adoption of ChatGPT in higher education: A UTAUT based comparative study of students and lecturers. Journal of Educational Technology Systems, 55(1), 6–39. https://doi.org/10.1177/00472395261439145

Becker, J.-M., Cheah, J.-H., Gholamzade, R., Ringle, C. M., & Sarstedt, M. (2023). PLS-SEM's most wanted guidance. International Journal of Contemporary Hospitality Management, 35(1), 321–346. https://doi.org/10.1108/IJCHM-04-2022-0474

Bekturova, M., Tulepova, S., & Zhaitapova, A. (2025). Predicting Kazakhstani TEFL students' continuance intention towards using ChatGPT in academic writing. Education and Information Technologies, 30(9), 11889– 11909. https://doi.org/10.1007/s10639-024-13306-6

Berger, A., & Golan, O. (2024). Online religious learning: Digital epistemic authority and self-socialization in religious communities. Learning, Media and Technology, 49(2), 274–289. https://doi.org/10.1080/17439884.2023.2169833

Bučinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. Proceedings of the ACM on Human-Computer Interaction, 5(CSCW1), 1–21. https://doi.org/10.1145/3449287

Campbell, H. A., & Bellar, W. (2023). Digital religion: The basics. Routledge.

Campbell, H. A., & Tsuria, R. (Eds.). (2022). Digital religion: Understanding religious practice in digital media (2nd ed.). Routledge.

Cheah, J.-H., Nitzl, C., Roldán, J. L., Cepeda-Carrión, G., & Gudergan, S. P. (2021). A primer on the conditional mediation analysis in PLS-SEM. ACM SIGMIS Database: The DATABASE for Advances in Information Systems, 52(SI), 43–100. https://doi.org/10.1145/3505639.3505645

Choung, H., David, P., & Ross, A. (2023). Trust in AI and its role in the acceptance of AI technologies. International Journal of Human–Computer Interaction, 39(9), 1727–1739. https://doi.org/10.1080/10447318.2022.2050543

Darmu'in, D., Nasikhin, N., Darnoto, D., Sofanudin, A., Maarif, M. A., & Haryanto, J. T. (2025). Perceived barriers to ChatGPT integration in Islamic studies education: Insights from Malaysian international students in Indonesia. Journal of International Students, 15(11), 141–164. https://doi.org/10.32674/h99a4p06

de Jong, S., Bos, M. W., van Berkel, N., & Lamers, M. H. (2025). Algorithm appreciation or aversion: The effects of accuracy disclosure on users' reliance on algorithmic suggestions. Behaviour & Information Technology. Advance online publication. https://doi.org/10.1080/0144929X.2025.2535732

Ding, Y., Facciani, M., Joyce, E., Poudel, A., Bhattacharya, S., Veeramani, B., Aguinaga, S., & Weninger, T. (2025). Citations and trust in LLM generated responses. Proceedings of the AAAI Conference on Artificial Intelligence, 39(22), 23787–23795. https://doi.org/10.1609/aaai.v39i22.34550

Erman, E., & Furendal, M. (2024). Artificial intelligence and the political legitimacy of global governance. Political Studies, 72(2), 421–441. https://doi.org/10.1177/00323217221126665

Faizin, N., Alfan, M., Basid, A., Ramadhan, M. R., Panatik, S. A., & Kawakip, A. N. (2025). Muslim students' acceptance of artificial intelligence in Islamic religious education: An extended TAM approach. Discover Education, 4(1), Article 304. https://doi.org/10.1007/s44217-025-00767-1

Faizin, N., Maarif, A. S., & Hanafi, Y. (2025). Considering religious moderation in Islamic law through AI ChatGPT and Bahsul Masail of Nahdlatul Ulama. Law, Innovation and Technology, 17(1), 271–288. https://doi.org/10.1080/17579961.2025.2469351

Fatmawati, Khoiriyah, N., Fauzan, Syam, A. R., Zh, M. H. R., Fauzi, A., & Ghani, M. F. A. (2026). Daʿwah in the digital age: Bridging al-Ghazālī's classical daʿwah pedagogy and digital television adoption among Indonesian Muslims. Journal of Islamic Thought and Civilization, 16(2). https://doi.org/10.32350/jitc.162.13

Fatmawati, F., Musyarofah, U., & Minangsih, K. (2026). Digital da'wah content in Islamic religious education: Critical reception and curriculum integration in Banten senior high schools. Reslaj: Religion Education Social Laa Roiba Journal, 8(7), 900–918. https://doi.org/10.47467/reslaj.v8i7.13058

Grimmelikhuijsen, S., & Meijer, A. (2022). Legitimacy of algorithmic decision-making: Six threats and the need for a calibrated institutional response. Perspectives on Public Management and Governance, 5(3), 232–242. https://doi.org/10.1093/ppmgov/gvac008

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

Habibi, A., Mukminin, A., Octavia, A., Wahyuni, S., Danibao, B. K., & Wibowo, Y. G. (2024). ChatGPT acceptance and use through UTAUT and TPB: A big survey in five Indonesian universities. Social Sciences & Humanities Open, 10, Article 101136. https://doi.org/10.1016/j.ssaho.2024.101136

Haditia, M., Abubakar, I. Y., Rokimin, & Fauzi, A. (2026). Design and validation of a quality assurance framework for AI educational technology in low-resource settings. *Educational Technology in Developing Countries, 1*(1), 83–102. https://doi.org/10.64929/etdc.v1i1.68

Hair, J. F., & Alamer, A. (2022). Partial least squares structural equation modeling (PLS-SEM) in second language and education research: Guidelines using an applied example. Research Methods in Applied Linguistics, 1(3), Article 100027. https://doi.org/10.1016/j.rmal.2022.100027

Hair, J. F., Hult, G. T. M., Ringle, C. M., & Sarstedt, M. (2022). A primer on partial least squares structural equation modeling (PLS-SEM) (3rd ed.). Sage.

Hair, J. F., Sarstedt, M., Ringle, C. M., Sharma, P. N., & Liengaard, B. D. (2024). Going beyond the untold facts in PLS-SEM and moving forward. European Journal of Marketing, 58(13), 81–106. https://doi.org/10.1108/EJM-08-2023-0645

Hillo, J., Vento, I., & Erkkilä, T. (2026). Algorithmic governance: Experimental evidence on citizens' and public administrators' legitimacy perceptions of automated decision-making. Public Administration, 104(3), 536– 551. https://doi.org/10.1111/padm.70028

Jahan, I., Kamruzzaman, S., Hossen, K. M. N., Nurain, N., & Islam, A. B. M. A. A. (2026). From mind mapping to dream interpretation: Understanding Muslim users' engagement with GenAI tools in religious context. In Human-computer interaction (Lecture Notes in Computer Science, Vol. 16702, pp. 353–372). Springer. https://doi.org/10.1007/978-3-032-29586-6_23

Jussupow, E., Benbasat, I., & Heinzl, A. (2024). An integrative perspective on algorithm aversion and appreciation in decision-making. MIS Quarterly, 48(4), 1575–1590. https://doi.org/10.25300/MISQ/2024/18512

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, J., Giroux, M., & Lee, J. C. (2021). When do you trust AI? The effect of number presentation detail on consumer trust and acceptance of AI recommendations. Psychology & Marketing, 38(7), 1140–1155. https://doi.org/10.1002/mar.21498

Kock, N. (2021). Harman's single factor test in PLS-SEM: Checking for common method bias. Data Analysis Perspectives Journal, 2(2), 1–6. https://scriptwarp.com/dapj/2021_DAPJ_2_2/Kock_2021_DAPJ_2_2_HarmansCMBTest.pdf

Kurniawan, E., Ahmad, H., & Muttaqin, M. Z. (2025). From content creators to zakat payers: The fatwa of Indonesian Ulema Council and the rise of digital fiqh. *Islamic Law and Social Issues in Society, 1*(2), 99–115. https://doi.org/10.64929/ilsiis.v1i2.9

Liengaard, B. D. (2024). Measurement invariance testing in partial least squares structural equation modeling. Journal of Business Research, 177, Article 114581. https://doi.org/10.1016/j.jbusres.2024.114581

Liengaard, B. D., Sharma, P. N., Hult, G. T. M., Jensen, M. B., Sarstedt, M., Hair, J. F., & Ringle, C. M. (2021). Prediction: Coveted, yet forsaken? Introducing a cross-validated predictive ability test in partial least squares path modeling. Decision Sciences, 52(2), 362–392. https://doi.org/10.1111/deci.12445

Lohlker, R., & Wahid, S. H. (2026). The transformation of Islamic religious authority. Religions, 17(4), Article 493. https://doi.org/10.3390/rel17040493

Mahendra, B. A., Munandar, A., & Nurzuliana, Z. (2026). Mahar as esteem, not price: Mediating classical fiqh, the Compilation of Islamic Law, and custom in Indonesia's plural legal order. *PFL: Plural Family Law Review, 1*(1), 59–80. https://www.ejournaltuah.id/index.php/pfl/article/view/13

Mahmud, H., Islam, A. K. M. N., Ahmed, S. I., & Smolander, K. (2022). What influences algorithmic decision-making? A systematic literature review on algorithm aversion. Technological Forecasting and Social Change, 175, Article 121390. https://doi.org/10.1016/j.techfore.2021.121390

Malhotra, P., & Shin, S. Y. (2025). Understanding information credibility evaluation on bounded social media places: A mixed methods study. Communication Monographs, 92(3), 653–676. https://doi.org/10.1080/03637751.2025.2455714

Milella, F., & Cabitza, F. (2026). Perceiving AI as an epistemic authority or algority: A user study on the human attribution of authority to AI. Machine Learning and Knowledge Extraction, 8(2), Article 36. https://doi.org/10.3390/make8020036

Mun, I. B., & Hwang, K.-H. (2025). Understanding ChatGPT continuous usage intention: The role of information quality, information usefulness, and source trust. Information Development, 41(3), 675–691. https://doi.org/10.1177/02666669241307595

Noh, M. A. C., Basri, M., & Panaemalae, A. (2026). Ethical leadership: Reframing moral complexity and power through a philosophical anthropological perspective in organizational practice. *Leadership in Muslim Societies, 1*(1), 19–34. https://doi.org/10.64929/lms.v1i1.57

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(1), 17–61. https://doi.org/10.1146/annurev-orgpsych-110721-040030

Ramadhan, M. A. I., Syam, A. R., & Barkah, S. A. M. (2026). Customay judges in digital space: Authority contestations, living fiqh, and legal pluralism in Bugis customary inheritance dispute resolution via WhatsApp groups. *Tuah Customary Law Review, 1*(1), 1–20. https://ejournal.tuah.or.id/index.php/tclr/article/view/74

Rasoolimanesh, S. M., Wang, M., Roldán, J. L., & Kunasekaran, P. (2021). Are we in right path for mediation analysis? Reviewing the literature and proposing robust guidelines. Journal of Hospitality and Tourism Management, 48, 395–405. https://doi.org/10.1016/j.jhtm.2021.07.013

Ringle, C. M., Wende, S., & Becker, J.-M. (2024). SmartPLS 4 [Computer software]. SmartPLS. https://www.smartpls.com

Rochmawati, I. (2026). Digitalizing female religious authority: KUPI's gender-responsive fatwas and the reconfiguration of Islamic discourse in Indonesia. Al-Tahrir: Jurnal Pemikiran Islam, 26(1). https://doi.org/10.21154/altahrir.v26i1.12384

Roemer, E., Schuberth, F., & Henseler, J. (2021). HTMT2 – An improved criterion for assessing discriminant validity in structural equation modeling. Industrial Management & Data Systems, 121(12), 2637–2650. https://doi.org/10.1108/IMDS-02-2021-0082

Rusli, R., & Nurdin, N. (2022). Understanding Indonesia millennial ulama online knowledge acquisition and use in daily fatwa making habits. Education and Information Technologies, 27(3), 4117–4140. https://doi.org/10.1007/s10639-021-10779-7

Sabol, M., Hair, J. F., Cepeda-Carrión, G., Roldán, J. L., & Chong, A. Y. L. (2023). PLS-SEM in information systems: Seizing the opportunity and marching ahead full speed to adopt methodological updates. Industrial Management & Data Systems, 123(12), 2997–3017. https://doi.org/10.1108/IMDS-07-2023-0429

Sarstedt, M., Hair, J. F., Pick, M., Liengaard, B. D., Radomir, L., & Ringle, C. M. (2022). Progress in partial least squares structural equation modeling use in marketing research in the last decade. Psychology & Marketing, 39(5), 1035–1064. https://doi.org/10.1002/mar.21640

Sarstedt, M., Hair, J. F., & Ringle, C. M. (2023). "PLS-SEM: Indeed a silver bullet" – Retrospective observations and recent advances. Journal of Marketing Theory and Practice, 31(3), 261–275. https://doi.org/10.1080/10696679.2022.2056488

Sarstedt, M., & Moisescu, O.-I. (2024). Quantifying uncertainty in PLS-SEM-based mediation analyses. Journal of Marketing Analytics, 12(1), 87–96. https://doi.org/10.1057/s41270-023-00231-9

Schemmer, M., Kuehl, N., Benz, C., Bartos, A., & Satzger, G. (2023). Appropriate reliance on AI advice: Conceptualization and the effect of explanations. In Proceedings of the 28th International Conference on Intelligent User Interfaces (pp. 410–422). Association for Computing Machinery. https://doi.org/10.1145/3581641.3584066

Schiedermair, I., Baumgartner, M., Kick, E., Kopp, T., & Kinkel, S. (2026). The effect of source disclosure on trust in an AI-based knowledge management system. International Journal of Human–Computer Interaction. Advance online publication. https://doi.org/10.1080/10447318.2026.2708454

Sharma, P. N., Liengaard, B. D., Hair, J. F., Sarstedt, M., & Ringle, C. M. (2023). Predictive model assessment and selection in composite-based modeling using PLS-SEM: Extensions and guidelines for using CVPAT. European Journal of Marketing, 57(6), 1662–1677. https://doi.org/10.1108/EJM-08-2020-0636

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

Sun, P., Li, L., Hossain, M. S., & Zabin, S. (2025). Investigating students' behavioral intention to use ChatGPT for educational purposes. Sustainable Futures, 9, Article 100531. https://doi.org/10.1016/j.sftr.2025.100531

Syafaq, H., Musyafa'ah, N. L., Al Banna, A. H., & Rosyidah, U. (2026). “Santri without pesantren” and the sectarian violence on Indonesian Muslim social media. Cogent Arts & Humanities, 13(1), Article 2612046. https://doi.org/10.1080/23311983.2025.2612046

Tawakkal, M. G., Kurniawan, M. I., & Faathima. (2026). Cash waqf linked sukuk and economic empowerment: A grounded theory from Indonesia. *Ethiconomics: Journal of Islamic Banking and Finance, 1*(01), 1–20. https://doi.org/10.64929/ethiconomics.v1i01.2

Tsuria, R., & Tsuria, Y. (2024). Artificial intelligence's understanding of religion: Investigating the moralistic approaches presented by generative artificial intelligence tools. Religions, 15(3), Article 375. https://doi.org/10.3390/rel15030375

Vaithilingam, S., Ong, C. S., Moisescu, O. I., & Nair, M. S. (2024). Robustness checks in PLS-SEM: A review of recent practices and recommendations for future applications in business research. Journal of Business Research, 173, Article 114465. https://doi.org/10.1016/j.jbusres.2023.114465

Vereschak, O., Bailly, G., & Caramiaux, B. (2021). How to evaluate trust in AI-assisted decision making? A survey of empirical methodologies. Proceedings of the ACM on Human-Computer Interaction, 5(CSCW2), 1–39. https://doi.org/10.1145/3476068

Wahid, S. H. (2025). Artificial intelligence in Islamic guidance: Assessing chatbot performance and jurisprudential adherence. Journal of Digital Islamicate Research, 3(1), 33–96. https://doi.org/10.1163/27732363-20250003

Whyte, S. A. (2022). Islamic religious authority in cyberspace: A qualitative study of Muslim religious actors in Australia. Religions, 13(1), Article 69. https://doi.org/10.3390/rel13010069

Zhou, T., & Ma, X. (2025). Examining generative AI user continuance intention based on the SOR model. Aslib Journal of Information Management, 78(4), 969–987. https://doi.org/10.1108/AJIM-08-2024-0620