JOURNAL OF POLICIES AND RECOMMENDATIONS

Volume: 5 Issue: 4

(October-December), 2026

ISSN: 2957-4129



KMF PUBLISHERS

www.kmf-publishers.com/jopar/

Contents lists available at KMF Publishers

KMF Journal of Policies and Recommendations

journal homepage: www.kmf-publishers.com/jopar/

Research Article

Public Trust in AI Governance and Support for Algorithmic Accountability in Bangladesh

Muhammad Anik Tanvir Faysal[1]; Shyma Islam[1]; Md. Mafuz[1]; Kazi Abdul Mannan[2]

[1] Department of Law; [2] Department of Business Administration
Shanto-Mariam University of Creative Technology, Dhaka, Bangladesh

Corresponding Author Muhammad Anik Tanvir Faysal, Email: anik420faysal@gmail.com

ARTICLE INFO

Keywords: algorithmic accountability; conditional trust; public governance; explanation and review; institutional trust; Bangladesh

ABSTRACT

Public support for algorithmic accountability can conceal substantial differences between agreement that organisations should answer for harm and support for particular safeguards. This study examines that distinction in a Bangladesh AI-governance database, separating conditional trust from the breadth of accountability endorsement. All 385 records enter the primary analysis. Three items address organisational accountability for significant harm, oversight of high-risk systems, and access to explanation and review. Eight response profiles and a zero-to-three breadth count preserve their distinct content. Analyses combine matched-response tests, respondent bootstrap intervals and threshold-specific logistic models with joint sandwich inference. Harm accountability is endorsed by 294 respondents (76.4%), compared with 215 (55.8%) for oversight and 208 (54.0%) for explanation and review. Only 93 (24.2%; 95% Wilson interval 20.1-28.7) endorse all three. Stronger agreement that effective accountability increases trust is associated with complete endorsement (adjusted odds ratio 1.55 per score point; 95% confidence interval 1.12-2.14; Holm-adjusted $p = 0.025$). Associations at lower breadth thresholds are uncertain; the joint three-threshold test yields $p = 0.069$, and slope heterogeneity is not established. Low block consistency supports item-level interpretation. Nineteen records have inconsistent experience/frequency responses; synthetic labels for identifiers and comments leave numerical provenance unresolved. Comments are excluded. The analysis identifies selective endorsement of accountability provisions rather than a validated general trust disposition. It proposes evaluating accountability through accessible oversight and review processes while distinguishing support for safeguards from evidence that existing governance is trustworthy.

Received: 20 September 2026 / Revised: 25 October 2026 / Accepted: 28 October 2026 / Published: 30 October 2026

Copyright: © 2026 by the authors. Licensee KMF Publishers (www.kmf-publishers.com). This open-access article is distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).
DOI: https://doi.org/10.64907/xkmf.v5i4.jopr.3

1 Introduction

Artificial intelligence increasingly enters public decisions through systems that classify information, recommend action and organise access to services. When a consequential decision involves AI, the accountability question extends beyond whether the output is accurate. People may need to know which organisation is responsible, how the system was assessed and whether an affected person can obtain an explanation or challenge an outcome. These questions concern the institutions surrounding a system as much as the system itself. Legal and public-administration scholarship accordingly treats algorithmic accountability as a problem of organisational responsibility, answerability and the conditions under which decisions can be examined.[1]

Public trust is relevant to these arrangements, but its meaning requires care. A person may trust an institution because they believe its decisions are competent and fair. Another may state that they would trust governance more if accountability were effective. The second statement identifies a condition for trust; it does not establish present confidence in a government, regulator or service provider. Treating the two as equivalent can produce an attractive but unsupported conclusion that public institutions already enjoy trust because respondents approve of accountability in principle.

A similar ambiguity concerns support for safeguards. Agreement that organisations should be accountable for significant AI harm is a broad normative judgement. Endorsement of oversight for high-risk systems and access to explanation and review concerns more specified arrangements. These positions can coexist, but they need not be equally salient to the same respondent. A mean score combining them may conceal whether someone supports only general responsibility, supports oversight without review, or endorses all three. The distribution of such combinations is therefore a substantive governance question rather than merely a measurement detail.

Accountability scholarship provides a reason to make that distribution visible. Bovens describes an accountability relationship involving an actor, a forum and the possibility of questioning, judgement and consequences. Algorithmic systems complicate those relationships when responsibility is distributed across developers, procurers and deploying organisations. Wieringa's review organises these concerns across a diverse literature. Both perspectives suggest asking what kind of accountability is supported and through which arrangement, rather than assuming that one positive response establishes a complete institutional preference.[2]

The empirical AI-trust literature also indicates that the relevant decision context matters. Aoki's experimental work examines

[1] Joshua A Kroll and others, 'Accountable Algorithms' (2017) 165 University of Pennsylvania Law Review 633; Madalina Busuioc, 'Accountable Artificial Intelligence: Holding Algorithms to Account' (2021) 81 Public Administration Review 825.

[2] Mark Bovens, 'Analysing and Assessing Accountability: A Conceptual Framework' (2007) 13 European Law Journal 447; Maranke Wieringa, 'What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability' in Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (ACM 2020) 1.

initial trust in public-sector chatbots, while Grimmelikhuijsen investigates how transparency affects perceived trustworthiness in automated decision-making. These studies give exposure or context a defined role. A questionnaire asking whether accountability increases trust addresses a different object: respondents' endorsement of a conditional relationship. Its interpretation should preserve that distinction when drawing comparisons with experimental findings.[3]

Bangladesh provides a relevant setting for examining these questions, without warranting assumptions about national opinion. UNESCO's 2025 account of an AI-ethics event in Bangladesh documents policy engagement with ethical and inclusive AI. Such engagement makes questions about public understanding and institutional safeguards pertinent. It does not establish how Bangladesh's population distributes across accountability preferences, nor does it validate the sample used here. The present article contributes a bounded analysis of recorded responses rather than a nationally representative estimate.[4]

The supplied Group 17 materials contain a public-trust and algorithmic-accountability block that is particularly suited to this distinction. Two items state that transparent rules or effective accountability increase trust. Three others address harm accountability, oversight, and explanation and review. The analysis retains these different roles. It focuses on the support combinations within the accountability block, while leaving ethical-risk support for national regulation and individual responsible-use behaviour to separate research questions. This scope gives the article a distinct institutional focus within the common database.

All 385 records enter the primary analysis because the topic concerns public attitudes across the respondent groups represented in the file. The database includes current law students, law graduates, legal professionals and other relevant respondents. It would be inappropriate to describe them all as students or as a random cross-section of Bangladesh's population. A student-only sensitivity analysis is useful for checking interpretation, but it does not define the main denominator for a public-governance topic.

The study asks three questions. First, how does endorsement differ between general harm accountability, high-risk oversight and explanation and review? Second, which combinations of these positions occur within respondents, and how common is endorsement of all three? Third, is stronger accountability-contingent trust associated with broader support after adjustment for transparent-rule conditional trust and recorded respondent characteristics? The third question is examined across cumulative breadth thresholds, allowing complete endorsement to be distinguished from support for at least one or two provisions.

[3] Naomi Aoki, 'An Experimental Study of Public Trust in AI Chatbots in the Public Sector' (2020) 37 Government Information Quarterly 101490; Stephan Grimmelikhuijsen, 'Explaining Why the Computer Says No: Algorithmic Transparency Affects the Perceived Trustworthiness of Automated Decision-Making' (2023) 83 Public Administration Review 241.

[4] UNESCO, 'AI for All: Shaping Bangladesh's Future with Ethical AI' (22 May 2025) <https://www.unesco.org/en/articles/ai-all-shaping-bangladeshs-future-ethical-ai> accessed 20 October 2026.

The contribution is an account of accountability support as a specified portfolio of positions. It avoids treating low-consistency items as a validated reflective trust scale and avoids interpreting a conditional belief as an observed institutional effect. The data's origin also remains a material boundary: the codebook labels identifiers and optional comments synthetic without documenting how the numerical responses were produced. Comments are excluded, and empirical interpretation remains conditional on provenance verification. The following sections develop the framework, describe the analysis, report the profile and threshold results, and consider what they imply for evaluating institutional safeguards.

2 Literature Review and Theoretical Foundation

This section connects institutional trust with accountability as a set of arrangements that can be supported separately. It distinguishes trust from trustworthiness, transparency from answerability, and broad normative approval from endorsement of a combination of safeguards. These distinctions guide the analytical model and limit the claims that the available questionnaire can sustain.

2.1 Trust and institutional trustworthiness

Trust research distinguishes a trusting party's expectations from the characteristics that make another actor worthy of trust. Mayer, Davis and Schoorman's organisational model separates trust from related constructs and its antecedents. Rousseau and colleagues identify common concerns across disciplinary approaches, including vulnerability and expectations about others' actions. These contributions make it important to specify whom a respondent trusts, in relation to what task and under which circumstances. A favourable view of an ethical principle cannot answer all those questions.[5]

Political and institutional trust further concern judgements about public organisations and their conduct. Levi and Stoker's review distinguishes citizens' evaluations of government from the institutional qualities those evaluations may reflect. For AI governance, the relevant trustee could be a regulator, a public agency, a contractor or an organisation deploying a system. A general questionnaire statement about AI governance does not identify one of those actors or establish a willingness to accept a specific decision.[6]

The Group 17 trust statements are expressly conditional. PTG1 connects trust with transparent, publicly understandable rules. PTG5 connects it with effective accountability. Agreement indicates

[5] Roger C Mayer, James H Davis and F David Schoorman, 'An Integrative Model of Organizational Trust' (1995) 20 Academy of Management Review 709; Denise M Rousseau and others, 'Not So Different After All: A Cross-Discipline View of Trust' (1998) 23 Academy of Management Review 393.

[6] Margaret Levi and Laura Stoker, 'Political Trust and Trustworthiness' (2000) 3 Annual Review of Political Science 475.

endorsement of those conditions as reasons for increased trust. It may reflect principled caution, prior experience or a general preference for governance safeguards. Those possible explanations are not separately measured. The article therefore uses conditional trust as a concise label for the recorded response, while avoiding claims about current confidence in institutions.

This interpretation is consistent with Jacovi and colleagues' distinction between the causes and goals of trust in AI. A governance arrangement should aim to make reliance appropriate to demonstrated capabilities and protections. Increasing expressed trust without improving institutional performance is not inherently a desirable outcome. The empirical question here is consequently whether respondents who endorse accountability as a trust condition also endorse a broader set of safeguards.[7]

2.2 Accountability as an institutional relationship

Algorithmic accountability cannot be reduced to releasing technical information. It involves questions about who must provide an account, to whom, at what point and with what consequences. Wieringa's review identifies the breadth of those questions, while Kroll and colleagues examine how technical and legal approaches can contribute to accountable automated processes. The implication for survey design is that general approval of accountability and support for particular arrangements should be observable separately.[8]

The first focal provision concerns organisational accountability for significant harm. It identifies an actor category and a reason for responsibility, but does not specify a forum, evidentiary standard or remedy. The second concerns oversight and assessment of high-risk systems. Its emphasis is preventive and supervisory. The third concerns affected persons' access to explanation and review. Its emphasis is the ability to obtain reasons and have a consequential decision examined. These descriptions interpret the supplied items; they do not claim that the items collectively cover every component of a legal accountability regime.

The distinction is useful because the functions can complement one another. Assessment may identify problems before

[7] Alon Jacovi and others, 'Formalizing Trust in Artificial Intelligence: Prerequisites, Causes and Goals of Human Trust in AI' in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (ACM 2021) 624.
[8] Wieringa (n 2); Kroll and others (n 1).

harm occurs, while review can address a particular contested outcome. Organisational accountability can provide a basis for assigning responsibility across those processes. Conversely, a person can favour one arrangement without understanding or endorsing the others. The questionnaire records such combinations without requiring the analyst to rank them as stages of moral development or legal sophistication.

Busuioc's analysis of accountable AI shows why the organisational setting remains central. Technical complexity and distributed responsibility can make formal answerability difficult to exercise. Raji and colleagues' auditing framework addresses processes across the development lifecycle, offering a complementary operational perspective. These approaches justify asking about oversight and review as concrete governance functions, while the survey cannot establish that respondents have experienced either process or that an institution has implemented them.[9]

2.3 Transparency and review provide different evidence

Transparency can make information available, but information availability does not guarantee understanding or effective challenge. Ananny and Crawford examine limitations of the transparency ideal in algorithmic accountability. Cobbe, Lee and Singh develop reviewability as a framework that attends to the broader decision process. Their work supports separating publicly understandable rules from opportunities to examine and contest a decision.[10]

Empirical trust studies offer a related caution. Grimmelikhuijsen's experiments distinguish aspects of algorithmic transparency and show that their implications for perceived trustworthiness vary with context. Shin examines explainability, causability, trust and acceptance from a user-understanding perspective. These studies do not imply that every explanation will produce greater trust, or that higher trust is proof of better accountability. The response depends on what is explained, how it is understood and the decision involved.[11]

Human review should also be specified rather than assumed effective. Green identifies problems with policies that invoke human oversight without adequate attention

[9] Busuioc (n 1); Inioluwa Deborah Raji and others, 'Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing' in Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (ACM 2020) 33.

[10] Mike Ananny and Kate Crawford, 'Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability' (2018) 20 New Media & Society 973; Jennifer Cobbe, Michelle Seng Ah Lee and Jatinder Singh, 'Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems' in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (ACM 2021) 598.

[11] Grimmelikhuijsen (n 3); Donghee Shin, 'The Effects of Explainability and Causability on Perception, Trust, and Acceptance: Implications for Explainable AI' (2021) 146 International Journal of Human-Computer Studies 102551.

to whether the proposed oversight works. Alon-Barkat and Busuioc examine automation bias and selective adherence in public-sector human-AI interaction. Together, these contributions challenge the idea that adding a human reviewer necessarily resolves accountability. For the present study, endorsement of oversight and review expresses a preference; it is not evidence that a particular safeguard is effective.[12]

2.4 Responsible AI and the limits of composite measures

Responsible-AI frameworks place accountability alongside other concerns rather than making it synonymous with trust. UNESCO's Recommendation on the Ethics of Artificial Intelligence includes responsibility, oversight, transparency and human-rights considerations. NIST's AI Risk Management Framework distinguishes characteristics of trustworthy systems and organises risk-management work through governance and related functions. These frameworks provide normative context, not a validated source for the questionnaire items or evidence that their requirements have been implemented in Bangladesh.[13]

Jobin, Ienca and Vayena's analysis of AI-ethics guidelines demonstrates why apparent agreement around broad principles does not settle their interpretation or implementation. Different organisations can use similar ethical language while assigning different institutional responsibilities. A survey that combines such language into a single favourable-attitude score may therefore exaggerate consensus about arrangements that have not been specified.[14]

The present study retains the five PTG items separately and constructs a transparent count of endorsement for PTG2-PTG4 only. The count measures how many of three named positions are endorsed. It does not assume they are interchangeable manifestations of one underlying trait. This approach follows the distinction between defining an index by its content and validating a reflective scale by its measurement properties. Equal counting is an analytical convention; the eight response profiles preserve information that the count alone would discard.[15]

Low internal consistency strengthens the case for avoiding a reflective latent model in these data. However, alpha cannot decide the validity of an item or establish that a content-defined count is substantively complete. Its role is diagnostic. The appropriate response is to narrow the interpretation, disclose the construction and show the component patterns, rather than to force the block into structural equation modelling.[16]

[12] Ben Green, 'The Flaws of Policies Requiring Human Oversight of Government Algorithms' (2022) 45 Computer Law & Security Review 105681; Saar Alon-Barkat and Madalina Busuioc, 'Human–AI Interactions in Public Sector Decision Making: "Automation Bias" and "Selective Adherence" to Algorithmic Advice' (2023) 33 Journal of Public Administration Research and Theory 153.

[13] UNESCO, Recommendation on the Ethics of Artificial Intelligence (UNESCO 2022); National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1, NIST 2023).

[14] Anna Jobin, Marcello Ienca and Effy Vayena, 'The Global Landscape of AI Ethics Guidelines' (2019) 1 Nature Machine Intelligence 389.

[15] Adamantios Diamantopoulos and Heidi M Winklhofer, 'Index Construction with Formative Indicators: An Alternative to Scale Development' (2001) 38 Journal of Marketing Research 269.

[16] Lee J Cronbach, 'Coefficient Alpha and the Internal Structure of Tests' (1951) 16 Psychometrika 297; Klaas Sijtsma, 'On the Use, the Misuse, and the Very Limited Usefulness of Cronbach's Alpha' (2009) 74 Psychometrika 107.

2.5 Theoretical expectations and analytical questions

Institutional trust reasoning motivates a positive relationship between accountability-contingent trust and support for a complete combination of safeguards. A respondent who regards effective accountability as a reason for trusting governance may also favour arrangements that make responsibility examinable and decisions reviewable. That reasoning concerns consistency between stated positions. It does not imply that trust causes policy support or that a governance reform would produce the same association.

The expected relationship is also vulnerable to common wording and evaluative tendencies. PTG5 expressly contains accountability language, which overlaps with the outcomes' normative domain. Respondents may interpret effective accountability as compensation after harm rather than oversight or review. A positive association with complete endorsement would therefore be informative but insufficient to validate a general institutional-trust mechanism. The article reports the association alongside component and threshold results to preserve that uncertainty.

H1 proposes that stronger endorsement of accountability-contingent trust is positively associated with endorsement of all three accountability provisions. This is an exploratory hypothesis developed for analysis of the supplied database, without preregistration. The full-profile outcome is the focal relationship; lower cumulative breadth thresholds provide a broader diagnostic. Multiplicity-adjusted results across all three thresholds are reported so that the strongest threshold is not presented without its analytical context.

A separate question concerns transparent-rule conditional trust. Theoretically, understandable rules could help people identify why oversight and review matter. They could also attract agreement without implying support for every safeguard. PTG1 therefore enters the adjusted model as a separate predictor. Its coefficients are assessed in a distinct exploratory family across the three thresholds. No mediation pathway from transparency through trust is estimated, because the instrument does not establish temporal order or independent measures for those stages.

The theory of planned behaviour supplies a further boundary. Ajzen distinguishes attitudes, subjective norms, perceived control and intentions in relation to a specified action. The present outcomes are policy-support judgements rather than observations of action, and the instrument lacks separate measures of those constructs. TPB therefore informs the need for specificity but is not claimed as a tested structural theory. Institutional trust and responsible-AI scholarship carry the principal explanatory and normative roles.[17]

[17] Icek Ajzen, 'The Theory of Planned Behavior' (1991) 50 Organizational Behavior and Human Decision Processes 179.

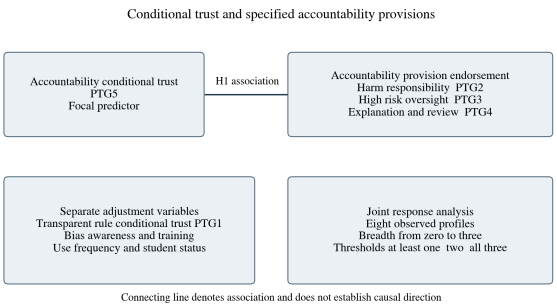

Figure 1. Proposed conceptual framework distinguishing conditional trust from accountability provisions. H1 concerns complete endorsement. The connecting line represents association, not a causal pathway.

Figure 1 represents an association-based design. Conditional-trust items remain distinct from accountability provisions; the breadth thresholds are alternative summaries of the same respondents' endorsement patterns. Connections in the figure do not establish causal direction. The framework is designed to make conceptual overlap visible and to prevent a favourable attitude towards accountability from being reported as proof of present institutional trustworthiness.

3 Methodology

The analysis uses the complete Group 17 database to examine how respondents combine support for accountability provisions and how those combinations relate to conditional trust. The methodology distinguishes the intended survey design from procedures documented in the available materials. It specifies item roles, denominators, multiplicity and data-quality boundaries before presenting the results.

3.1 Design and analytical population

The design is a cross-sectional analysis of the supplied Group-17-Law-SDG-Database.xlsx, interpreted alongside the Group 17 questionnaire and research plan. The unit of analysis is one record. All N = 385 records enter the primary analysis because the question concerns governance attitudes across the recorded respondent groups. The sample contains 133 current law students, 59 law graduates, 47 advocates or legal professionals and 146 other relevant respondents. The last category is not described more precisely in the source, limiting population interpretation.

The questionnaire records adult eligibility, consent and AI awareness for every row. These recorded indicators are not independent verification of administration. Institutional identifiers, collection dates and an achieved sampling frame are unavailable. The analysis consequently does not infer a nationally representative public-opinion distribution, calculate a response rate or apply weights based on undocumented inclusion probabilities. STROBE informs reporting of these design limits.[18]

[18] Erik von Elm and others, 'The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) Statement: Guidelines for Reporting Observational Studies' (2007) 4 PLOS Medicine e296.

3.2 Sample planning and source verification

The research plan specifies approximately 385 completed responses using the conventional large-population calculation: n■ = z²p(1−p)/e², with z = 1.96, p = 0.50 and e = 0.05. The result is 384.16, rounded up to 385. The preferred probability-sampling approach and approximately 428 planned contacts are planning statements; the materials do not establish that either was implemented. The achieved record count matches the target but does not establish the target's assumed precision.[19]

The codebook labels identifiers a synthetic respondent sequence and optional comments synthetic contextual comments. It does not document how the substantive numerical responses were generated. The comments are excluded. Numerical provenance remains unresolved, so empirical interpretations require confirmation against original collection records. The questionnaire recommends a pilot with 30-50 participants, but no pilot dataset, revision log or implementation record is supplied; no completed pilot is claimed.

3.3 Measures and response profiles

PTG5 is the focal predictor, scored one through five and interpreted as agreement that effective accountability increases trust. PTG1 separately measures trust conditional on transparent, understandable rules. PTG2-PTG4 are the three accountability provisions. Endorsement means a score of four or five; non-endorsement includes scores one through three and therefore includes neutrality. The coding must not be interpreted as opposition to accountability or disregard for rights.

Table 1. Measurement content and analytical roles

Code	Content abbreviated	Analytical role
PTG5	Accountability increases trust	Focal conditional-trust predictor
PTG1	Transparent rules increase trust	Separate conditional-trust predictor
PTG2	Organisations accountable for harm	Provision endorsement indicator
PTG3	High-risk oversight and assessment	Provision endorsement indicator
PTG4	Access to explanation and review	Provision endorsement indicator
B = E■ + E■ + E■	Number of provisions endorsed	Content-defined breadth count
AIL3 / D8 / D7 / D3	Bias awareness / training / frequency / status	Recorded adjustment variables

Source: supplied Group 17 questionnaire and codebook. Theory informs interpretation; validated item adaptation is not established. Endorsement means a score of four or five. The breadth count is not a

[19] William G Cochran, Sampling Techniques (3rd edn, John Wiley & Sons 1977).

reflective scale.

Table 1 records the analytical distinctions. The breadth count is B = E■ + E■ + E■, where each E equals one for endorsement and zero otherwise. B ranges from zero to three. The complete profile has B = 3. Each of the eight binary patterns is also reported so that substantively different positions with the same count remain distinguishable. This is a content-defined descriptive index, not a validated latent scale.

Five additional predictors accompany PTG5: PTG1, awareness of AI bias or incomplete data (AIL3), any formal AI training, AI-use frequency and current-law-student status. Item scores are centred at three. Training combines orientation and course/certified training versus none. Frequency is coded zero for never through four for daily or almost daily, centred at two. Student status contrasts current law students with other statuses; a sensitivity model separates all four categories. The models estimate conditional associations rather than identified causal effects.

3.4 Data checks and descriptive inference

All twenty substantive items are complete and within the one-to-five range. Twenty missing entries concern only AI-tool identification, which is not modelled. There are no duplicate identifiers, exact duplicate twenty-item vectors or complete twenty-item straight-line responses. Sixteen records report never using AI in the experience field while indicating some use in the frequency field. Three different records report never in frequency despite non-never experience. These inconsistent descriptions are retained in the primary attitude analysis and examined through sensitivity exclusion rather than silently corrected.

Cronbach's alpha is calculated for each five-item block as a diagnostic. Low consistency and heterogeneous wording do not support a reflective SEM analysis. Item proportions and response-profile proportions receive pointwise 95% Wilson intervals. These describe uncertainty under respondent-level sampling assumptions; they do not repair undocumented recruitment or provenance.[20]

Cochran's Q tests equality of the three correlated endorsement proportions. Three follow-up exact paired binomial tests use the discordant responses, with Holm adjustment across the three comparisons. Ten thousand respondent bootstrap resamples provide pointwise percentile intervals for paired percentage-point differences, using seed

[20] Edwin B Wilson, 'Probable Inference, the Law of Succession, and Statistical Inference' (1927) 22 Journal of the American Statistical Association 209.

20261020. Resampling preserves the within-record relationship among the items.[21]

3.5 Threshold models and sensitivity analysis

Three separate unpenalised logistic models estimate Pr(B ≥ k) for k = 1, 2 and 3 using the same predictors. These are threshold-specific binary models, without an imposed proportional-odds constraint. They are not used to reconstruct a complete ordinal probability distribution or classify individual respondents. Full coefficient estimates are reported in Appendix C. HC1 sandwich intervals use the respondent score contributions and multiplier n/(n–p), where p includes the intercept. Model convergence is checked through the numerical score.[22]

The three PTG5 coefficients receive Holm adjustment as one family; the PTG1 coefficients form a separate exploratory family. Stacking respondent influence contributions across the models supplies their cross-model covariance. A three-degree-of-freedom Wald test examines whether all PTG5 coefficients are zero, and a two-degree-of-freedom Wald test examines equality across thresholds. These tests address different alternatives; neither is treated as a prerequisite for reporting a threshold estimate. Pointwise intervals are not multiplicity-adjusted simultaneous intervals.

Sensitivity analyses examine complete endorsement in the 133 law-student records, the 382 frequency-defined current users, and the 366 records with neither experience nor frequency coded never. A fourth model expands current status to three indicators. A strict complete profile requires strong agreement with all three provisions; only three records satisfy it, so no multivariable strict-profile model is fitted. Rank associations retain the original item scores as a descriptive check. All analyses are exploratory, with no stepwise selection, imputation, mediation or individual prediction claims.

3.6 Ethics and responsible reporting

This analysis involves no new recruitment or contact. The questionnaire specifies voluntary participation, anonymity, aggregate analysis and the option to skip questions or stop. No formal approval document, exemption determination or verified administration record is supplied. The original collectors must establish the applicable review pathway; absence of a document is not evidence that review was unnecessary. Aggregate reporting and exclusion of the labelled synthetic comments minimise disclosure risk. No attempt is made to identify respondents or infer an individual's political views from a profile. Consent implementation and provenance require confirmation before empirical journal submission.

[21] William G Cochran, 'The Comparison of Percentages in Matched Samples' (1950) 37 Biometrika 256; Quinn McNemar, 'Note on the Sampling Error of the Difference between Correlated Proportions or Percentages' (1947) 12 Psychometrika 153; Sture Holm, 'A Simple Sequentially Rejective Multiple Test Procedure' (1979) 6 Scandinavian Journal of Statistics 65; Bradley Efron and Robert J Tibshirani, An Introduction to the Bootstrap (Chapman & Hall 1993).

[22] Leonard A Stefanski and Dennis D Boos, 'The Calculus of M-Estimation' (2002) 56 American Statistician 29.

4 Results

The results use N = 385 throughout the primary analysis, with smaller denominators stated only for sensitivity checks. They first describe the sample and item responses, then examine the joint accountability profiles and adjusted threshold associations. The findings concern recorded endorsement; they establish neither national opinion nor actual institutional performance.

4.1 Respondent composition and measurement diagnostics

The sample contains a substantial current-law-student group, but students do not constitute the full database. Respondents also include graduates, professionals and the broadly labelled other group. Table 2 shows their distribution alongside age, location and AI experience. Weekly or daily use is recorded for 279 respondents, and 191 report some formal training.

Table 2. Respondent profile in the complete database (N = 385)

Characteristic	Categories in order	Counts in order
Age	18-20 / 21-23 / 24-26	72 / 158 / 74
Age	27-30 / 31+	39 / 42
Gender	Male / female / other / prefer not to say	194 / 175 / 5 / 11
Current status	Law student / law graduate	133 / 59
Current status	Legal professional / other relevant	47 / 146
Education	Undergraduate / graduate / postgraduate / professional	245 / 60 / 33 / 47
Location	Dhaka / other city / district town	162 / 89 / 82
Location	Upazila-rural / other	47 / 5
Use frequency	Never / less than monthly / monthly	3 / 43 / 60
Use frequency	Weekly / daily or almost daily	152 / 127
Training	None / orientation / course or certified	194 / 128 / 63

Counts describe the supplied records and are unweighted. The other relevant respondent category is not further specified. Education level and current legal status are different variables.

The profile defines the scope of the observations rather than a population weighting scheme. In particular, the 245 undergraduate records must not be mistaken for 245 current law students: education level and current status are different fields. The primary denominator is 385, while the student-only sensitivity denominator is 133. Institutional composition and recruitment probabilities cannot be recovered from these categories.

Block alpha is 0.193 for AIL, 0.260 for PTG, 0.389 for ERS and 0.265 for RRB. These values do not support reporting a validated general trust factor or a structural path between composite block scores. The

analysis therefore retains the item roles specified in Table 1 and uses a transparent breadth count only for the three named provisions. The experience-field inconsistencies affect nineteen records in total; they do not create missing outcome data.

4.2 Uneven endorsement across accountability provisions

Table 3 presents the five focal PTG responses. Transparency-contingent trust is endorsed by 284 respondents (73.8%) and accountability-contingent trust by 283 (73.5%). Organisational accountability for harm attracts 294 endorsements (76.4%), compared with 215 (55.8%) for high-risk oversight and 208 (54.0%) for explanation and review. The last two proportions are close despite addressing different arrangements.

Table 3. Distinct trust and accountability responses (N = 385)

Item abbreviated	Mean	SD	Endorsed n (%)	95% interval %
PTG1 Transparent-rule conditional trust	3.97	0.76	284 (73.8%)	69.2-77.9
PTG2 Harm accountability	4.03	0.79	294 (76.4%)	71.9-80.3
PTG3 High-risk oversight	3.63	0.77	215 (55.8%)	50.9-60.7
PTG4 Explanation and review	3.58	0.79	208 (54.0%)	49.0-58.9
PTG5 Accountability conditional trust	3.98	0.76	283 (73.5%)	68.9-77.7

All item scores range from one to five. Intervals are pointwise Wilson intervals for endorsement proportions; they do not establish representativeness. PTG1 and PTG5 state conditions for trust rather than current institutional confidence.

The omnibus matched-response test rejects equal endorsement across the three provisions: Cochran’s Q = 51.45, two degrees of freedom, $p < 0.001$. This finding establishes variation in recorded endorsement across item wording. It does not determine why respondents distinguish the provisions or whether a more detailed description would change their answers. The component distributions should therefore be read alongside the joint profiles and paired contrasts.

4.3 Joint profiles and breadth of support

Table 4 lists every combination in the order harm accountability, high-risk oversight, explanation and review. The largest single profile endorses all three, with 93 records (24.2%; 95% interval 20.1-28.7). However, the remaining profiles together comprise more than three quarters of the sample. The complete profile’s status as the largest individual category should not be confused with majority support for the combination.

Table 4. Eight joint endorsement profiles (N = 385)

Pattern	Provision combination	n	%	95% interval %
000	None endorsed	26	6.8	4.6-9.7
001	Review only	23	6.0	4.0-8.8
010	Oversight only	20	5.2	3.4-7.9
011	Oversight and review	22	5.7	3.8-8.5
100	Harm accountability only	51	13.2	10.2-17.0
101	Harm accountability and review	70	18.2	14.6-22.3
110	Harm accountability and oversight	80	20.8	17.0-25.1
111	All three provisions	93	24.2	20.1-28.7

Pattern order: PTG2 harm accountability, PTG3 oversight, PTG4 explanation and review. One denotes endorsement; zero

includes neutrality and disagreement. Wilson intervals are pointwise, not simultaneous. Profiles exhaust the 385 records.

Exactly two provisions are endorsed by 172 respondents (44.7%), exactly one by 94 (24.4%) and none by 26 (6.8%). The two most common partial profiles combine harm accountability with oversight but not review (80 records), and harm accountability with review but not oversight (70). Only 93 of the 294 harm-accountability endorsers also endorse both other provisions, or 31.6%. The profiles reveal selective combinations that an average score would conceal.

Support for at least one provision is recorded for 359 respondents (93.2%), support for at least two for 265 (68.8%), and support for all three for 93 (24.2%). These thresholds summarise breadth rather than a sequence respondents passed through. Non-endorsement at a threshold includes neutral responses and does not establish hostility to safeguards. Figure 2 displays the profile composition beside the conditional-trust associations.

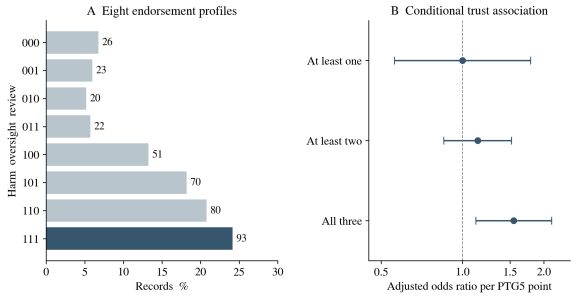

Figure 2. Joint accountability profiles and adjusted conditional-trust associations (N = 385). Panel A gives profile percentages and count labels; one denotes endorsement in the order harm, oversight, review. Panel B shows PTG5 odds ratios and pointwise 95% HC1 intervals across breadth thresholds.

Figure 2 separates the descriptive profiles from model-based associations. The profile panel retains the three specific provisions; the coefficient panel uses the cumulative breadth definitions. The confidence intervals in the latter are pointwise robust intervals, while multiplicity-adjusted p-values appear in Table 6. Neither panel establishes a causal effect of trust or accountability reform.

4.4 Matched differences between provisions

Table 5 reports paired differences, preserving the fact that all three items were answered by the same respondents. Harm accountability exceeds oversight by 20.5 percentage points, with bootstrap interval 14.3-26.5, and exceeds explanation and review by 22.3 points, with interval 15.8-28.6. Both exact paired tests remain below 0.001 after Holm adjustment.

Table 5. Matched provision contrasts with multiplicity adjustment

First minus second	First only / second only	Gap pp	95% bootstrap interval	Holm p
Harm accountability – oversight	121 / 42	20.5	14.3-26.5	< .001
Harm accountability – review	131 / 45	22.3	15.8-28.6	< .001
Oversight – review	100 / 93	1.8	-5.2-8.6	0.666

N = 385 paired records; pp = percentage points. Exact paired tests use discordant cells, with Holm correction across three comparisons. Intervals use 10,000 complete-record bootstrap resamples and are pointwise.

The oversight-review difference is 1.8 percentage points, with interval −5.2-8.6 and adjusted p = 0.666. Similar marginal proportions do not mean the same people endorse each item: 100 endorse oversight without review and 93 endorse review without oversight. This substantial discordance explains why the eight profiles remain informative even when two marginal percentages are close. It also precludes inferring a simple hierarchy in which one procedural preference contains the other.

4.5 Conditional trust across breadth thresholds

Table 6 reports the focal PTG5 and separate PTG1 associations. For accountability-contingent trust, the adjusted odds ratios are 1.00 for at least one provision, 1.14 for at least two and 1.55 for all three. The complete-profile estimate has a pointwise 95% interval of 1.12-2.14, raw p = 0.008 and Holm-adjusted p = 0.025. The corresponding adjusted p-values for the lower thresholds are 0.998 and 0.766.

Table 6. Conditional-trust associations across breadth thresholds

Threshold	Events	PTG5 OR	PTG5 95% CI	Holm p	PTG1 OR	Holm p
At least one	359	1.00	0.56-1.78	0.998	1.75	0.073
At least two	265	1.14	0.85-1.52	0.766	1.31	0.134
All three	93	1.55	1.12-2.14	0.025	1.07	0.695

Each model uses N = 385 and all six non-intercept predictors. PTG5 is accountability-contingent trust; PTG1 is transparent-rule conditional trust. Intervals use HC1 covariance. Holm correction is separate for the three coefficients of each predictor. Joint PTG5 test: p = 0.069; equality across thresholds: p = 0.200.

The complete-profile association supports H1 within the stated exploratory model and remains below 0.05 after adjustment across the three threshold coefficients. It does not establish a uniform association throughout the breadth distribution. The joint three-coefficient Wald statistic is 7.09, three degrees of freedom, p = 0.069. The threshold-equality statistic is 3.22, two degrees of freedom, p = 0.200. Thus, the

numerically larger complete-profile coefficient is not sufficient evidence that the true coefficients differ by threshold.

Transparent-rule conditional trust has odds ratios of 1.75, 1.31 and 1.07 across the three thresholds. Their Holm-adjusted p-values are 0.073, 0.134 and 0.695. No coefficient in that separate family remains below 0.05. The two conditional-trust statements therefore do not yield interchangeable association patterns, although the data do not directly establish a difference between their effects. All fitted-model maximum absolute scores are below 0.000001.

4.6 Sensitivity and interpretation limits

Table 7 reports complete-profile models under alternative specifications. The unadjusted PTG5 odds ratio is 1.57 (95% interval 1.14-2.17). Detailed status adjustment gives 1.51 (1.08-2.09), frequency-defined users give 1.58 (1.14-2.19), and the consistent-user subset gives 1.58 (1.14-2.20). These checks retain the direction of the association, but use overlapping records and are not independent replications.

Table 7. Sensitivity of the complete-profile PTG5 association

Specification or subset	n	OR	95% CI	Raw p
Unadjusted complete profile	385	1.57	1.14-2.17	0.006
Detailed status adjustment	385	1.51	1.08-2.09	0.015
Frequency-defined AI users	382	1.58	1.14-2.19	0.005
Neither field coded never	366	1.58	1.14-2.20	0.007
Current law students only	133	1.25	0.75-2.11	0.392

All rows concern endorsement of all three provisions. Models are adjusted as specified except the first row. The student model omits invariant status; detailed status uses three indicators. Sensitivity p-values are descriptive and not multiplicity-adjusted. Strict complete endorsement has three events and is not modelled.

The student-only estimate is 1.25, with interval 0.75-2.11 and p = 0.392. That wider interval does not establish disagreement between students and other respondents; no status interaction is tested. Strict complete endorsement occurs only three times, preventing a stable multivariable analysis at that threshold. Spearman’s correlation between PTG5 and breadth is 0.111 (p = 0.029), while its rank correlations with the three original provision scores range from 0.059 to 0.074. These modest associations caution against interpreting the complete-profile odds ratio as evidence of a strong general trust disposition.

5 Discussion

The discussion interprets selective endorsement of accountability provisions alongside the conditional-trust association. It connects the profile results with institutional trust and responsible-AI scholarship, while preserving distinctions between stated preferences, operational safeguards and present institutional trustworthiness. The strongest inference concerns the structure of recorded responses, with empirical generalisation dependent on verified provenance and recruitment.

5.1 General responsibility and specified safeguards

The central descriptive result is a difference between support for organisational accountability after significant harm and support for oversight or explanation and review. The difference persists when responses are compared within the same record, so it is not an artefact of using different respondent groups for each percentage. Yet it remains a contrast between questionnaire statements. Different wording, familiarity or interpretations may contribute, and the survey does not identify which explanation accounts for the pattern.

Accountability scholarship helps clarify its significance without supplying an unmeasured cause. Broad responsibility identifies who should answer when harm occurs, while assessment and review specify ways through which decisions can be examined. Institutional accountability requires more than favourable ethical language. The present profiles show why measuring only the broadest statement would overstate agreement with the entire combination of safeguards. This is an empirical measurement contribution, not a demonstration that respondents have inconsistent values.[23]

The result also fits the concern that ethical-principle agreement can conceal differences in institutional interpretation. Jobin and colleagues examine variation across AI-ethics guidelines, and Ananny and Crawford question whether transparency alone delivers accountability. Here, the different objects are public-attitude statements rather than organisational guidelines. The shared implication is that a general principle should not stand in for agreement about the processes needed to realise it.[24]

5.2 What the complete profile adds

Only approximately one quarter of records endorse all three provisions. This is more informative than treating three majority proportions as evidence that most respondents endorse the complete combination. Marginal support and joint support answer different questions. The complete profile identifies simultaneous endorsement; the other seven profiles show where combinations vary. Because the three provisions are not exhaustive, complete means complete within this measured set, not complete commitment to every requirement of accountable AI.

The profile analysis also shows that similar oversight and review percentages conceal different respondents. Those who endorse one without the other may distinguish preventive assessment from an affected person’s opportunity to obtain reasons. They may instead interpret one statement more clearly than the other. The questionnaire cannot decide between these possibilities. A future study should use concrete decisions and ask respondents to explain what they believe assessment or review would involve.

[23] Bovens (n 2); Busuioc (n 1).

[24] Jobin, Ienca and Vayena (n 14); Ananny and Crawford (n 10).

Cobbe, Lee and Singh's reviewability framework provides a useful institutional connection. Reviewing an automated decision requires attention to the surrounding process and relevant evidence, rather than only the internal workings of a model. The profile results therefore motivate more specific public-engagement questions: who should review, what records should be accessible and what can change after review? These are proposed follow-up questions, not mechanisms demonstrated by the present survey.[25]

5.3 Conditional trust and cautious support for H1

The positive complete-profile association indicates that respondents more strongly endorsing accountability as a reason for increased trust are more likely to endorse all three provisions, conditional on the recorded predictors. The estimate remains positive under detailed status adjustment and user-subset exclusions. This pattern is consistent with an institutional-trust interpretation in which safeguards provide reasons for confidence. It does not establish whether greater trust precedes policy support, whether support shapes the trust response or whether both reflect a common evaluative tendency.

The wording overlap is especially relevant. PTG5 mentions effective accountability, and the outcome includes agreement with organisational accountability. The resulting relationship is not a clean test between independent, fully validated constructs. Its value lies in showing how a conditional trust statement corresponds with a specified response combination. The small original-item correlations and low block consistency argue against presenting the result as a strong general pathway from trust to responsible governance.

The broader threshold tests also matter. One threshold association survives Holm adjustment, while the joint test does not cross 0.05 and heterogeneity is not established. These findings are compatible because the tests examine different alternatives. They should prevent a stronger narrative that trust progressively increases every level of support or becomes more important at a particular threshold. A threshold-specific result is informative when presented with the full set of estimates, rather than as evidence for an untested developmental sequence.[26]

5.4 Transparency does not substitute for effective review

The transparent-rule item attracts considerable endorsement, but its adjusted associations do not remain below 0.05 after correction across thresholds. This does not show that transparency lacks importance. The item concerns a condition for trust, not exposure to understandable rules or comprehension of a particular process. It therefore differs from experimental transparency manipulations and from evaluations of actual information access.

[25] Cobbe, Lee and Singh (n 10).

[26] Ronald L Wasserstein and Nicole A Lazar, 'The ASA Statement on p-Values: Context, Process, and Purpose' (2016) 70 American Statistician 129.

Grimmelikhuijsen's work demonstrates the importance of distinguishing forms and contexts of transparency, while Shin focuses on the user's ability to make sense of an explanation. The present uncertain coefficients do not contradict those studies. They address a less specific survey statement and a different outcome: the breadth of safeguard endorsement. The appropriate implication is to improve the precision of measurement rather than to accept or reject transparency as a governance principle.[27]

Effective safeguards also need institutional capacity. Green's critique of human-oversight policies and Alon-Barkat and Busuioc's findings about human-AI interaction caution against assuming that an announced review process will work as intended. The study therefore treats oversight and review as objects of support, not demonstrated solutions. Their implementation would require evidence about reviewer authority, competence, access to information and the practical ability to change a decision.[28]

5.5 Bangladesh and the boundary of public-opinion claims

The Bangladesh contribution is contextual rather than representative. The file combines legally connected respondent categories with a broad other category, and recruitment is undocumented. The results can identify a pattern worth investigating in civic and legal-policy research, but they cannot establish that a specified percentage of Bangladesh's public supports a particular regulatory design. The distinction is important when a sample reaches a conventional target of 385 but lacks the design information needed to support the calculation's assumptions.

The unresolved numerical provenance imposes an additional limit. Synthetic labels for identifiers and comments do not prove that all numerical responses were generated, but they require clarification. Internal checks cannot authenticate a survey. The sixteen inconsistent experience responses and three additional frequency inconsistencies are reported for the same reason: the analytical response is to document and test their influence, not silently reclassify them. The complete-profile association remains positive after those exclusions, while the origins of the records remain a separate question.

The article consequently provides a reproducible way to separate accountability principles, procedural preferences and conditional trust. Its contribution can guide a verified follow-up survey and institutional consultation. It should not be converted into

[27] Grimmelikhuijsen (n 3); Shin (n 11).
[28] Green (n 12); Alon-Barkat and Busuioc (n 12).

a claim that Bangladesh already possesses trustworthy AI governance, that citizens oppose review, or that a measured safeguard has been implemented successfully. Those questions require different evidence.

6 Theoretical and Practical Implications

The implications concern the design of concepts, public engagement and institutional evaluation. They follow from the distinction between broad harm-accountability endorsement and a more selective combination of safeguards. The survey supports investigating these arrangements more precisely; it does not demonstrate that a particular policy intervention will improve trust.

6.1 Theoretical implications

Institutional trust should be measured with an identified trustee and decision context. The present conditional items reveal which governance features respondents associate with greater trust, while leaving current confidence and willingness to rely unmeasured. A stronger future model would distinguish those objects and establish their temporal ordering. This would make it possible to assess whether trustworthy processes precede calibrated confidence, rather than treating an affirmative conditional statement as evidence that the processes already exist.

Accountability support should also retain its component content. The eight-profile approach shows how respondents combine responsibility, oversight and review. A breadth count provides a useful summary only when its components and limitations remain visible. Treating it as a latent psychological trait would require separate validation and a different theoretical claim. This distinction between content-defined indices and reflective measurement prevents an advanced statistical method from obscuring what the questionnaire actually asks.[29]

The complete-profile association identifies a testable proposition: conditional trust may correspond more closely with simultaneous safeguard endorsement than with one broad approval statement. The present threshold-equality test does not establish that difference, so the proposition remains prospective. Future research can compare trust responses after respondents evaluate concrete accountability packages, keeping the possibility of informed scepticism as well as greater confidence in view.

[29] Diamantopoulos and Winklhofer (n 15).

6.2 Practical implications for public engagement

Policy consultation should ask about specified arrangements rather than only whether AI should be accountable. Participants can be shown a consequential public-service decision, the responsible organisation, a proposed assessment process and a route for explanation or review. Questions can then examine which arrangement they understand, support and consider usable. This approach would make different preferences interpretable without labelling respondents as lacking an ethical commitment.

Communication should explain how a person obtains an account and what can follow from it. An accessible notice could identify whether AI contributed to a decision, whom to contact, which information can be reviewed and whether a reviewer can revise the outcome. These are proposed design features, not findings that such processes are absent in Bangladesh. Their evaluation should include comprehension and practical accessibility, rather than public confidence alone.

Oversight should have an evidentiary purpose. Organisations can assess whether reviewers have sufficient information, time and authority to examine a contested recommendation. Internal audit records should be connected to responsibility for addressing findings. Raji and colleagues' lifecycle approach and NIST's risk-management framework offer starting points for such work. Their applicability requires adaptation and evaluation in the relevant institution; neither supplies an effectiveness estimate for the arrangements considered here.[30]

6.3 Policy and societal implications

UNESCO's Recommendation provides a normative basis for connecting accountability with rights, oversight and public understanding. The present study suggests that consultation should make these connections concrete, while respecting differences between public preference and the safeguards a rights-protecting system requires. A majority vote on a short survey item cannot determine the full extent of an affected person's protection.[31]

A prospective Bangladesh study could recruit across clearly defined communities and service contexts, provide understandable descriptions of governance alternatives and collect explanations for preferences. Institutions could evaluate whether people can access and use a review process, not only whether they approve of it in principle. That programme would connect the article's profile distinction with observable institutional performance and a more defensible assessment of public trustworthiness judgements.

[30] Raji and others (n 9); NIST (n 13).
[31] UNESCO, Recommendation (n 13).

7. Conclusion and Future Research

This study examines public trust in AI governance through a distinction between conditional trust and support for three named accountability provisions. All 385 records enter the primary analysis. Organisational accountability for significant harm attracts substantially more endorsement than high-risk oversight or explanation and review. The difference remains in matched-response comparisons. Although each provision attracts majority endorsement separately, only 93 respondents, approximately one quarter of the database, endorse all three together.

The joint profiles constitute the central contribution. They show why a high overall approval percentage should not be treated as support for an entire combination of institutional safeguards. Partial profiles are common, and similar oversight and review proportions conceal substantial disagreement within respondents. These findings describe the scope and composition of recorded support without classifying neutrality as opposition or interpreting the profiles as stages of ethical maturity.

Accountability-contingent trust is positively associated with complete endorsement in the adjusted model, with odds ratio 1.55 and pointwise 95% interval 1.12-2.14. The association remains below 0.05 after Holm adjustment across the three breadth thresholds. However, the joint threshold test and the test of slope equality do not establish a uniform relationship or a difference between thresholds. Wording overlap, modest rank associations and low block consistency further limit a general trust-mechanism interpretation. The findings concern association among self-reported positions, not a policy effect.

Several limitations shape the conclusions. Numerical provenance is unresolved, recruitment and collection dates are undocumented, and institution identifiers are absent. The broad other-status category limits population interpretation. Self-reported conditional statements do not measure present institutional trust or actual experience of accountability. The breadth index covers only three provisions and uses an equal-count convention. Dichotomisation combines neutral and disagreeing responses, while the strict complete profile is too sparse for stable multivariable modelling. Experience-field inconsistencies are disclosed and assessed through exclusions.

Future research should first establish authentic collection records and the applicable ethics pathway, then recruit an explicitly defined population with documented coverage. Measures should separate present trust, perceived trustworthiness, conditional expectations and willingness to rely on a particular decision. Scenario experiments can vary responsibility, assessment and review arrangements independently, while longitudinal research can examine whether experience of an effective process changes calibrated confidence. Qualitative follow-up can explain why respondents support some provisions without others. External institutional evidence should assess whether review is accessible and capable of changing outcomes. Such a design would move beyond recorded approval towards a

substantive evaluation of accountable AI governance.

Declarations

Ethics and consent

No new participant recruitment was undertaken for this numerical analysis. The questionnaire describes voluntary participation, anonymity and the option to stop, and all records contain affirmative consent and adult-eligibility fields. The codebook's provenance labels require clarification before publication as an empirical human-participant study.

Data availability

The analysis uses the numerical Group 7 database, excluding optional contextual comments. Aggregate results and analytical specifications are reported in the article and appendices. Access to row-level records is subject to confirmation of data origin, ownership, confidentiality obligations and applicable participant permissions.

Author contributions

Muhammad Anik Tanvir Faysal: Lead authorship; development of the topic-specific research argument; manuscript preparation. The work was developed under the supervision of Prof. Dr Kazi Abdul Mannan.

Kazi Abdul Mannan: Supervision; conceptual and methodological guidance; interpretation of findings; critical review and editing of the manuscript.

Generative AI disclosure

Generative AI assisted literature organisation, manuscript drafting, language editing and analytical coding. Numerical estimates were computed programmatically from the database. The authors remain responsible for the accuracy of the analysis, interpretation and reporting.

Bibliography

Ajzen I, ‘The Theory of Planned Behavior’ (1991) 50 Organizational Behavior and Human Decision Processes 179

Alon-Barkat S and Busuioc M, ‘Human-AI Interactions in Public Sector Decision Making: “Automation Bias” and “Selective Adherence” to Algorithmic Advice’ (2023) 33 Journal of Public Administration Research and Theory 153

Ananny M and Crawford K, ‘Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability’ (2018) 20 New Media & Society 973

Aoki N, ‘An Experimental Study of Public Trust in AI Chatbots in the Public Sector’ (2020) 37 Government Information Quarterly 101490

Bovens M, ‘Analysing and Assessing Accountability: A Conceptual Framework’ (2007) 13 European Law Journal 447

Busuioc M, ‘Accountable Artificial Intelligence: Holding Algorithms to Account’ (2021) 81 Public Administration Review 825

Cobbe J, Lee MSA and Singh J, ‘Reviewable Automated Decision-Making:

A Framework for Accountable Algorithmic Systems' in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM 2021) 598

Cochran WG, 'The Comparison of Percentages in Matched Samples' (1950) 37 Biometrika 256

——, *Sampling Techniques* (3rd edn, John Wiley & Sons 1977)

Cronbach LJ, 'Coefficient Alpha and the Internal Structure of Tests' (1951) 16 Psychometrika 297

Diamantopoulos A and Winklhofer HM, 'Index Construction with Formative Indicators: An Alternative to Scale Development' (2001) 38 Journal of Marketing Research 269

Efron B and Tibshirani RJ, *An Introduction to the Bootstrap* (Chapman & Hall 1993)

Green B, 'The Flaws of Policies Requiring Human Oversight of Government Algorithms' (2022) 45 Computer Law & Security Review 105681

Grimmelikhuijsen S, 'Explaining Why the Computer Says No: Algorithmic Transparency Affects the Perceived Trustworthiness of Automated Decision-Making' (2023) 83 Public Administration Review 241

Holm S, 'A Simple Sequentially Rejective Multiple Test Procedure' (1979) 6 Scandinavian Journal of Statistics 65

Jacovi A and others, 'Formalizing Trust in Artificial Intelligence: Prerequisites, Causes and Goals of Human Trust in AI' in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM 2021) 624

Jobin A, Ienca M and Vayena E, 'The Global Landscape of AI Ethics Guidelines' (2019) 1 Nature Machine Intelligence 389

Kroll JA and others, 'Accountable Algorithms' (2017) 165 University of Pennsylvania Law Review 633

Levi M and Stoker L, 'Political Trust and Trustworthiness' (2000) 3 Annual Review of Political Science 475

Mayer RC, Davis JH and Schoorman FD, 'An Integrative Model of Organizational Trust' (1995) 20 Academy of Management Review 709

McNemar Q, 'Note on the Sampling Error of the Difference between Correlated Proportions or Percentages' (1947) 12 Psychometrika 153

National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1, NIST 2023)

Raji ID and others, 'Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing' in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (ACM 2020) 33

Rousseau DM and others, 'Not So Different After All: A Cross-Discipline View of Trust' (1998) 23 Academy of Management Review 393

Shin D, 'The Effects of Explainability and Causability on Perception, Trust, and Acceptance: Implications for Explainable AI' (2021) 146 International Journal of

Human-Computer Studies 102551

Sijtsma K, 'On the Use, the Misuse, and the Very Limited Usefulness of Cronbach's Alpha' (2009) 74 Psychometrika 107

Stefanski LA and Boos DD, 'The Calculus of M-Estimation' (2002) 56 American Statistician 29

UNESCO, 'AI for All: Shaping Bangladesh's Future with Ethical AI' (22 May 2025) <https://www.unesco.org/en/articles/ai-all-shaping-bangladeshs-future-ethical-ai> accessed 20 October 2026

——, *Recommendation on the Ethics of Artificial Intelligence* (UNESCO 2022)

von Elm E and others, 'The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) Statement: Guidelines for Reporting Observational Studies' (2007) 4 PLOS Medicine e296

Wasserstein RL and Lazar NA, 'The ASA Statement on p-Values: Context, Process, and Purpose' (2016) 70 American Statistician 129

Wieringa M, 'What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability' in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (ACM 2020) 1

Wilson EB, 'Probable Inference, the Law of Succession, and Statistical Inference' (1927) 22 Journal of the American Statistical Association 209

Appendix A Questionnaire Wording

The focal wording is reproduced in Table A1 to distinguish conditional trust from the three accountability provisions and the adjustment item.

Table A1. Exact wording of focal questionnaire statements

Code	Statement
PTG1	I trust AI governance more when rules are transparent and publicly understandable.
PTG2	Organizations should be accountable for significant harms caused by their AI systems.
PTG3	High-risk AI systems should be subject to appropriate oversight and assessment.
PTG4	People affected by significant AI decisions should have access to explanation and review.
PTG5	Effective accountability increases my trust in AI governance.
AIL3	I understand that AI outputs may reflect bias or incomplete data.

Common agreement scale: one strongly disagree, two disagree, three neutral, four agree, five strongly agree. Demographic and experience categories are shown in Table 2.

The exact statements delimit the findings. None independently establishes current government trustworthiness, legal knowledge or experience of an institutional review process.

Appendix B Analytical Specification

Read the Data sheet of Group-17-Law-SDG-Database.xlsx and retain all 385 records for primary analysis. Define E■, E■ and E■ as one when PTG2, PTG3 and PTG4 respectively are at least four, and zero otherwise. Enumerate the eight profiles in the order E■E■E■, and define breadth B as their sum. The three outcomes are B ≥ 1, B ≥ 2 and B ≥ 3. Each fit uses an intercept, PTG5–3, PTG1–3, AIL3–3, training coded one for any orientation/course, frequency coded zero through four minus two, and current-law-student status coded one versus all others. No weights or imputation are used.

Fit each binary outcome by unpenalised logistic maximum likelihood. With fitted probability p, calculate information A = X′diag[p(1−p)]X. Each respondent's influence contribution is A■¹x■(y■−p■). Sum the influence outer products and multiply by n/(n−k), with k including the intercept, for HC1 covariance. For cross-model inference, stack the PTG5 influence contributions for each respondent before forming covariance. Test all three coefficients against zero with three degrees of freedom. Test the contrasts β■−β■ and β■−β■ jointly against zero with two degrees of freedom. Holm adjustment is applied separately to the three PTG5 tests, the three PTG1 tests and the three paired provision tests.

For paired differences, use 10,000 bootstrap resamples of complete records, seed 20261020. Within each resample, compute each provision's endorsement proportion and subtract the paired proportions before taking the 2.5th and 97.5th percentiles. Exact paired tests condition on the two discordant cells. Wilson intervals use a 1.95996398454 normal quantile and are pointwise. Sensitivity complete-profile models use the same predictor definitions; omit the invariant student indicator within the 133-student subset, and replace it with three status indicators in the detailed-status model. Exclusion of any record coded Never in either experience or frequency gives 366 records. Python calculations use pandas, NumPy and SciPy; numerical score convergence is checked independently of the optimiser's status message.

Appendix C Complete Adjusted Coefficients

Table C1 provides all non-intercept predictor estimates for each breadth threshold. These covariate associations are exploratory; their pointwise intervals should not be read as simultaneous evidence of separate causal effects.

Table C1. Full adjusted threshold model coefficients (N = 385)

Predictor	At least one OR (95% CI)	At least two OR (95% CI)	All three OR (95% CI)
Accountability conditional trust PTG5	1.00 (0.56-1.78)	1.14 (0.85-1.52)	1.55 (1.12-2.14)
Transparency conditional trust PTG1	1.75 (1.08-2.86)	1.31 (0.98-1.74)	1.07 (0.76-1.50)
Bias awareness AIL3	1.24 (0.72-2.15)	1.02 (0.78-1.34)	0.98 (0.71-1.36)
Any formal AI training	0.66 (0.28-1.56)	1.17 (0.75-1.83)	0.91 (0.55-1.50)
AI use frequency step	1.20 (0.78-1.83)	1.39 (1.11-1.74)	1.15 (0.88-1.50)
Current law student	0.48 (0.21-1.08)	0.78 (0.49-1.23)	0.82 (0.49-1.37)

All models include an intercept and the six listed predictors. Intervals use HC1 sandwich covariance and are pointwise. Score variables enter per one-point increment; frequency per ordered category; training contrasts any training with none and status contrasts students with all other respondents. Intercepts are omitted.

The full table shows the recorded adjustment specification without assigning substantive importance solely from whether an interval includes one. The current-status sensitivity in Table 7 replaces the single student indicator with three separate status indicators.