

AI Enabled Cybercrime Awareness and Preventive Behaviour among Law Students in Bangladesh

Dr Kazi Abdul Mannan

drkaziabdulmannan@gmail.com

Center for Academic & Professional Career Development and Research (CAPCDR)

<https://orcid.org/0000-0002-7123-132X>

Aungon Saha

Shanto-Mariam University of Creative Technology <https://orcid.org/0009-0006-3179-8416>

Research Article

Keywords: cybercrime awareness, preventive practices, law students, protection motivation, generative artificial intelligence, Bangladesh, measurement validity

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11390521/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Abstract

AI-assisted impersonation and social engineering make it necessary to distinguish threat recognition from protective practices. This study examines associations between AI-enabled cybercrime awareness and self-reported prevention among law students in Bangladesh. Protection Motivation Theory informs the interpretation of awareness and protective responses. Of 385 survey records, 253 respondents identified as law students and constituted the primary sample. Following reliability diagnostics, awareness and preventive-practice items were analysed separately. Ordinary least squares models used HC3 robust inference, with ordered logistic models, Holm correction and full-sample analyses as sensitivity checks. Among students, 49.0% endorsed recognition of AI misuse for phishing or impersonation, 80.6% reported verifying unusual requests, and 82.2% endorsed security practices. Adjusted awareness coefficients were small and uncertain for both outcomes; all four confidence intervals included zero, and all Holm-adjusted p values were 1.000. Daily AI use was positively associated with verification ($b = 0.308$, 95% CI 0.115–0.502) and security practices ($b = 0.255$, 95% CI 0.067–0.444). Ordered logistic models showed a comparable pattern. The findings support distinguishing recognition, perceived detection capability and behaviour-specific responses in cybersecurity research and university education. They identify contextual associations rather than causal effects and highlight the need for validated measures and prospective behavioural assessments. Generalisation is limited by the available recruitment documentation and the self-reported cross-sectional design.

1 Introduction

Generative artificial intelligence changes the communication environment in which university students assess authenticity. A message can appear coherent, personalised and institutionally credible while requesting an unsafe action. The practical question for university cybersecurity is therefore more specific than whether students have heard about artificial intelligence. It concerns whether recognising possible misuse is associated with checking an unusual request and maintaining everyday security practices. These actions take place at different points in a possible attack. Verification concerns a suspicious interaction; account protection concerns a configuration or repeated routine. Treating them as interchangeable can obscure useful differences.

The NIST Generative Artificial Intelligence Profile provides a risk-management reference for considering the misuse of generative systems and the need for evaluation and governance¹. The present study focuses on perceptions of AI-enabled cybercrime and reported protective practices. It does not classify individual attacks or estimate the proportion of cybercrime involving AI. The outcomes are questionnaire responses rather than incident records or direct observations of account-security settings.

Law students are a relevant population because legal education involves evaluating digital information and communicating about consequential matters. Apparent institutional instructions, fabricated representations of another person and requests involving credentials or payment provide examples of

situations in which verification may be important. This study examines awareness and preventive-practice responses in a Bangladesh law-student sample; it does not estimate the frequency of these incidents or national prevalence.

Awareness interventions often assume a sequence from information to safer conduct. Yet knowing that phishing exists, believing that one can identify a scam, checking unusual requests and using account safeguards are different propositions. A student can endorse one without endorsing the others. Protection Motivation Theory draws attention to the evaluations connecting a threat to a protective response². This makes it an appropriate foundation for asking why awareness alone might be insufficient, but it also creates a measurement obligation. Awareness should not be renamed perceived vulnerability, and reported security practices should not be mistaken for verified protection.

The questionnaire's five-item Cybercrime Awareness and Prevention block combines threat recognition, perceived detection capability, reported practices and a belief that awareness improves behaviour. Combining these statements into one score would blur the distinction between the proposed antecedents and outcomes. The analysis therefore begins with content mapping and reliability assessment and retains separate indicators where the evidence does not support reflective composite measurement.

Information-security scholarship offers models of threat avoidance, policy compliance and protective motivation³. Their relevance here lies in distinguishing beliefs about danger from reasons and opportunities for responding. However, an employee following an organisational security policy differs from a student reporting personal verification practices. The present study does not assume that a mechanism established in one setting transfers unchanged to another. It asks a narrower question using the measures actually available.

The study contributes a population-specific and behaviour-specific assessment of AI-enabled cybercrime awareness. First, it distinguishes the law-student sample from respondents with other current-status categories. Second, it separates two awareness indicators from two practice outcomes, while treating endorsement of the awareness-behaviour proposition descriptively. Third, it examines uncertainty using robust inference, multiplicity adjustment and ordinal sensitivity analysis. Together, these choices clarify the empirical relationship between recognition, perceived capability and specific protective practices.

The objective is to examine associations between AI-enabled cybercrime awareness and reported preventive practices among law students. Four questions organise the analysis. How do recognition, perceived scam-detection capability and reported practices differ descriptively? Are the awareness items associated with verification and security practices after adjustment for general AI knowledge, training exposure and use frequency? Are the conclusions sensitive to an ordinal outcome specification? Does including the wider mixed-status sample alter the pattern of associations? The study is exploratory and was not preregistered.

The study analyses coded numerical survey responses. Its scope is conditional association within the analytical sample, with no claim of nationally representative prevalence or causal identification. The optional contextual-text field was excluded from all analyses. Eligibility, measurement and sampling limitations are addressed in the methodology and limitations sections.

2 Literature Review and Theoretical Foundation

The literature provides reasons to expect a connection between recognising cybercrime and taking protective action, but also reasons to expect that connection to be conditional. The framework developed here focuses on the observable indicators rather than imposing an unmeasured motivational chain. Threat recognition and perceived detection capability are treated as related research topics, not automatically as manifestations of one latent trait. The two preventive outcomes remain separate because an interaction-specific action and an account-security practice may depend on different circumstances.

2.1 AI-enabled cybercrime awareness

For this study, AI-enabled cybercrime awareness refers narrowly to acknowledging the possible misuse of AI for phishing, impersonation or social engineering. CAP1 captures that acknowledgement. It is neither an objective knowledge test nor a measure of personal exposure. CAP2 concerns perceived ability to recognise common warning signs. Its wording places it closer to confidence about detection, although it does not establish actual ability. A respondent could correctly understand that AI-assisted scams exist while feeling uncertain about recognising one in practice.

Threat recognition and perceived detection capability are conceptually distinct. A respondent may understand that AI-assisted scams exist while remaining uncertain about recognising their warning signs. Content assessment and empirical measurement diagnostics are therefore necessary before combining the items. The distinction between measurement specification and statistical estimation provides the basis for selecting an appropriately bounded analytical model⁴.

2.2 Preventive practices

CAP3 asks whether respondents verify unusual requests before sharing money, credentials or sensitive information. CAP4 asks about security measures such as strong passwords and multi-factor authentication where available. Both are self-reports in an agreement response format. They do not record how often verification occurs, which communication channel is used or whether an authentication method is configured correctly. CAP4 also combines examples, so agreement cannot establish that both password quality and multi-factor authentication are present.

Cybersecurity measurement research illustrates the importance of defining behavioural domains explicitly. The Security Behavior Intentions Scale and subsequent smartphone-scale research provide relevant examples of systematic instrument development⁵. The present items are study-specific

indicators rather than validated adaptations of those scales. Their content and empirical relationships are assessed directly. No reverse-coding was applied because the selected statements are positively keyed.

2.3 Protection Motivation Theory

Protection Motivation Theory links protective responses to appraisals of danger and coping ⁶. In this article it provides a reasoning framework: recognising a threat may contribute to protection, but taking action also requires believing that an appropriate response can work and can be performed. Recognition alone therefore need not map closely onto reported conduct. The distinction is especially useful for an educational setting, where a lecture can increase familiarity without necessarily changing a repeated routine.

The available questionnaire does not separately measure severity, vulnerability, response efficacy, response costs, fear and protection intention with validated multi-item scales. CAP2 may overlap with perceived detection capability, but it is not a complete self-efficacy instrument. CAP5 states that greater awareness improves preventive behaviour; it is an endorsement of a proposition rather than an independent estimate of the underlying causal effect. Consequently, this article is PMT-informed and not a full structural test of PMT. A failure to observe a strong awareness coefficient cannot be interpreted as a rejection of the theory.

The distinction between theoretical interpretation and direct measurement also prevents circular reasoning. If CAP5 were included in the awareness score and the combined score used to explain prevention, a respondent's belief that awareness helps could be presented as evidence that it actually does. Keeping CAP5 descriptive protects against that error. It also preserves a useful contrast between endorsing an educational principle and reporting a particular action.

2.4 Threat avoidance and contextual constraints

Technology Threat Avoidance Theory helps frame security behaviour as a response to danger rather than simply adoption of a useful technology ⁷. This is relevant because frequent AI use can coexist with either protective routines or risky exposure. Frequency itself has no necessary protective meaning. It may represent experience, opportunity, confidence, access to tools or a pattern of digital engagement. The present data cannot distinguish those explanations.

Research on security compliance additionally brings routines and organisational influences into view ⁸. These contributions caution against expecting awareness items to explain most behavioural variation. A protective practice may depend on interface defaults, the availability of a second authentication factor, support from peers or the perceived inconvenience of checking a request. None of those variables is directly observed here. Their possible relevance motivates future research rather than functioning as a post hoc explanation that the current models can verify.

2.5 Training and experience

Formal AI training is recorded as no training, short orientation, or course/certified training. The main model combines the latter two categories into an exposure indicator to maintain a compact specification. This variable does not identify cybersecurity-specific content, duration, timing, quality or completion. It cannot be treated as a randomly assigned intervention. A non-significant coefficient would show uncertainty about the adjusted association of the recorded exposure, not evidence that carefully designed training cannot improve security.

Daily AI use is coded as a contrast with all less-frequent categories, including non-use. Its coefficient therefore represents a daily-versus-nondaily difference, conditional on the included covariates. It should not be interpreted as a linear exposure effect or as a comparison between verified experts and novices. The coding provides a compact indicator of use intensity while leaving differences within the nondaily group unresolved.

2.6 General AI knowledge as a covariate

AIL1 addresses plausible but inaccurate AI output, and AIL3 addresses bias or incomplete data. These knowledge-oriented items enter the models separately. AIL2 and AIL4 describe verification and confidentiality practices and are excluded to reduce overlap with the outcomes. AIL5 concerns an oversight norm. The analysis therefore uses two specific knowledge covariates rather than assuming a validated general AI literacy score.

The focal model concerns cybercrime recognition and preventive practices. Digital Evidence Reliability and Deepfake Awareness and Evidentiary Trust items fall outside this behaviour-specific question and are excluded from the explanatory models. This delimitation distinguishes everyday protective responses from judgments about evidentiary reliability while maintaining a common numerical data source.

2.7 Analytical expectations

Four exploratory directional hypotheses are examined. H1a proposes that recognition of AI misuse is positively associated with request verification, and H1b proposes the corresponding association with security practices. H2a proposes that perceived scam-detection capability is positively associated with request verification, and H2b proposes the corresponding association with security practices. These hypotheses describe conditional associations and do not imply causal identification.

There are reasonable counterarguments. A respondent may recognise scams but lack an actionable verification route. Confidence may exceed skill. Strong account safeguards may already be present through service requirements, independently of awareness. These possibilities justify reporting effect sizes and intervals rather than requiring every expectation to obtain support. Inference is strengthened by considering whether the overall pattern survives alternative representations of the ordinal outcomes.

2.8 Conceptual framework

Figure 1 represents four focal associations from two awareness items to two separate practice outcomes. General AI knowledge, training exposure and daily use enter both models as covariates. The arrows define conditional statistical relationships to investigate; they do not assert temporal order. Neither mediation nor moderation forms part of the primary item-level specification. Trust and evidence reliability remain contextual concepts because the selected measures do not justify estimating those mechanisms.

Arrows indicate exploratory associations. The framework is PMT-informed; the full PMT mechanism is not measured.

3 Methodology

The study uses a quantitative cross-sectional design to examine associations between AI-enabled cybercrime awareness and preventive practices. The broader survey database includes 385 respondents across several law-related current-status categories. The primary analysis is restricted to law students, with the full sample used for sensitivity analysis. Individual inclusion probabilities are unavailable; the analysis does not assume that the sample constitutes a probability sample of Bangladeshi university students.

3.1 Data source and analytical population

The numerical survey database contains 385 observations, background and eligibility variables, and 20 five-point agreement items. The wider survey covers AI literacy, cybercrime awareness and prevention, digital evidence reliability, and deepfake-related perceptions. This article focuses on awareness and preventive practices. Current-status responses identify the relevant law-student population within the wider survey.

The primary sample comprises 253 respondents whose current status is recorded as Law student. The remaining 132 respondents are retained only for the full-sample sensitivity analysis. Student status is self-reported. Education-level categories are not used to recode current status, so the analytical population includes the recorded law-student group rather than an assumed undergraduate-only sample.

3.2 Sample size and sampling considerations

The planning target follows Cochran's large-population calculation, $n = z^2 p(1 - p)/e^2$. With $z = 1.96$, $p = 0.50$ and $e = 0.05$, the calculation yields 384.16, rounded upward to 385⁹. This is a sample-planning benchmark. Because selection probabilities are unavailable, it is not used to claim a five-percentage-point national margin of error. Student-specific inference is based on the 253 eligible current-status records.

The analytical database does not include an invitation denominator, recruitment log or sampling frame. A response rate cannot therefore be estimated. Collection dates, institutional coverage and survey mode are not documented in the materials available for this analysis, limiting assessment of selection and

coverage. No completed pilot or instrument-validation procedure is reported because supporting records were unavailable.

3.3 Measurement and coding

Table 1 maps the selected indicators to their analytical roles. Every item uses a five-point agreement scale, from strongly disagree to strongly agree. Higher values indicate greater endorsement. Training exposure equals one for either recorded training category and zero for no training. Daily use equals one for Daily or almost daily and zero otherwise. The two AI knowledge items are entered in their original one-to-five units. The model uses six predictors for each outcome.

Table 1
Measurement roles and coding

Role	Code	Meaning
Awareness	CAP1	Recognition of AI misuse
Capability	CAP2	Perceived scam detection
Outcome 1	CAP3	Request verification
Outcome 2	CAP4	Security practices
Knowledge controls	AIL1; AIL3	Inaccuracy; bias knowledge
Exposure controls	Training; daily use	Binary indicators
Descriptive only	CAP5	Belief that awareness improves prevention

CAP5 is retained for descriptive reporting because it states a belief about the awareness-behaviour relationship rather than measuring either component independently. DER and DET items are outside the focal model. The primary regressions use the individual indicators, consistent with the content assessment and reliability diagnostics described below.

3.4 Data screening

All 20 agreement items were screened for missing values, out-of-range values, exact duplication of complete response vectors and uniform responding across the entire item set. All responses fall within the specified one-to-five range. There are no missing agreement responses, exact duplicate 20-item vectors or records selecting one category across all 20 items. These checks identify selected data-quality issues but do not establish respondent independence or exclude all forms of careless responding.

All 385 records contain affirmative consent, adult-eligibility and AI-awareness fields. The optional contextual-text field is labelled as synthetic in the data dictionary and was excluded from all quantitative and qualitative interpretation. The AI-tool field is not used in the models, so its missing or non-applicable entries do not affect sample inclusion. No numerical response was imputed or altered.

3.5 Measurement assessment and model choice

Cronbach's alpha and the inter-item Pearson correlation are computed for candidate awareness and practice pairs and for the five-item CAP block. For two-item measures, the Spearman-Brown coefficient is also relevant, with $SB = 2r/(1 + r)$ ¹⁰. The student CAP block alpha is 0.079; the awareness pair has $r = -0.078$ and $\alpha = -0.168$; the practice pair has $r = 0.029$ and $\alpha = 0.056$. These results do not support the planned reflective composite interpretation.

The diagnostics favour separate observed-item models rather than reflective composite estimation. Request verification and security practices are therefore analysed as distinct outcomes, and the two awareness indicators enter as separate predictors. Latent-variable loadings, composite reliability, average variance extracted and HTMT are not reported because a reflective latent model is not estimated. This specification preserves item-specific interpretation while acknowledging limited content coverage and measurement error.

3.6 Statistical strategy

For each student outcome, an ordinary least squares model regresses the response on CAP1, CAP2, AIL1, AIL3, training exposure and daily use. HC3 heteroskedasticity-robust standard errors produce two-sided Wald p values and normal-approximation 95% confidence intervals. Slopes are unstandardised differences in agreement points, conditional on the other included predictors. Variance inflation factors assess linear collinearity. The four focal awareness coefficients form one family for Holm correction¹¹. Other coefficients are exploratory context estimates.

Ordered logistic models with identical predictors assess sensitivity to treating the outcomes as ordered categories. Their odds ratios refer to higher response categories under a proportional-odds specification. A formal proportional-odds assessment was not undertaken, so the ordinal estimates are interpreted as sensitivity checks. The same predictor coding is applied to the 385-record full-sample models to assess the influence of sample definition.

3.7 Reproducibility and interpretation

Estimates were calculated programmatically from the numerical dataset, with counts, distributions, coefficients and confidence intervals checked against the variable coding. Since institutional identifiers are unavailable, adjustment for within-institution dependence is not possible. HC3 inference addresses heteroskedasticity but does not account for unobserved clustering. The study estimates exploratory conditional associations; it does not identify causal effects or estimate cybercrime victimisation risk.

3.8 Research ethics and participant protection

Ethical conduct was grounded in respect for participants, minimisation of foreseeable harm and fair treatment. These principles are consistent with the ethical foundation articulated in the Belmont Report¹². The study involved an adult questionnaire about AI-related knowledge, perceptions and routine

protective practices. It involved no clinical procedure, experimental manipulation, collection of account credentials or requirement to disclose identifiable cybercrime incidents. The design addressed participant protection through the questionnaire content, consent information, anonymous recording and aggregate analysis.

Before answering, participants were informed of the academic purpose, voluntary nature and anonymous format of the survey and their option to stop. Consent and age eligibility were recorded before inclusion; all 385 numerical records indicate agreement to participate and age 18 years or above. Participation was based on a voluntary decision rather than a compulsory research requirement. Only the 253 self-identified law students were included in the primary student analysis. The recorded eligibility and consent fields were checked before analysis.

The questionnaire did not request participants' names, passwords or account identifiers. The numerical dataset uses respondent codes and categorical background variables, and results are presented in aggregate. No attempt was made to identify individuals from their responses. Because combinations of background categories may still carry disclosure risks, access to row-level data is subject to confidentiality considerations. Anonymity also limits the feasibility of locating an individual's response after submission; the questionnaire's option to stop applies during participation.

Risk minimisation extended to analytical reporting. Questions focused on general perceptions and practices rather than identifiable allegations, sensitive account contents or details of victimisation. Optional contextual text was excluded from empirical interpretation. Numerical responses were retained in their recorded form. Sampling uncertainty, measurement limitations and non-significant estimates are reported explicitly.

At the time of the study, the authors' institutional setting did not provide an accessible formal review mechanism applicable to this type of non-interventional adult questionnaire research. Formal approval was therefore not obtained, and no approval number, exemption decision or waiver is represented as having been issued. The absence of an accessible mechanism did not remove the researchers' responsibility for informed consent, confidentiality and proportionate risk minimisation. Supervisory responsibility for ethical conduct was distinct from independent ethics approval.

4 Results

Results are presented for the law-student sample, followed by measurement diagnostics, separate outcome models and sensitivity analyses. Agreement percentages combine response categories four and five. All estimates describe the analytical sample and are interpreted within the limits of the cross-sectional self-report design.

4.1 Respondent profile

The 253 students comprise 65.7% of the full sample. One hundred are aged 21–23, 57 are aged 24–26 and 48 are aged 18–20. Thirty report age 31 or above and 18 report age 27–30. Men account for 129 records and women for 115; seven prefer not to disclose gender and two select other. Dhaka is recorded for 101 students, followed by district towns, other city corporations and upazila/rural areas. Table 2 summarises the respondent profile.

Table 2
Student sample profile

Variable	Category	n	%
Age	21–23	100	39.5
Age	24–26	57	22.5
Age	18–20	48	19.0
Age	31+	30	11.9
Age	27–30	18	7.1
Gender	Male	129	51.0
Gender	Female	115	45.5
Gender	Prefer not to say	7	2.8
Gender	Other	2	0.8
Location	Dhaka	101	39.9
Location	District town	54	21.3
Location	Other city corporation	52	20.6
Location	Upazila-rural area	41	16.2
Location	Other	5	2.0
Training	No	145	57.3
Training	Yes, short orientation	77	30.4
Training	Yes, course or certified training	31	12.3
AI use	Weekly	108	42.7
AI use	Daily or almost daily	83	32.8
AI use	Monthly	36	14.2
AI use	Less than monthly	17	6.7
AI use	Never	9	3.6

The sample includes respondents across several location categories, although population representativeness cannot be inferred. Daily or almost daily AI use is reported by 83 students, and 108 report weekly use. Training exposure is recorded for 108 students: 77 report a short orientation and 31 report a course or certified training. These variables capture recorded exposure rather than the quality or cybersecurity content of instruction.

4.2 Item distributions

Table 3 shows a marked difference between recognition-related endorsement and reported practices. Recognition of AI misuse has mean 3.490, while perceived scam-detection capability has mean 3.514. The respective agreement counts are 124 and 132. Request verification has mean 4.123 and agreement count 204; security practices have mean 4.162 and agreement count 208. The latter outcomes are concentrated toward agreement and strong agreement, which limits the dispersion available to explain.

Table 3
Student item statistics and agreement

Item	Mean	SD	Agree n	Agree %
CAP1	3.490	0.805	124	49.0
CAP2	3.514	0.871	132	52.2
CAP3	4.123	0.727	204	80.6
CAP4	4.162	0.752	208	82.2
CAP5	3.581	0.781	135	53.4
AIL1	3.553	0.803	138	54.5
AIL3	4.146	0.744	201	79.4

Practice statements receive greater endorsement than recognition and detection capability. Differences in wording, response thresholds and social desirability may contribute to this pattern; it is not an objective comparison of knowledge and competence. CAP5, the proposition that awareness improves prevention, is endorsed by 135 students, or 53.4%, and remains outside the explanatory models.

4.3 Measurement findings

The five CAP items do not behave as a coherent reflective scale in the student subset. Alpha is 0.079, and the two recognition-related items are slightly negatively correlated. The practice items also have almost no inter-item correlation. The corresponding full-sample alphas remain low: 0.246 for all five CAP items, – 0.025 for the awareness pair and 0.084 for the practice pair. Thus, the student restriction alone does not account for the measurement problem.

Table 4
Reliability diagnostics for candidate composites

Candidate	Student α	Full α	Student pair r
awareness	-0.168	-0.025	-0.078
prevention	0.056	0.084	0.029
CAP	0.079	0.246	Not applicable

These diagnostics support separately interpretable item models. The findings remain specific to the questionnaire statements, and single-item measurement error cannot be separated from substantive variation. The analysis therefore does not make construct-level claims about a validated cybercrime-prevention scale.

4.4 Request verification model

Tables 5a and 5b report HC3 estimates for the separate outcomes. For request verification, recognition has coefficient 0.037, with a 95% interval from -0.078 to 0.153 and $p = 0.529$. Perceived detection capability has coefficient 0.044, with an interval from -0.071 to 0.159 and $p = 0.451$. Neither estimate provides clear evidence of an adjusted association. The intervals remain compatible with modest positive associations as well as values near zero.

Daily AI use has coefficient 0.308, with a 95% interval from 0.115 to 0.502 and $p = 0.002$. This represents the adjusted difference between daily and nondaily users in reported verification. AIL1 and AIL3 have positive estimates with intervals including zero. Training exposure has a negative estimate with substantial uncertainty. Model R^2 is 0.066, indicating limited explanatory power for the included predictors.

4.5 Security practices model

For security practices, the recognition coefficient is 0.067, with a 95% interval from -0.053 to 0.187 and $p = 0.276$. Perceived detection capability has coefficient 0.024, with an interval from -0.078 to 0.126 and $p = 0.646$. These estimates again do not provide clear support for the directional expectations. The two AI knowledge indicators have different coefficient signs, but neither is estimated precisely enough to support a directional conclusion.

Daily AI use is associated with a 0.255-point higher security-practice response, conditional on the model, with a 95% interval from 0.067 to 0.444 and $p = 0.008$. Training exposure has coefficient 0.028 and $p = 0.769$. Model R^2 is 0.045. The models therefore share a pattern of limited explanatory power, uncertain awareness associations and a positive daily-use association. They do not establish whether the same mechanisms underlie the two reported practices.

Table 5
a. Request verification model with HC3 inference

Predictor	b	SE	95% CI	p
CAP1	0.037	0.059	-0.078 to 0.153	0.529
CAP2	0.044	0.058	-0.071 to 0.159	0.451
AIL1	0.090	0.059	-0.026 to 0.207	0.129
AIL3	0.046	0.061	-0.075 to 0.166	0.457
trained	-0.048	0.093	-0.229 to 0.133	0.604
daily	0.308	0.099	0.115 to 0.502	0.002

$n = 253$. $R^2 = 0.066$. Intercept = 3.248. Daily = daily or almost daily use; trained = any recorded training.

Table 5
b. Security practices model with HC3 inference

Predictor	b	SE	95% CI	p
CAP1	0.067	0.061	-0.053 to 0.187	0.276
CAP2	0.024	0.052	-0.078 to 0.126	0.646
AIL1	0.083	0.060	-0.036 to 0.201	0.171
AIL3	-0.061	0.065	-0.188 to 0.066	0.348
trained	0.028	0.096	-0.160 to 0.217	0.769
daily	0.255	0.096	0.067 to 0.444	0.008

$n = 253$. $R^2 = 0.045$. Intercept = 3.708. Daily = daily or almost daily use; trained = any recorded training.

4.6 Multiplicity and collinearity

All four Holm-adjusted p values for the focal awareness coefficients are 1.000. H1a, H1b, H2a and H2b are therefore not supported in this exploratory specification. The confidence intervals indicate uncertainty about the magnitude and direction of the conditional associations; the findings do not establish that the underlying associations are exactly zero.

Student-model variance inflation factors range from 1.032 to 1.060, indicating little linear redundancy among the included predictors. Collinearity is therefore unlikely to account for the uncertainty in the focal estimates. This diagnostic does not address construct coverage, measurement error, selection bias or common-method effects.

4.7 Ordinal and full-sample sensitivity analyses

The ordered logistic models preserve the broad student-sample conclusions. Recognition odds ratios are 1.115 for verification and 1.175 for security practices; perceived detection-capability odds ratios are 1.136 and 1.103. Their p values exceed 0.05. Daily-use odds ratios are 2.313 for verification and 1.796 for security practices, with $p = 0.002$ and $p = 0.024$ respectively. These odds ratios are not probability ratios or changes in actual protection.

Table 6
Student ordered logistic sensitivity results

Outcome	Predictor	Odds ratio	p
CAP3	CAP1	1.115	0.480
CAP3	CAP2	1.136	0.376
CAP3	daily	2.313	0.002
CAP4	CAP1	1.175	0.292
CAP4	CAP2	1.103	0.483
CAP4	daily	1.796	0.024

Table 7
Full-sample awareness sensitivity estimates

Outcome	Predictor	b	95% CI	Raw p
CAP3	CAP1	0.087	0.001 to 0.174	0.049
CAP3	CAP2	0.047	-0.047 to 0.141	0.326
CAP4	CAP1	0.096	-0.001 to 0.192	0.052
CAP4	CAP2	0.056	-0.030 to 0.143	0.203

$n = 385$. Holm-adjusted p values are 0.194 for CAP1 and 0.406 for CAP2 in both outcome models.

In the full 385-record sample, recognition has coefficient 0.087 for verification, with unadjusted $p = 0.049$, and coefficient 0.096 for security practices, with $p = 0.052$. After Holm correction across the four awareness paths, both recognition p values are 0.194. The full-sample models therefore do not provide multiplicity-adjusted support for the focal awareness associations. The sensitivity analysis also shows why the larger mixed-status sample should be distinguished from the student-specific population.

5 Discussion

The principal finding is an uncertain adjusted relationship between the two awareness indicators and the two reported practices among law students. The estimates are bounded by item-specific measurement and the cross-sectional design. They suggest that the recognition and perceived capability captured by

these questions explain little variation in the two practice responses after adjustment. This conclusion concerns the present operationalisation rather than the general value of cybersecurity awareness.

5.1 Recognising a threat is not the same as enacting a response

The distinction between recognition and action is compatible with a PMT-informed account, but the current data do not identify the missing motivational stages. A student may know that impersonation is possible while relying on an existing account-security routine. Another may feel unable to identify sophisticated scams while consistently checking unusual payment requests through another channel. These examples illustrate logically possible pathways; they are not findings about particular respondents.

The contribution lies in examining specific associations across conceptually distinct domains. Separating the recognition and practice items avoids interpreting a combined awareness-and-prevention score as evidence of a relationship between its components. Future confirmatory research can extend this approach by distinguishing threat knowledge, response knowledge, vulnerability, response efficacy, verification opportunity and observed conduct. These mechanisms require dedicated measures rather than proxies assigned after analysis.

5.2 Recognition and reported practice endorsement

Practice statements receive substantially greater endorsement than recognition and detection capability. The pattern suggests that a general awareness-action deficit should not be assumed. Respondents may follow protective routines despite uncertainty about identifying sophisticated scams, while favourable self-reporting or differences in item difficulty may also contribute. The available measures cannot distinguish these explanations.

The practical implication is methodological before it is educational. An intervention should diagnose which specific capability is missing instead of assuming that a general awareness score captures the problem. Students who already report verification may need support in choosing trustworthy channels. Students who express confidence in detection may still benefit from realistic tasks that test calibration. These are candidate programme designs to evaluate, not interventions whose efficacy has been established by this survey.

5.3 Daily use as a contextual association

Daily AI users report higher verification and security-practice endorsement in both student models. This may reflect digital experience, routines, access to services or unobserved differences in motivation. Reverse direction is also possible: more security-conscious students may feel comfortable using AI regularly. Since the measures are contemporaneous and the model has limited covariates, none of these mechanisms can be preferred as a demonstrated explanation.

The frequency association motivates further investigation of AI-use contexts and protective routines. Because use frequency is neither randomly assigned nor an objective competence measure, the result does not establish that increasing AI use would improve cybersecurity. Longitudinal research could examine whether verification routines change as students gain experience with uncertain outputs while separately observing account-security practices.

5.4 Training exposure and interpretation

The training variable has no clear adjusted association with either outcome. This does not establish the ineffectiveness of university training. The recorded categories combine a short AI orientation with a course or certification, and neither confirms that scam detection, verification or authentication was taught. The timing may also differ across respondents. These measurement limitations make a causal training claim particularly inappropriate.

Training evaluation requires clear information about content, delivery, duration and assessment timing. Practical verification instruction may address different competencies from a general introduction to generative AI. Greater understanding of task difficulty may also reduce self-confidence without reducing performance. The broad exposure indicator cannot resolve these possibilities, so its uncertain coefficients should inform better measurement rather than conclusions about training effectiveness.

5.5 Measurement structure and analytical interpretation

The low internal consistency of the CAP block indicates that its items should not be treated as interchangeable indicators of one reflective factor. Their content spans recognition, capability, action and belief about an association. Measurement-specification research shows why such distinctions matter for substantive inference¹³. Separate outcome models provide a more directly interpretable specification for these data.

Item-level analysis nevertheless leaves important limitations. Single items have restricted content coverage and do not permit separate estimation of measurement error. The findings remain tied to the wording of request verification and security practices. Replication should combine multiple behaviour-specific indicators with objective tasks or independently observed practices. The present study supports a bounded item-level interpretation rather than validation of a replacement scale.

5.6 Common methods and alternative explanations

All selected predictors and outcomes come from the same questionnaire. Common-method research identifies reasons that such a design can distort associations¹⁴. The present study has no marker variable, repeated measurement occasion or independent behavioural source that would permit a convincing decomposition of method effects. Low internal consistency and small correlations do not prove that method bias is absent. Desirable practices may also be endorsed more strongly than knowledge statements because respondents recognise the socially preferred answer.

A separate concern is unobserved opportunity. Multi-factor authentication is not available in every interaction, and verification can be easy or difficult depending on the communication channel and organisational support. The item wording acknowledges availability in one case, but the dataset does not measure the relevant opportunities. The unexplained variance should not be attributed automatically to personal irresponsibility. An institution can make a safe response easier or harder regardless of the student's general knowledge.

5.7 Population definition and substantive contribution

The primary models address law students, while the full-sample models include other current-status categories. Differences in the estimates may reflect sample size, respondent composition or both. The sensitivity comparison does not establish which factor accounts for the differences, but it reinforces the importance of specifying the population before interpreting model estimates.

The article focuses on request verification and everyday security practices, distinguishing these outcomes from judgments about digital-evidence reliability or deepfake authenticity. This behaviour-specific scope provides a separate research question within the wider AI and law survey. The contribution concerns population alignment, measurement separation and uncertainty-aware estimation. Any related publications using the same survey should identify their shared data source and distinct analytical scope.

6 Theoretical and Practical Implications

The implications follow from the observed boundaries of the data. They are framed as priorities for measurement, educational design and future evaluation rather than as proven consequences of an intervention. The most useful lesson is to specify the response being promoted and the evidence that would demonstrate improvement.

6.1 Theoretical implications

First, the study distinguishes a PMT-informed question from a complete PMT test. Recognition of possible AI misuse does not operationalise the full threat-appraisal process, while reported practices do not reveal the intervening protection motivation. Keeping those distinctions explicit avoids both overstating support and incorrectly claiming theoretical rejection. The weak awareness coefficients identify the limits of the current operationalisation rather than a final boundary on the theory.

Second, prevention should not automatically be represented by one broad reflective factor. Verifying a request and using authentication safeguards occur under different constraints and may vary independently. The very low practice-pair correlation supports analysing them separately in these records. A future model could test whether they share a higher-order disposition only after developing adequate indicators and obtaining measurement evidence.

Third, the study highlights the importance of the target population in theoretical interpretation. Pooling students with professionals changes the composition of experience, responsibilities and likely opportunities for security action. A larger sample does not necessarily provide a better test of a student-specific proposition. Theory development should therefore identify whose threat appraisal and whose protective response are being explained before selecting an estimation method.

6.2 Educational and institutional implications

Universities could organise instruction around specific decisions. One module could ask students to verify an unusual payment or credential request using an independently obtained contact channel. Another could demonstrate account safeguards and recovery procedures. Assessment should record whether students complete the action correctly, not only whether they agree that the action is desirable. These recommendations are design proposals motivated by the measurement gaps; the current survey has not evaluated their impact.

Confidence should be assessed alongside performance. A task-based exercise can distinguish accurate recognition from overconfidence and excessive suspicion. The NIST Phish Scale illustrates an established approach to considering message difficulty and recipient context in phishing assessment¹⁵. Any local adaptation would require testing and should avoid equating one exercise score with a permanent personal characteristic. Safe handling of simulated messages and clear educational objectives would also be necessary.

Institutions can support behaviour through clear reporting routes, accessible help and realistic verification procedures. Students should know how to check a purported university instruction without relying on the suspicious message itself. Account-security guidance should identify the services and settings to which it applies. A broad warning that AI can be dangerous gives less actionable information than a concrete procedure, but the relative effectiveness of these approaches should be evaluated prospectively rather than assumed.

6.3 Research and governance implications

Future evaluations should use a documented dataset with informed consent and proportionate confidentiality safeguards. Potential outcomes include task accuracy, configuration checks or recorded verification steps, collected without exposing passwords or private communications. The study identifies a need for observable evidence of protective action alongside self-report measures; it does not determine the effectiveness of a particular intervention.

Transparent documentation should distinguish empirical responses from contextual examples and describe recruitment, consent, coding and analysis. Keeping non-empirical contextual material outside the analytical dataset improves interpretability and reproducibility. Future data dictionaries should also separate threat-recognition items, capability judgments and reported practices to support clearer measurement decisions.

7 Conclusion Limitations and Future Research

This study examined AI-enabled cybercrime awareness and preventive-practice endorsement among 253 self-identified law students within a 385-respondent Bangladesh survey. Item-level robust regression and ordered logistic sensitivity models were used because measurement diagnostics did not support the CAP composite. Neither awareness indicator showed a clear adjusted association with either outcome in the student subset. Daily AI use was positively associated with both reported practices, although overall explanatory power was limited.

The awareness indicators provided limited explanatory information about the two practice responses under the selected specification. The findings remain associative and do not establish whether awareness interventions, training or increased AI use change security behaviour. The principal contribution is to distinguish recognition, perceived detection capability and behaviour-specific responses while making uncertainty explicit. This approach identifies priorities for instrument development and subsequent behavioural assessment.

The main limitations concern sampling documentation and measurement. The analysis used an existing numerical survey database without access to original recruitment logs, a sampling frame, collection dates or complete institutional coverage information. Selection probabilities, response rates and independent respondent authentication could not be assessed. Student status is self-reported, and within-institution dependence cannot be examined. The optional contextual-text field was excluded. These limitations restrict population generalisation and assessment of the collection process.

The design is cross-sectional, self-reported and exploratory. Practice endorsement is not observed conduct, and the measures omit central PMT mechanisms. Very low block reliability prevents broad construct-level claims; separate items still contain measurement error. HC3 inference addresses heteroskedasticity, not selection bias or omitted-variable bias. The ordered logistic sensitivity models assume proportional odds without a formal assumption test. Multiple-comparison correction covers the four focal awareness paths but does not transform the wider exploratory analysis into a confirmatory study.

A documented replication should define the sampling frame, eligible population, institutional coverage and fieldwork timeline. It should develop distinct multi-item measures of recognition, detection efficacy, response efficacy, costs and practices, followed by cognitive interviewing and pilot evaluation. Direct recruitment of the target student population would improve alignment between the research question and sampling design.

Longitudinal or experimental work could then compare actionable verification instruction with general awareness information. Outcomes should include ethically designed tasks and observable practices alongside self-reports. Training content, duration and timing should be recorded so that exposure has a clear meaning. A separate external sample would permit prediction assessment without relying on the

same data used to select the model ¹⁶. Such designs would address the mechanisms that the current questionnaire cannot identify.

Declarations

Ethical approval and oversight

Formal ethical approval was not obtained because, at the time of the study, the authors' institutional setting did not provide an accessible formal review mechanism applicable to this non-interventional adult questionnaire research. The study nevertheless followed participant-protection safeguards, including informed consent, voluntary participation, the opportunity to discontinue, an anonymous questionnaire and aggregate reporting. The methodological safeguards are detailed in Section 3.8. The absence of an accessible review mechanism describes the institutional context; responsibility for ethical conduct remained with the researchers.

Informed consent

Informed consent was obtained from adult participants before questionnaire completion. The questionnaire explained the study purpose, voluntary participation, anonymity and the option to discontinue. All analysed records contain affirmative consent and adult-eligibility responses.

Data availability

The numerical data supporting the reported analyses are held by the research team. Requests for de-identified data and analytical specifications may be directed to the corresponding author. Any sharing is subject to participant confidentiality and applicable institutional requirements. The contextual-text field was excluded from the analysis.

Author contributions

Aungon Saha: Lead student researcher; conceptualisation; preparation of study materials; organisation and curation of survey data; manuscript development. The project was conducted under the supervision of Prof. Dr Kazi Abdul Mannan.

Kazi Abdul Mannan: Conceptual and methodological guidance; supervision; interpretation of statistical findings; critical review and editing of the manuscript.

Generative AI disclosure

Generative AI assisted manuscript drafting, language editing, organisation and the development of analytical code. Statistical estimates were computed programmatically from the numerical survey data. Responsibility for the study's analysis, interpretation and reporting rests with the authors.

References

1. National Institute of Standards and Technology Artificial Intelligence Risk Management Framework Generative Artificial Intelligence Profile (NIST AI 600-1, 2024), 10.6028/NIST.AI.600-1
2. Ronald W, Rogers, 'A Protection Motivation Theory of Fear Appeals and Attitude Change' (1975) 91 J Psychol 93–114. 10.1080/00223980.1975.9915803
3. Maddux JE, Ronald W, Rogers 'Protection Motivation and, Self-Efficacy A, Revised Theory of Fear Appeals and Attitude Change' (1983) 19 Journal of Experimental Social Psychology 469–479. 10.1016/0022-1031(83)90023-9
4. Patricia A, Rippetoe, Ronald W, Rogers (1987) Effects of Components of Protection-Motivation Theory on Adaptive and Maladaptive Coping with a Health Threat. 52 J Personality Social Psychol 596–604. 10.1037/0022-3514.52.3.596
5. Liang H, Xue Y, 'Avoidance of Information Technology Threats A Theoretical Perspective' (2009) 33 MIS Q 71–90. 10.2307/20650279
6. Liang H, Xue Y, 'Understanding Security Behaviors in Personal Computer Usage A Threat Avoidance Perspective' (2010) 11 Journal of the Association for Information Systems 394–413. <https://aisel.aisnet.org/jais/vol11/iss7/1/>
7. Vance A, Siponen M, Pahnla S, 'Motivating IS Security Compliance Insights from Habit and Protection Motivation Theory' (2012) 49 Inform Manage 190–198. 10.1016/j.im.2012.04.002
8. Princely, Ifinedo, 'Understanding Information Systems Security Policy Compliance An Integration of the Theory of Planned Behavior and the Protection Motivation Theory' (2012) 31 Computers Secur 83–95. 10.1016/j.cose.2011.10.007
9. Mikko Siponen MA, Mahmood, Pahnla S, 'Employees' Adherence to Information Security Policies An Exploratory Field Study' (2014) 51 Inform Manage 217–224. 10.1016/j.im.2013.08.006
10. Scott B, MacKenzie PM, Podsakoff, Cheryl, Burke, Jarvis (2005) 'The Problem of Measurement Model Misspecification in Behavioral and Organizational Research and Some Recommended Solutions' 90 Journal of Applied Psychology 710–730. 10.1037/0021-9010.90.4.710
11. Serge Egelman and Eyal Peer 'Scaling the Security Wall Developing a Security Behavior Intentions Scale SeBIS' in Proceedings of CHI 2015 (ACM 2015). 10.1145/2702123.2702249
12. Huang H-Y, Demetriou S, Hassan M (2023) Güliz Seray Tuncay, Carl A Gunter and Masooda Bashir, 'Evaluating User Behavior in Smartphone Security: A Psychometric Perspective' in Proceedings of the Nineteenth Symposium on Usable Privacy and SecurityUSENIX Association 509–524. <https://www.usenix.org/conference/soups2023/presentation/huang>
13. William G, Cochran (1977) Sampling Techniques, 3rd edn. Wiley

14. Rob, Eisinga (2013) Manfred te Grotenhuis and Ben Pelzer, 'The Reliability of a Two-Item Scale Pearson Cronbach or Spearman-Brown'. 58 Int J Public Health 637–642. 10.1007/s00038-012-0416-3
15. Sture, Holm, 'A Simple Sequentially Rejective Multiple Test Procedure' (1979) 6 Scandinavian Journal of Statistics 65–70
16. Philip M, Podsakoff SB, MacKenzie J-Y, Lee, Podsakoff NP (2003) Common Method Biases in Behavioral Research A Critical Review of the Literature and Recommended Remedies. J Appl Psychol 88:879–903. 10.1037/0021-9010.88.5.879
17. Philip M, Podsakoff, Scott B, MacKenzie, Podsakoff NP, 'Sources of Method Bias in Social Science Research and Recommendations on How to Control It' (2012) 63 Annual Rev Psychol 539–569. 10.1146/annurev-psych-120710-100452
18. Michelle Steves, Kristen K, Greene, Theofanos M (2020) Categorizing Human Phishing Difficulty A Phish Scale. 6 J Cybersecur tyaa009. 10.1093/cybsec/tyaa009
19. Shanee Dawkins and Jody Jacobs, NIST Phish Scale User Guide (2023) NIST Tech Note 2276. 10.6028/NIST.TN.2276
20. Gavin C, Cawley, Nicola LC, Talbot 'On Over-fitting in Model Selection and Subsequent Selection Bias in Performance Evaluation' (2010) 11 Journal of Machine Learning Research 2079–2107, <https://www.jmlr.org/papers/v11/cawley10a.html>
21. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research (1979) <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html>


Footnotes

1. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework Generative Artificial Intelligence Profile (NIST AI 600-1, 2024), doi:10.6028/NIST.AI.600-1.
2. Ronald W Rogers, 'A Protection Motivation Theory of Fear Appeals and Attitude Change' (1975) 91 Journal of Psychology 93–114, doi:10.1080/00223980.1975.9915803.; James E Maddux and Ronald W Rogers, 'Protection Motivation and Self-Efficacy A Revised Theory of Fear Appeals and Attitude Change' (1983) 19 Journal of Experimental Social Psychology 469–479, doi:10.1016/0022-1031(83)90023-9.; Patricia A Rippetoe and Ronald W Rogers, 'Effects of Components of Protection-Motivation Theory on Adaptive and Maladaptive Coping with a Health Threat' (1987) 52 Journal of Personality and Social Psychology 596–604, doi:10.1037/0022-3514.52.3.596.
3. Huigang Liang and Yajiong Xue, 'Avoidance of Information Technology Threats A Theoretical Perspective' (2009) 33 MIS Quarterly 71–90, doi:10.2307/20650279.; Huigang Liang and Yajiong Xue, 'Understanding Security Behaviors in Personal Computer Usage A Threat Avoidance Perspective' (2010) 11 Journal of the Association for Information Systems 394–413, <https://aisel.aisnet.org/jais/vol11/iss7/1/>.; Anthony Vance, Mikko Siponen and Seppo Pahnla,

- 'Motivating IS Security Compliance Insights from Habit and Protection Motivation Theory' (2012) 49 Information and Management 190–198, doi:10.1016/j.im.2012.04.002.; Princely Ifinedo, 'Understanding Information Systems Security Policy Compliance An Integration of the Theory of Planned Behavior and the Protection Motivation Theory' (2012) 31 Computers and Security 83–95, doi:10.1016/j.cose.2011.10.007.; Mikko Siponen, M Adam Mahmood and Seppo Pahnla, 'Employees' Adherence to Information Security Policies An Exploratory Field Study' (2014) 51 Information and Management 217–224, doi:10.1016/j.im.2013.08.006.
4. Scott B MacKenzie, Philip M Podsakoff and Cheryl Burke Jarvis, 'The Problem of Measurement Model Misspecification in Behavioral and Organizational Research and Some Recommended Solutions' (2005) 90 Journal of Applied Psychology 710–730, doi:10.1037/0021-9010.90.4.710.
 5. Serge Egelman and Eyal Peer, 'Scaling the Security Wall Developing a Security Behavior Intentions Scale SeBIS' in Proceedings of CHI 2015 (ACM 2015), doi:10.1145/2702123.2702249.; Hsiao-Ying Huang and others, 'Evaluating User Behavior in Smartphone Security A Psychometric Perspective' (2023) Symposium on Usable Privacy and Security, <https://www.usenix.org/conference/soups2023/presentation/huang>.
 6. Ronald W Rogers, 'A Protection Motivation Theory of Fear Appeals and Attitude Change' (1975) 91 Journal of Psychology 93–114, doi:10.1080/00223980.1975.9915803.; James E Maddux and Ronald W Rogers, 'Protection Motivation and Self-Efficacy A Revised Theory of Fear Appeals and Attitude Change' (1983) 19 Journal of Experimental Social Psychology 469–479, doi:10.1016/0022-1031(83)90023-9.; Patricia A Rippetoe and Ronald W Rogers, 'Effects of Components of Protection-Motivation Theory on Adaptive and Maladaptive Coping with a Health Threat' (1987) 52 Journal of Personality and Social Psychology 596–604, doi:10.1037/0022-3514.52.3.596.
 7. Huigang Liang and Yajiong Xue, 'Avoidance of Information Technology Threats A Theoretical Perspective' (2009) 33 MIS Quarterly 71–90, doi:10.2307/20650279.; Huigang Liang and Yajiong Xue, 'Understanding Security Behaviors in Personal Computer Usage A Threat Avoidance Perspective' (2010) 11 Journal of the Association for Information Systems 394–413, <https://aisel.aisnet.org/jais/vol11/iss7/1/>.
 8. Anthony Vance, Mikko Siponen and Seppo Pahnla, 'Motivating IS Security Compliance Insights from Habit and Protection Motivation Theory' (2012) 49 Information and Management 190–198, doi:10.1016/j.im.2012.04.002.; Princely Ifinedo, 'Understanding Information Systems Security Policy Compliance An Integration of the Theory of Planned Behavior and the Protection Motivation Theory' (2012) 31 Computers and Security 83–95, doi:10.1016/j.cose.2011.10.007.; Mikko Siponen, M Adam Mahmood and Seppo Pahnla, 'Employees' Adherence to Information Security Policies An Exploratory Field Study' (2014) 51 Information and Management 217–224, doi:10.1016/j.im.2013.08.006.
 9. William G Cochran, Sampling Techniques (3rd edn, Wiley 1977).
 10. Rob Eisinga, Manfred te Grotenhuis and Ben Pelzer, 'The Reliability of a Two-Item Scale Pearson Cronbach or Spearman-Brown' (2013) 58 International Journal of Public Health 637–642, doi:10.1007/s00038-012-0416-3.

11. Sture Holm, 'A Simple Sequentially Rejective Multiple Test Procedure' (1979) 6 Scandinavian Journal of Statistics 65–70.
12. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, The Belmont Report (1979), <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html>. Used here as an ethical-principles reference, not as a source of an exemption from review.
13. Scott B MacKenzie, Philip M Podsakoff and Cheryl Burke Jarvis, 'The Problem of Measurement Model Misspecification in Behavioral and Organizational Research and Some Recommended Solutions' (2005) 90 Journal of Applied Psychology 710–730, doi:10.1037/0021-9010.90.4.710.
14. Philip M Podsakoff, Scott B MacKenzie, Jeong-Yeon Lee and Nathan P Podsakoff, 'Common Method Biases in Behavioral Research A Critical Review of the Literature and Recommended Remedies' (2003) 88 Journal of Applied Psychology 879–903, doi:10.1037/0021-9010.88.5.879.; Philip M Podsakoff, Scott B MacKenzie and Nathan P Podsakoff, 'Sources of Method Bias in Social Science Research and Recommendations on How to Control It' (2012) 63 Annual Review of Psychology 539–569, doi:10.1146/annurev-psych-120710-100452.
15. Michelle Steves, Kristen K Greene and Mary Theofanos, 'Categorizing Human Phishing Difficulty A Phish Scale' (2020) 6 Journal of Cybersecurity tyaa009, doi:10.1093/cybsec/tyaa009.; Shanee Dawkins and Jody Jacobs, NIST Phish Scale User Guide (NIST Technical Note 2276, 2023), doi:10.6028/NIST.TN.2276.
16. Gavin C Cawley and Nicola LC Talbot, 'On Over-fitting in Model Selection and Subsequent Selection Bias in Performance Evaluation' (2010) 11 Journal of Machine Learning Research 2079–2107, <https://www.jmlr.org/papers/v11/cawley10a.html>.

Figures



Covariates in both outcome models
 AI inaccuracy knowledge • AI bias knowledge • Training exposure • Daily use

Figure 1

Proposed item-level analytical framework

Arrows indicate exploratory associations. The framework is PMT-informed; the full PMT mechanism is not measured.