

DOI: https://doi.org/10.48009/2_iis_2026_146

The role of personalized training in influencing information security attitudes and behavioral intentions

Jacob Donaldson, *Industry Scholar*, jddonaldson04@gmail.com

Gabriel Giordano, *Ohio University*, giordano@ohio.edu

Abstract

Despite decades of investment in technical defenses and mandatory awareness programs, the most consequential information security (IS) failures continue to originate in human behavior. This study examines whether personalizing IS awareness and training content to a learner's functional domain produces measurable improvements in two psychological precursors of secure behavior: attitude toward IS compliance and behavioral intention to comply. The investigation is anchored in three complementary theories of behavior change - Protection Motivation Theory (PMT), the Theory of Planned Behavior (TPB), and the Elaboration Likelihood Model (ELM) - with PMT, motivational affordance, and role identity informing the design of the personalization manipulation and ELM and TPB grounding the tested hypotheses. Using a randomized between-subjects experiment with 317 undergraduate marketing students, we compared a personalized training condition that embedded role-specific examples, quotations from marketing professionals, and domain-relevant Excel scenarios against a generic training condition that delivered the same core content in role-neutral framing. A multivariate analysis of variance (MANOVA) on pre-to-post change scores yielded a significant overall effect (Wilks' $\lambda = 0.981$, $F(2, 314) = 3.12$, $p = .045$, partial $\eta^2 = .019$). Follow-up univariate tests revealed a significant improvement in attitude for the personalized group ($F(1, 315) = 6.05$, $p = .014$, partial $\eta^2 = .019$) but no significant difference in behavioral intention ($F(1, 315) = 1.37$, $p = .243$, partial $\eta^2 = .004$). Findings suggest that role-aligned personalization is a feasible, low-cost lever for shifting employee attitudes toward secure behavior, but that intention, a more proximal predictor of action, likely requires sustained exposure, normative reinforcement, or behavioral practice opportunities to shift. Because the proposed mediating mechanisms were not directly measured, the study should be read as a test of personalization's net effect on attitude and intention rather than of the specific cognitive pathways that produce it. Implications for IS program design, theoretical refinement, and future longitudinal and field research are discussed.

Keywords: information security, security awareness training, personalization, security attitude, behavioral intention, elaboration likelihood model

Introduction

In March 2021, a highly publicized breach at the global security-camera company Verkada laid bare a familiar vulnerability - one rooted not in cryptography or network design, but in everyday human behavior. Attackers gained access to more than 150,000 live camera feeds across hospitals, schools, police stations, prisons, and corporate sites including Tesla and Cloudflare. They did so not through novel malware or a zero-day exploit, but by reusing administrator credentials that had been left exposed online (Greenberg, 2021; Verkada, 2021). Verkada was not a security-naïve organization. The company maintained a written

IS policy, deployed mandatory awareness training, and operated within a regulatory environment that incentivized compliance. Yet the training did not produce the secure credential-handling behaviors it was designed to instill.

The Verkada incident is emblematic rather than exceptional. Similar patterns characterize the breaches at Equifax, Target, the Colonial Pipeline, MGM Resorts, and SolarWinds, where the ultimate point of failure was human - an unpatched system, a phished credential, a social-engineered help-desk agent, a weak or reused password (Sen & Borle, 2015; Crossler et al., 2013). The consistent moral of these incidents is that technical defenses, while imperfect, are typically adequate to the threats they were designed to address; the dominant residual risk lives in the behavior of authorized users (Bulgurcu, Cavusoglu, & Benbasat, 2010). Verizon's annual Data Breach Investigations Report has documented for over a decade that a large majority of breaches involve a human element, whether through error, misuse, or social engineering.

If most failures are behavioral, then most defenses must also be behavioral. Yet conventional security awareness and training programs remain stubbornly generic. They are typically delivered as compliance artifacts - annual, video-based, identical across roles - and they emphasize policy recall over psychological engagement (Karjalainen & Siponen, 2011; Puhakainen & Siponen, 2010). Employees are aware that policies exist; they often cannot, or will not, translate that awareness into action. This persistent gap between knowledge and behavior is what makes attitude and intention important constructs of study. Attitude - an individual's overall evaluation of secure behavior as beneficial, relevant, and worth enacting - frames whether security is internalized as a value or merely tolerated as a constraint. Intention, the conscious self-instruction to perform a behavior, stands between attitude and action and is widely regarded as the most proximal psychological predictor of behavior (Ajzen, 1991; Sheeran & Webb, 2016).

The present study examines whether a comparatively modest intervention, personalizing the framing of IS training to a learner's functional domain, can shift these two psychological precursors. We focus on personalization for three reasons. First, personalization is increasingly feasible at scale, given content-management systems and learning platforms that can dynamically assemble role-relevant modules. Second, personalization is theoretically well-motivated: PMT, TPB, and ELM all predict that role-aligned content should activate stronger threat appraisals, stronger normative pressures, and deeper cognitive engagement than generic content. Third, despite this theoretical plausibility, empirical evidence on the effects of role-based personalization in IS contexts is comparatively thin. The contribution of this paper is therefore both theoretical and practical: we test a clearly specified personalization manipulation against a generic control and report differential effects on attitude and intention.

Background and Literature Review

Three theoretical traditions - Protection Motivation Theory, the Theory of Planned Behavior, and the Elaboration Likelihood Model dominate the behavioral IS literature and motivate the present study. We summarize each and then integrate them through the complementary lenses of motivational affordance and role identity. As we develop below, these traditions play two distinct roles in this paper. ELM and TPB supply the causal logic for the two formal hypotheses tested here. PMT, motivational affordance, and role identity instead serve as the design rationale for the personalization manipulation: they explain why role-aligned examples, role-relevant consequences, and role-credible quotations were expected to function as effective personalization cues, even though the appraisal constructs they name were not themselves measured as mediators in this study. Distinguishing the theories that justify the manipulation from the theories that generate the testable predictions is intended to keep the scope of the empirical claims clear.

Protection Motivation Theory (PMT)

Protection Motivation Theory was originally developed to explain how fear appeals influence health behavior (Rogers, 1975) and was subsequently refined to emphasize the cognitive appraisals that mediate

behavioral response (Maddux & Rogers, 1983). In its modern form, PMT models protective behavior as a function of two parallel appraisal processes. Threat appraisal weighs perceived severity (how bad the outcome would be) and perceived vulnerability (how likely the outcome is to befall me), and is offset by intrinsic and extrinsic rewards associated with maladaptive behavior. Coping appraisal weighs response efficacy (whether the recommended behavior would actually prevent the threat) and self-efficacy (whether I am able to perform the behavior), and is offset by the response costs of compliance.

PMT has been widely adopted in IS research because it accommodates the practical realities of cybersecurity decision making: users are routinely asked to enact behaviors whose costs are immediate and visible (longer passwords, multi-factor prompts, slower workflows) and whose benefits are probabilistic and counterfactual (a breach that did not occur). Boss, Galletta, Lowry, Moody, and Polak (2015) provide a comprehensive synthesis of PMT in IS, arguing that fear appeals are most effective when paired with credible efficacy information. Herath and Rao (2009) extended PMT to organizational settings, demonstrating that perceived effectiveness and self-efficacy explain compliance even after controlling for penalties and peer pressure. In the present study, PMT is used to inform the construction of the personalized condition rather than as a source of measured mediators: role-specific examples were written to raise the salience of perceived vulnerability (“this is the kind of thing that could happen to me, in my job”) and role-relevant remediation steps were written to raise the salience of response efficacy. We did not measure these appraisals directly, and we therefore treat PMT as design rationale, returning to this scope boundary when the hypotheses are developed.

Theory of Planned Behavior (TPB)

The Theory of Planned Behavior (Ajzen, 1991) models behavioral intention as a weighted function of three antecedents: attitude toward the behavior, subjective norms (perceived social pressure), and perceived behavioral control. Attitude reflects the overall favorable or unfavorable evaluation of the behavior in question and is itself shaped by beliefs about outcomes and the value placed on those outcomes. Subjective norms reflect beliefs about what salient referents (peers, supervisors, professionals one respects) think one ought to do. Perceived behavioral control reflects beliefs about the ease or difficulty of performing the behavior and overlaps conceptually with self-efficacy.

TPB has been productively applied to IS compliance. Bulgurcu et al. (2010) showed that attitude toward compliance is the strongest predictor of intention to comply with information security policies, with the attitude itself shaped by beliefs about cost of compliance, cost of non-compliance, and the intrinsic benefits of secure behavior. Ifinedo (2012) integrated TPB with PMT and found that personal relevance and self-interest predicted compliance better than awareness alone. Pahnla, Siponen, and Mahmood (2007) similarly demonstrated that normative pressure and habit shape compliance intention. TPB supplies the causal logic for the second hypothesis tested in this study: it locates intention downstream of attitude and subjective norm, which implies that an intervention that improves attitude should, other things equal, propagate into intention. The embedded quotations from marketing professionals were included as a subjective-norm cue consistent with this logic, importing role-relevant referents into the learner’s salient referent set.

Elaboration Likelihood Model (ELM)

The Elaboration Likelihood Model (Petty & Cacioppo, 1986) addresses a question that PMT and TPB do not: under what conditions are persuasive messages processed deeply enough to produce lasting attitudinal change? ELM identifies two routes to persuasion. Central-route processing engages systematic, effortful evaluation of message arguments and produces attitude change that is more durable, more resistant to counterpersuasion, and more predictive of behavior. Peripheral-route processing relies on superficial cues - source attractiveness, message length, surface markers of expertise - and yields attitude change that is more fragile and less behaviorally consequential. Two factors govern which route dominates: motivation and ability. Motivation rises with personal relevance; ability rises with comprehension.

ELM has been imported into IS research with substantial success. Bhattacharjee and Sanford (2006) showed that job relevance and information quality drive central-route processing of IT adoption messages, while source credibility primarily engages peripheral processing. Johnston and Warkentin (2010) demonstrated that fear-appeal communications produce significant changes in IS-related attitude and intention, and that source credibility and self-efficacy moderate these effects. Posey, Roberts, and Lowry (2013) found that organizational commitment and felt obligation to the organization, both highly identity-laden constructs, shape compliance behavior over and above cognitive appraisals. ELM supplies the causal logic for the first hypothesis tested here: role-aligned content increases perceived personal relevance, which raises motivation to process the message via the central route, which in turn should produce more durable change in attitude.

Motivational Affordance and Role Identity

The PMT–TPB–ELM trio captures the cognitive architecture of behavior change but leaves underspecified the design properties of training content that activate that architecture. Two adjacent constructs, motivational affordance and role identity, help bridge that gap and, like PMT, are used here as design rationale rather than as measured mechanisms. Motivational affordance, introduced by Zhang (2008), refers to the properties of an artifact (here, a training module) that enable it to satisfy a learner’s underlying motivations, such as competence, autonomy, and relatedness. Personalized content is, in essence, a motivational affordance: it offers the learner an opportunity to see security behavior as connected to their professional identity rather than imposed from outside it.

Role identity theory (Stryker, 1980) and the closely related concept of professional identification (Ashforth, Harrison, & Corley, 2008) argue that individuals organize behavior around socially situated identities and that identity-congruent behaviors are pursued more spontaneously and more durably than identity-incongruent ones. Renaud and Goucher (2014) have argued specifically that IS awareness initiatives are more effective when they recruit identity, not only cognition. Personalization in this study was operationalized to activate the marketing-professional identity through both content (analytics dashboards in Excel, customer-data scenarios) and source (quotations from marketing instructors and practitioners). The theoretical expectation is that role-congruent framing makes the security message feel like “what people like me do,” rather than “what the IT department wants me to do.” Because role identity and motivational affordance describe properties of the manipulation rather than constructs we measured, they are not represented in the hypotheses or the analysis; they explain how the personalized condition was built.

Integration

These traditions converge on a single design-level prediction. PMT-informed examples were written to raise the salience of vulnerability and response-efficacy concerns; TPB-informed referents were included to raise subjective-norm pressure; and ELM specifies that role relevance raises motivation to process centrally, yielding more durable attitudinal change. Motivational affordance and identity-based frameworks specify that role-aligned framing is the design property that is expected to activate these mechanisms. Two of these traditions, ELM and TPB, yield the directional, testable predictions formalized in the next section. The remaining traditions justify why the manipulation was expected to work but are not themselves tested, a boundary we make explicit in the hypotheses.

Hypothesis Development

Two hypotheses follow from the theoretical integration above, and each is grounded in a single, clearly identified mechanism. Before stating them, we make the scope of the empirical test explicit. The constructs that PMT, motivational affordance, and role identity introduce, perceived vulnerability, response efficacy, depth of processing, and identity salience, are treated in this study as theoretically inferred consequences of the manipulation, not as measured mediators. The design varies role alignment and observes its net effect

on attitude and intention; it does not instrument the intervening appraisals. The hypotheses below should therefore be read as predictions about the observable endpoints (attitude change and intention change), with the mechanism named only to clarify the expected direction of effect. We return to the implications of not measuring these mediators in the Limitations section, where we treat their absence as a substantive constraint on causal interpretation rather than a minor caveat.

The first hypothesis concerns attitude change, the construct most directly implicated by ELM's central-route mechanism. Because role-aligned content raises perceived personal relevance, learners in the personalized condition should engage more deeply with the message, evaluate its arguments more carefully, and emerge with attitudes that are more favorable. Both Bhattacharjee and Sanford (2006) and Johnston and Warkentin (2010) provide empirical precedent for this prediction in adjacent contexts.

H1: *Personalized information security training that reflects participants' functional roles will produce a greater pre-to-post increase in positive attitude toward information security than will generic training.*

The second hypothesis concerns behavioral intention. TPB locates intention downstream of attitude, subjective norm, and perceived behavioral control. Because the personalization manipulation was designed to target attitude (via central-route processing) and subjective norm (via role-relevant referents), it should, on TPB logic, propagate into intention. Workman, Bommer, and Straub (2008) showed that omission of security behavior is shaped by both knowledge and motivation, with motivation more closely tied to intention. Johnston and Warkentin (2010) showed that fear appeals incorporating personal relevance significantly raise IS compliance intention. We therefore expect a parallel effect on intention, while noting that intention is conceptually further downstream than attitude and may require longer or more situated interventions to shift reliably (Sheeran & Webb, 2016).

H2: *Participants receiving personalized information security training will exhibit a greater pre-to-post increase in behavioral intention to comply with information security practices than will participants receiving generic training.*

Methodology

Participants and Design

Of 338 students who began the study, 12 responses were missing more than 20% of their values and were removed as incomplete, and an additional 9 responses gave the same answer to every item and were removed as unengaged, leaving a final analytic sample of 317 undergraduate marketing students recruited from a large Midwestern public university. The retained sample was 53% male ($n = 169$) and 47% female ($n = 148$), with a mean age of approximately 21 years. All participants accessed the study through the university's online research participation system and provided informed consent prior to enrollment. The study used a between-subjects randomized design with two conditions - personalized training and generic training - delivered through the Qualtrics survey platform. Random assignment to condition was managed by Qualtrics, and pre- and post-intervention measures were collected within the same session to minimize attrition.

Manipulation

Both conditions delivered the same core IS content drawn from Microsoft's published Excel protection and security training materials, covering topics such as file-level password protection, workbook and worksheet protection, the principle of least access in shared spreadsheets, and recognition of common credential-related risks. The conditions differed only in framing. In the personalized condition, training materials were

embedded with marketing-relevant examples (customer analytics dashboards, campaign-tracking spreadsheets, audience-segmentation worksheets), screenshots that reflected realistic marketing artifacts, and brief quotations from marketing professors and practitioners articulating why the security practice in question mattered for marketing work. In the generic condition, the same security practices were presented with neutral, organizationally generic examples and no role-specific quotations. Module length, number of screenshots, and total reading time were held constant across conditions to isolate the personalization treatment from confounding differences in dosage.

Measures

Attitude toward IS (ATTI) and behavioral intention to comply with IS protocols (BINT) were measured pre- and post-intervention using multi-item seven-point Likert scales adapted from Johnston and Warkentin (2010) and Bulgurcu et al. (2010). Sample attitude items included “Following recommended information security practices is wise,” and “Following information security guidelines is beneficial.” Sample intention items included “I intend to follow recommended information security practices in my work,” and “I plan to comply with the information security guidelines presented.” Control variables included prior Excel experience, gender, age, and academic standing.

Because the model includes both reflective and formative measures, measurement quality was evaluated separately for each. Behavioral intention was modeled reflectively; following Gefen and Straub (2005), convergent validity was assessed through standardized item loadings and average variance extracted (AVE), and reliability through composite reliability. As reported in Table 1, all behavioral-intention item loadings exceeded the conventional .70 threshold, AVE exceeded .50 (AVE = .92), and composite reliability exceeded .70 (CR = .972), with the square root of the AVE (.959) exceeding inter-construct correlations, supporting convergent and discriminant validity. Attitude was modeled formatively; following Loch, Straub, and Kamel (2003) and Johnston and Warkentin (2010), a modified multitrait–multimethod (MTMM) matrix of item-to-item and item-to-construct correlations was inspected (Table 2). All attitude items correlated significantly with their composite construct value ($p < .01$), supporting convergent validity for the formative measure.

Table 1. Reliability and Validity for the Reflective Construct (Behavioral Intention)

Item / Statistic	Value
BINT1 loading	0.953
BINT2 loading	0.957
BINT3 loading	0.968
Average Variance Extracted (AVE)	0.92
Composite Reliability (CR)	0.972
Square root of AVE	0.959

Table 2. Inter-Item and Item-to-Construct Correlation Matrix for Attitude (ATTI)

	ATTI1	ATTI2	ATTI3	ATTI4	ATTI
ATTI1	–				
ATTI2	0.848	–			
ATTI3	0.672	0.747	–		
ATTI4	0.771	0.884	0.809	–	
ATTI	0.893	0.947	0.883	0.946	–

Note. Bold values in the bottom row are item-to-construct correlations; all are significant at $p < .01$.

Analysis Strategy

Hypotheses were tested using a multivariate analysis of variance (MANOVA) on pre-to-post change scores for the two dependent variables, conducted in SPSS. MANOVA was selected because attitude and intention are conceptually correlated and a multivariate test controls for inflated familywise error. We report Wilks' Lambda for the multivariate effect and follow up with univariate F-tests for each dependent variable. Effect sizes are reported as partial eta squared. A manipulation check confirmed that participants in the personalized condition rated the training as more job-relevant than did participants in the generic condition; the manipulation-check items were collected on the same seven-point format as the outcome measures.

Results

The MANOVA revealed a statistically significant overall effect of training type on the joint outcome of attitude and intention (Wilks' $\lambda = 0.981$, $F(2, 314) = 3.12$, $p = .045$, partial $\eta^2 = .019$). The full multivariate test results are reported in Table 3, and the univariate between-subjects effects in Table 4. Univariate follow-up tests indicated a significant effect of training type on attitude change ($F(1, 315) = 6.05$, $p = .014$, partial $\eta^2 = .019$), with participants in the personalized condition showing a larger improvement in pre-to-post attitude scores than participants in the generic condition. The corresponding test on behavioral-intention change was not significant ($F(1, 315) = 1.37$, $p = .243$, partial $\eta^2 = .004$).

Table 3. Multivariate Tests for Awareness and Training Type

Statistic	Value	F	Hyp. df	Error df	Sig.	Partial η^2
Pillai's Trace	0.019	3.122	2	314	0.045	0.019
Wilks' Lambda	0.981	3.122	2	314	0.045	0.019
Hotelling's Trace	0.020	3.122	2	314	0.045	0.019
Roy's Largest Root	0.020	3.122	2	314	0.045	0.019

Table 4. Tests of Between-Subjects Effects (Awareness and Training Type)

Dependent Variable	Type III SS	df	Mean Square	F	Sig.	Partial η^2
Attitude change (ATTIDif)	3.229	1	3.229	6.046	0.014	0.019
Behavioral intention (BINT)	0.758	1	0.758	1.368	0.243	0.004

Note. Error degrees of freedom = 315. ATTIDif $R^2 = .019$ (adjusted $R^2 = .016$); BINT $R^2 = .004$ (adjusted $R^2 = .001$).

Estimated marginal means for the two conditions are reported in Table 5. For attitude change, the personalized condition ($M = 0.355$, $SE = 0.058$, 95% CI [0.241, 0.469]) showed roughly twice the pre-to-post gain of the generic condition ($M = 0.153$, $SE = 0.058$, 95% CI [0.039, 0.268]), and the confidence intervals overlap only marginally, consistent with the significant univariate test. For behavioral intention, the two conditions were close and their confidence intervals overlapped substantially (personalized $M = 4.547$, $SE = 0.059$, 95% CI [4.431, 4.663]; generic $M = 4.449$, $SE = 0.059$, 95% CI [4.333, 4.566]). Two features of the intention means are worth noting for interpretation. First, both intention means sit well above the scale midpoint of 4 on the seven-point measure, indicating high baseline intention in both groups and raising the possibility of a ceiling effect. Second, the manipulation-check contrast on perceived job relevance favored the personalized condition, indicating that the framing manipulation was perceived as intended even though it did not move intention. We note that the present analysis files report change-score marginal means rather than separate pre- and post-test raw means and standard deviations by group; the raw pre/post descriptive statistics and the formal baseline-equivalence test are therefore not reproduced here and are flagged in the Limitations section as material we recommend reporting in full in subsequent work.

Table 5. Estimated Marginal Means by Awareness and Training Type

Dependent Variable	Condition	Mean	Std. Error	95% CI Lower	95% CI Upper
Attitude change (ATTIDif)	Generic	0.153	0.058	0.039	0.268
Attitude change (ATTIDif)	Personalized	0.355	0.058	0.241	0.469
Behavioral intention (BINT)	Generic	4.449	0.059	4.333	4.566
Behavioral intention (BINT)	Personalized	4.547	0.059	4.431	4.663

Hypothesis 1 was therefore supported: personalized training produced a larger attitudinal improvement than generic training. Hypothesis 2 was not supported: the two conditions did not differ significantly in pre-to-post change in behavioral intention.

Discussion of Findings

The central finding is asymmetric: a single, content-equivalent personalization manipulation produced a reliable shift in attitude toward information security, but did not produce a detectable shift in behavioral intention. This pattern is consistent with the ELM-derived expectation that role-aligned framing increases central-route processing and yields measurable attitudinal change. It is less consistent with the simplest TPB-derived expectation that any improvement in attitude should propagate proximally into intention. Three interpretations of the asymmetry are worth taking seriously, and they are not mutually exclusive. First, the intention measure may have been subject to a ceiling effect. As Table 5 shows, both groups reported intention means well above the scale midpoint at the outset, and marketing students in a classroom-adjacent research context tend to report high baseline intentions to follow recommended behavior, particularly behavior framed in the language of professional development. When pre-test intentions are already near the top of the scale, even a real treatment effect will be statistically difficult to detect. Second, attitude and intention may have different sensitivity profiles to a single brief intervention. Attitude can update from a single persuasive encounter; intention, especially intention to perform a behavior at some unspecified future point in a context the learner has not yet entered, may require richer cues - an anticipated situation, an implementation intention, a behavioral commitment. Implementation-intention research (Gollwitzer & Sheeran, 2006) suggests that intention formation benefits substantially from situated “if-then” planning that single training videos rarely afford. Third, the intervention itself may have been too brief and too decontextualized to generate transferable intention. Field interventions that combine personalization with on-the-job practice opportunities (Albrechtsen & Hovden, 2010) have produced more robust behavioral effects.

Interpreting the Effect Size

The observed effects are statistically reliable but small in magnitude: the overall multivariate effect and the univariate attitude effect both correspond to a partial eta squared of approximately .019, meaning that training type accounted for roughly 2% of the variance in attitude change. Several considerations bear on how much practical weight this should carry. On one hand, a 2% variance contribution from a single, low-cost framing change, with core content, length, and dosage held constant, is a meaningful return on a manipulation that adds essentially no marginal cost once a role-relevant exemplar library exists, and small per-exposure effects can compound across an organization’s full workforce and across repeated annual cycles. On the other hand, an effect of this size should not be oversold: it indicates that role-aligned framing nudges attitude rather than transforming it, and it leaves the large majority of variance in attitude change unexplained by the manipulation. The estimated marginal means put the effect in concrete terms, roughly a 0.2-point larger gain on a seven-point attitude change metric for the personalized group, which is modest at the individual level. We therefore frame personalization as one inexpensive lever among several rather

than as a stand-alone solution, and we caution against extrapolating from a single-session attitudinal nudge to durable behavioral change.

Implications for Research and Practice

Theoretical Contributions

The study contributes to the behavioral IS literature in three ways. First, it provides direct experimental evidence that role-aligned framing of identical core content is sufficient to shift attitudes toward secure behavior. Most prior IS personalization research has either varied content substantively across conditions (confounding personalization with information quality) or has relied on cross-sectional self-report. By holding core content constant and varying only role framing, this study isolates personalization as a design property and not merely a content property.

Second, it elevates attitude as an outcome of interest in its own right. Much of the IS compliance literature operationalizes attitude only as a predictor of intention and treats intention or compliance as the terminal dependent variable (Ifinedo, 2012; Herath & Rao, 2009). This is theoretically defensible but practically narrow. Attitude is not only a predictor of intention; it is also a meaningful early indicator of an intervention's psychological purchase. A program that fails to shift attitude is unlikely to shift behavior; a program that shifts attitude has, at minimum, established the cognitive precondition for subsequent change. Treating attitude as a primary outcome offers practitioners an earlier, more sensitive read on whether an intervention is doing any psychological work at all.

Third, the study reinforces the value of integrating cognitive-behavioral theories with identity- and affordance-based perspectives at the design stage. PMT, TPB, and ELM specify mechanisms; motivational affordance and role identity specify design properties that activate those mechanisms. The growing IS culture literature (Hovav & D'Arcy, 2012; Furnell & Thomson, 2009; Karjalainen & Siponen, 2011) increasingly insists that secure behavior is sustained by identification and shared values rather than by deterrence alone. Our findings are consistent with that broader reorientation, particularly given that the personalized framing succeeded in shifting attitude without any change in deterrent threats or sanctions.

Practical Implications for IS Program Design

For practitioners, the principal implication is that personalization is an accessible, low-cost lever for improving the psychological yield of awareness training. The manipulation used in this study - role-relevant examples, role-aligned screenshots, and a small number of role-credible quotations did not require new core content, additional instructor time, or expensive learning-technology investment. Organizations with established awareness programs can implement comparable personalization by maintaining a modest library of role-relevant exemplars (for marketing, finance, HR, engineering, customer-facing operations) and assembling modules dynamically based on the learner's department or job title.

Three operational design principles follow. First, exemplars should be drawn from the artifacts and workflows employees actually encounter. For marketing teams, a customer-segmentation spreadsheet is a more credible carrier of a least-access lesson than is a generic financial template. Second, role-credible voices should be recruited where feasible. A quotation from a senior marketing colleague, a marketing professor, or a respected industry practitioner has normative weight that a quotation from the CISO does not, particularly for employees whose primary identification is with their function rather than with security as such. Third, training should be modularized so that personalized framings can be updated as roles, tooling, and threats evolve. A static, year-over-year awareness program is structurally hostile to personalization; a modular content library is structurally hospitable to it.

A second implication concerns measurement. Many organizations evaluate awareness programs solely on completion rates, quiz scores, or downstream behavioral metrics such as phishing-simulation click-through rates. The asymmetric finding here suggests that attitude is worth tracking explicitly. A brief post-training

attitude measure can detect whether content is doing psychological work in advance of any observable change in behavior, and can be used to iterate program design before the lag of behavioral change has resolved. A third implication concerns expectations: the finding that intention did not move in a single brief session suggests that practitioners should not expect a single training module to alter intention reliably and should instead design for repeated, situated, and identity-affirming encounters with the secure behavior, such as just-in-time prompts and peer-to-peer security champions.

Limitations

Several limitations qualify the inferences drawn here, and we group them by the type of caution they impose. The most consequential limitation concerns the absence of measured mediators. The theoretical model draws on PMT, TPB, and ELM, each of which names intervening constructs - perceived vulnerability, response efficacy, subjective-norm strength, and depth of processing - that are hypothesized to carry the personalization effect. None of these mediators was measured in this study. As a result, while the design supports the claim that role-aligned framing changed attitude, it does not identify which mechanism within the PMT-TPB-ELM family produced the change, and it cannot rule out alternative explanations such as simple novelty or increased engagement unrelated to central-route processing. This is a substantive constraint on causal interpretation rather than a minor caveat, and it should temper any mechanism-level reading of the results. We frame the PMT, motivational affordance, and role-identity content of this paper accordingly: as the rationale that motivated the manipulation, not as mechanisms the data adjudicate.

A second set of limitations concerns generalizability, and it bears specifically on the attitude-to-intention relationship. The sample consisted of undergraduate marketing students, not working marketing professionals. Students differ from employees in ways that are directly relevant to the intention measure: they are not embedded in a corporate hierarchy, do not face real organizational sanctions or incentives for (non)compliance, and may treat professional-development-framed items as normatively easy to endorse, which plausibly inflated baseline intention and compressed the attitude-to-intention link that TPB predicts. The null result for intention should therefore be read with the sample in mind; it is consistent both with a genuine ceiling effect and with the possibility that the attitude-to-intention propagation TPB describes is weaker or slower to materialize in a population without authentic organizational stakes. Replication with working professionals across functional domains, finance, engineering, healthcare, manufacturing, using role-appropriate personalization is needed before the findings, and especially the intention null, are generalized to organizational settings.

A last set of limitations concerns the fact that the intervention was also brief and delivered in a single session, which limits inferences about the durability of attitudinal change; ELM predicts that central-route attitudes are more durable, but durability was not directly tested here. Also, outcome measures were self-reported, raising standard concerns about social desirability and common-method variance, though the pre-post within-subjects structure mitigates between-subjects social-desirability bias.

Future Research

Several extensions follow naturally. First, longitudinal field replication in working organizations would test whether role-aligned personalization produces durable change in observed behavior, not only in self-reported attitude. Second, and most directly responsive to the limitations above, mediation tests are needed: a study that directly measures perceived job relevance, perceived vulnerability, perceived response efficacy, subjective-norm strength, and self-reported depth of processing could disentangle the relative contributions of the PMT, TPB, and ELM pathways and convert the present net-effect finding into a mechanism-level account. Third, the boundary conditions of personalization deserve attention, since excessive personalization may signal surveillance and backfire when learners detect that their data has been used to

target them (Aguirre, Mahr, Grewal, de Ruyter, & Wetzels, 2015). Fourth, the interaction between personalization and other behavioral levers, gamification, social-proof prompts, implementation intentions, just-in-time nudges, deserves systematic study, as the strongest interventions likely combine identity-aligned framing with situated practice and timely reinforcement.

Conclusions

Information security failures continue to be driven less by the limits of cryptography or network defense than by the limits of human attention, motivation, and identification. Conventional awareness programs largely treat employees as generic recipients of standard content, which limits both their psychological purchase and their behavioral impact. This study tested whether personalization, holding core content constant and varying only the role-alignment of examples, screenshots, and quotations, could shift the psychological precursors of secure behavior. The intervention produced a reliable, though small, improvement in attitude toward information security and a non-significant change in behavioral intention. Taken together with the broader behavioral IS literature, the finding implies that attitude is a tractable, near-term target for intervention design and a useful indicator of whether an awareness program is doing psychological work, even when intention has not yet moved. The broader implication is that security professionals should design interventions that not only inform employees, but recruit their identification with their roles and their organization, while recognizing that a single low-cost framing change is one modest lever rather than a complete solution.

Acknowledgements

The authors used a generative AI tool to assist with formatting the manuscript to the IIS template and with copy-editing for clarity during revision. All substantive content, analysis, and interpretation are the authors' own, and the authors reviewed and take responsibility for the final text, consistent with the IIS AI Use Policy.

References

- Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing*, 91(1), 34–49.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211.
- Albrechtsen, E., & Hovden, J. (2010). Improving information security awareness and behaviour through dialogue, participation and collective reflection: An intervention study. *Computers & Security*, 29(4), 432–445.
- Anderson, C. L., & Agarwal, R. (2010). Practicing safe computing: A multimethod empirical examination of home computer user security behavioral intentions. *MIS Quarterly*, 34(3), 613–643.
- Ashforth, B. E., Harrison, S. H., & Corley, K. G. (2008). Identification in organizations: An examination of four fundamental questions. *Journal of Management*, 34(3), 325–374.

- Bhattacharjee, A., & Sanford, C. (2006). Influence processes for information technology acceptance: An elaboration likelihood model. *MIS Quarterly*, 30(4), 805–825.
- Boss, S. R., Galletta, D. F., Lowry, P. B., Moody, G. D., & Polak, P. (2015). What do systems users have to fear? Using fear appeals to engender threats and fear that motivate protective security behaviors. *MIS Quarterly*, 39(4), 837–864.
- Bulgurcu, B., Cavusoglu, H., & Benbasat, I. (2010). Information security policy compliance: An empirical study of rationality-based beliefs and information security awareness. *MIS Quarterly*, 34(3), 523–548.
- Crossler, R. E., Johnston, A. C., Lowry, P. B., Hu, Q., Warkentin, M., & Baskerville, R. (2013). Future directions for behavioral information security research. *Computers & Security*, 32, 90–101.
- D’Arcy, J., Hovav, A., & Galletta, D. (2009). User awareness of security countermeasures and its impact on information systems misuse: A deterrence approach. *Information Systems Research*, 20(1), 79–98.
- Furnell, S., & Thomson, K.-L. (2009). From culture to disobedience: Recognising the varying user acceptance of IT security. *Computer Fraud & Security*, 2009(2), 5–10.
- Gefen, D., & Straub, D. (2005). A practical guide to factorial validity using PLS-Graph: Tutorial and annotated example. *Communications of the Association for Information Systems*, 16(1), 91–109.
- Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, 38, 69–119.
- Greenberg, A. (2021, March 9). Hackers breach thousands of security cameras, exposing Tesla, jails, and hospitals. *Wired*. <https://www.wired.com/story/verkada-hack-security-cameras-tesla-cloudflare/>
- Herath, T., & Rao, H. R. (2009). Encouraging information security behaviors in organizations: Role of penalties, pressures and perceived effectiveness. *Decision Support Systems*, 47(2), 154–165.
- Hovav, A., & D’Arcy, J. (2012). Applying an extended model of deterrence across cultures: An investigation of information systems misuse in the U.S. and South Korea. *Information & Management*, 49(2), 99–110.
- Ifinedo, P. (2012). Understanding information systems security policy compliance: An integration of the theory of planned behavior and the protection motivation theory. *Computers & Security*, 31(1), 83–95.
- Johnston, A. C., & Warkentin, M. (2010). Fear appeals and information security behaviors: An empirical study. *MIS Quarterly*, 34(3), 549–566.
- Karjalainen, M., & Siponen, M. (2011). Toward a new meta-theory for designing information systems (IS) security training approaches. *Journal of the Association for Information Systems*, 12(8), 518–555.

- Loch, K. D., Straub, D. W., & Kamel, S. (2003). Diffusing the Internet in the Arab world: The role of social norms and technological cultururation. *IEEE Transactions on Engineering Management*, 50(1), 45–63.
- Maddux, J. E., & Rogers, R. W. (1983). Protection motivation and self-efficacy: A revised theory of fear appeals and attitude change. *Journal of Experimental Social Psychology*, 19(5), 469–479.
- Microsoft. (n.d.). Protection and security in Excel. <https://support.microsoft.com/en-us/office/protection-and-security-in-excel-be0b34db-8cb6-44dd-a673-0b3e3475ac2d>
- Pahnila, S., Siponen, M., & Mahmood, A. (2007). Employees' behavior towards IS security policy compliance. In *Proceedings of the 40th Hawaii International Conference on System Sciences* (pp. 156b–156b). IEEE.
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19, 123–205.
- Posey, C., Roberts, T. L., & Lowry, P. B. (2013). The impact of organizational commitment on insiders' motivation to protect organizational information assets. *Journal of Management Information Systems*, 29(4), 239–274.
- Puhakainen, P., & Siponen, M. (2010). Improving employees' compliance through information systems security training: An action research study. *MIS Quarterly*, 34(4), 757–778.
- Renaud, K., & Goucher, W. (2014). The curious incidence of security breaches by knowledgeable employees and the pivotal role of security culture. In *Human Aspects of Information Security, Privacy, and Trust* (pp. 361–372). Springer.
- Rogers, R. W. (1975). A protection motivation theory of fear appeals and attitude change. *The Journal of Psychology*, 91(1), 93–114.
- Sen, R., & Borle, S. (2015). Estimating the contextual risk of data breach: An empirical approach. *Journal of Management Information Systems*, 32(2), 314–341.
- Sheeran, P., & Webb, T. L. (2016). The intention–behavior gap. *Social and Personality Psychology Compass*, 10(9), 503–518.
- Siponen, M., & Vance, A. (2010). Neutralization: New insights into the problem of employee information systems security policy violations. *MIS Quarterly*, 34(3), 487–502.
- Stryker, S. (1980). *Symbolic interactionism: A social structural version*. Benjamin/Cummings.
- Verkada Inc. (2021, March 9). Security incident report. <https://www.verkada.com/security-update/report/>
- Witte, K. (1992). Putting the fear back into fear appeals: The extended parallel process model. *Communication Monographs*, 59(4), 329–349.

Workman, M., Bommer, W. H., & Straub, D. (2008). Security lapses and the omission of information security measures: A threat control model and empirical test. *Computers in Human Behavior*, 24(6), 2799–2816.

Zhang, P. (2008). Motivational affordances: Reasons for ICT design and use. *Communications of the ACM*, 51(11), 145–147.