American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Social Engineering Vulnerability in Multigenerational Workforces

Dr. Ananya Rao

Assistant Professor, Savitribai Phule Pune University, India

Abstract

Social engineering remains a persistent organizational cybersecurity challenge because attackers exploit trust, authority, urgency, familiarity, distraction, fear, reciprocity, and other characteristics of human decision-making rather than relying exclusively on technical vulnerabilities. Contemporary workplaces complicate this problem because employees from different age cohorts work together while possessing different combinations of technological experience, occupational seniority, communication preferences, digital routines, cybersecurity self-efficacy, and exposure to particular attack channels. This study examines social engineering vulnerability within multigenerational workforces while challenging the assumption that chronological age alone provides a reliable basis for predicting employee susceptibility. Recent empirical evidence is inconsistent: some studies have reported greater phishing susceptibility among older employees, other organizational simulations have identified comparatively greater vulnerability among younger workers, and further research has found that age does not independently predict susceptibility after behavioral and psychological factors are considered. The present study therefore develops a context-sensitive model in which vulnerability is shaped by attack type, security self-efficacy, workload, employment context, reporting behavior, training relevance, and organizational safeguards.

Because no authentic workforce dataset was supplied, a transparent simulation involving 320 hypothetical employees across four age-cohort categories was used for methodological demonstration. Five social-engineering dimensions were modeled: credential phishing, voice and SMS pretexting, authority and business-email-compromise manipulation, multifactor-authentication prompt fatigue, and delayed incident reporting. The simulated profiles reveal distinct rather than uniformly hierarchical vulnerability patterns. Early-career employees display comparatively higher simulated vulnerability to credential phishing and authentication-prompt fatigue, whereas the oldest modeled group shows higher simulated exposure to voice/SMS pretexting and reporting delay. The findings demonstrate why organizations should avoid age-based cybersecurity stereotypes and instead adopt adaptive security awareness, phishing-resistant technical controls, context-sensitive training, psychologically safe reporting, and behavioral risk assessment. The study contributes a multigenerational human-risk framework suitable for empirical validation in organizational cybersecurity research.

Keywords: social engineering; multigenerational workforce; phishing susceptibility; human factors; cybersecurity awareness; employee behavior; security training; human risk management

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

1. Introduction

Social engineering represents a distinctive cybersecurity problem because its effectiveness depends upon manipulation of human judgment. Rather than directly defeating a technical control, an attacker may persuade an employee to disclose credentials, approve an authentication request, open a malicious attachment, reveal organizational information, transfer funds, change payment details, or provide access to a supposedly trusted individual. CISA identifies phishing as a form of social engineering and emphasizes that attackers commonly impersonate trusted entities to persuade targets to disclose information or take unsafe actions. Modern social engineering also extends beyond conventional email. Attackers may combine telephone calls, SMS messages, social-media contact, identity impersonation, help-desk manipulation, multifactor-authentication fatigue, and other channels. CISA's advisory concerning Scattered Spider, for example, documents the use of several social-engineering techniques, including repeated authentication prompts and SIM-related attacks.

The workforce in which these attacks occur is increasingly multigenerational. Employees may differ substantially in occupational experience, familiarity with legacy and emerging technologies, preferred communication channels, confidence in digital systems, managerial authority, job responsibilities, and frequency of exposure to particular forms of digital interaction. It is tempting to interpret these differences through a simple age-risk model, but empirical research does not justify such a conclusion. Greitzer and colleagues' large-scale employee study noted that prior evidence concerning age and phishing susceptibility had been inconsistent and demonstrated that demographic and organizational variables can interact with attack outcomes. More recent evidence has strengthened the need for caution when making age-based generalizations.

Ribeiro, Guedes, and Cardoso's empirical investigation found that individual demographic characteristics such as age and gender did not independently predict phishing susceptibility, whereas phishing-detection self-efficacy showed a meaningful protective relationship. In contrast, Oner, Cetin, and Savas reported in a 2020 organizational study that older employees showed greater phishing susceptibility and were less likely to report phishing under some conditions. Yet another recent organizational investigation reported comparatively greater vulnerability among employees aged approximately 20–34. These apparently contradictory findings are not necessarily methodological failures. They indicate that social-engineering vulnerability is likely to be conditional upon attack design, organizational context, job role, technological familiarity, reporting culture, and employee experience rather than determined by age alone.

Other research supports this contextual interpretation. Frank and colleagues found that phishing susceptibility was shaped by social, task, and physical work contexts, including help-desk reliance, team size, job experience, and job status. Beu and colleagues similarly found in a naturalistic phishing simulation that shorter organizational tenure and lower employee satisfaction and loyalty were associated with increasingly unsafe behavior. Frauenstein and Flowerday found that heuristic information processing can increase phishing susceptibility, reinforcing the importance of how employees process messages rather than merely which demographic category they occupy. These studies suggest that attackers succeed by exploiting situational vulnerabilities that may appear differently across workforce groups.

The present study therefore examines Social Engineering Vulnerability in Multigenerational Workforces through a human-centered and non-stereotypical framework. The paper does not assume that one generation constitutes an organization's "weakest link." Instead, it proposes that different workforce cohorts may present different vulnerability profiles across different attack vectors. Because real organizational participant data were not supplied, the quantitative component is explicitly simulation-based. A synthetic sample of 320 hypothetical employees is used to illustrate how multigenerational vulnerability could be measured across credential phishing, voice/SMS pretexting, authority-based manipulation, multifactor-authentication fatigue, and incident-reporting behavior. The study's contribution lies in shifting cybersecurity analysis from the question of which generation is most vulnerable toward the more useful question of which contextual and behavioral conditions make particular employees vulnerable to particular forms of social engineering.

2. Objectives of the Study

2.1 To Examine Social Engineering Vulnerability Across Workforce Cohorts

The first objective is to evaluate whether different age-cohort groups may exhibit different patterns of susceptibility across multiple social-engineering vectors rather than a single uniform vulnerability hierarchy. The analysis distinguishes credential phishing, voice/SMS pretexting, authority manipulation, multifactor-authentication fatigue, and delayed reporting.

2.2 To Identify Behavioral and Organizational Drivers Beyond Chronological Age

The second objective is to develop a framework in which social-engineering vulnerability is influenced by cybersecurity self-efficacy, job context, workload, organizational tenure, communication habits, security culture, reporting confidence, and the relevance of awareness training. This objective responds directly to mixed empirical evidence concerning the predictive value of age.

2.3 To Develop Inclusive Multigenerational Cybersecurity Interventions

The third objective is to translate the analysis into practical organizational measures that improve security without stigmatizing particular age groups. The study emphasizes adaptive training, technical protections, contextual simulations, accessible reporting mechanisms, immediate feedback, and differentiated but equitable cybersecurity support.

3. Review of Literature

3.1 Human Factors and Social Engineering Susceptibility

Social engineering exploits cognitive and social processes that employees routinely depend upon for productive work. Trust in colleagues, responsiveness to authority, willingness to assist others, rapid communication, and compliance with urgent requests are generally desirable organizational behaviors. Attackers can convert these same behaviors into security vulnerabilities by designing messages and interactions that appear legitimate. Washo's interdisciplinary treatment of social engineering emphasizes that manipulation and deception should be understood through technological, psychological, organizational, and ethical perspectives rather than as an exclusively technical phenomenon.

Phishing research provides substantial evidence of this behavioral complexity. Ribeiro and colleagues found that phishing-detection self-efficacy reduced susceptibility while demographic variables, including age, did not independently predict outcomes in their empirical model. Their findings are significant for multigenerational workforce research because they suggest that an employee's confidence and competence in recognizing suspicious communications may provide more useful intervention targets than demographic classification.

Frauenstein and Flowerday's analysis of phishing on social-network sites similarly found that heuristic processing increased susceptibility. Employees under time pressure or dealing with familiar-looking messages may rely on rapid cognitive shortcuts rather than deliberate verification. This mechanism can affect any age group, although the communication context in which it occurs may differ between employees.

Carroll and colleagues' experimental work further showed that participants generally found modern phishing attacks difficult to detect and tended to rely on visible cues such as suspicious requests and inconsistencies in communication. This evidence is particularly important in contemporary environments where technically polished or AI-assisted messages may contain fewer traditional warning signs.

3.2 Multigenerational Differences and the Problem of Age Stereotyping

Research specifically examining age and phishing behavior produces mixed results. Greitzer and colleagues analyzed phishing susceptibility and awareness among more than 8,000 employees and concluded that demographic and organizational factors influenced outcomes, while also noting that previous studies concerning age and gender were inconsistent. Such inconsistency indicates that chronological age should not automatically be interpreted as a stable cybersecurity-risk factor.

Ribeiro and colleagues later found no predictive effect of age in their empirical phishing-susceptibility study. By comparison, Oner, Cetin, and Savas' 2020 research within a large European financial organization reported comparatively greater susceptibility among older employees and differences in phishing-reporting behavior. Another recent integrated organizational audit found greater vulnerability among younger employees aged approximately 20–34.

These findings strongly caution against describing any generation as inherently cybersecure or inherently vulnerable. A younger employee may possess extensive familiarity with digital interfaces while simultaneously demonstrating greater platform switching, higher exposure to mobile prompts, or stronger confidence that encourages rapid interaction. An older employee may possess substantial organizational knowledge and experience detecting unusual business requests while being less familiar with particular authentication mechanisms or communication channels. These are plausible contextual mechanisms rather than universal generational traits, and they require empirical testing in each organization.

Lin and colleagues' spear-phishing research also found that age-related patterns may change over repeated exposure and that different attack designs can affect age groups differently. This strengthens the argument that vulnerability is attack-specific. An employee who successfully detects email credential phishing may still be vulnerable to voice impersonation, help-desk manipulation, malicious QR codes, MFA fatigue, or authority-based business-email compromise.

3.3 Training, Reporting, and Organizational Human-Risk Management

Cybersecurity education remains an important defense, but research demonstrates that its effectiveness depends upon design. Prümmer, van Steen, and van den Berg's systematic review found that most evaluated cybersecurity training interventions reported beneficial effects, although methodological limitations remained common and many studies involved populations other than employees. Jampen and colleagues likewise concluded that awareness and training can strengthen phishing recognition, while also documenting that previously trained individuals may still remain vulnerable.

Marshall and colleagues' review found that phishing training was associated with a substantial reduction in susceptibility, but trained participants still fell for a meaningful proportion of phishing messages. Training should therefore be treated as one layer within a broader defense system rather than as evidence that employees have become immune to manipulation.

More recent experimental work by Madeira and colleagues found that immediate feedback can reduce phishing susceptibility and highlighted the importance of training format, organizational participation, and motivation. CISA similarly recommends organizational anti-phishing programs that combine employee education, simulated attacks, analysis of results, reporting practices, and technical protections.

4. Research Methodology

4.1 Research Design and Synthetic Sample

The study adopts a simulation-based comparative research design. No employees were recruited, no actual phishing exercise was conducted, and no organizational security records were processed. The quantitative analysis is therefore a methodological demonstration rather than a report of empirical workforce behavior.

A synthetic sample of 320 hypothetical employees was constructed and divided evenly among four analytical age cohorts: early-career employees aged 21–29, established-career employees aged 30–44, mid/late-career employees aged 45–60, and an older-workforce group aged 61 years and above. Each group contained 80 hypothetical observations.

The age categories are used only as research strata and should not be interpreted as deterministic generational identities. Individuals within the same chronological group can differ widely in technology experience, educational background, job tenure, role, cybersecurity capability, communication style, and exposure to attack vectors.

Five simulated vulnerability domains were created: credential phishing, voice/SMS pretexting, authority or business-email-compromise manipulation, multifactor-authentication prompt fatigue, and delayed incident reporting. Scores were expressed on a 0–100 vulnerability index, where a higher value indicates greater hypothetical susceptibility or behavioral exposure.

4.2 Proposed Measurement Instrument

A future empirical study should use a combination of validated behavioral measures, simulated social-engineering exercises, objective reporting data, and employee survey instruments. The following five questions are proposed as a concise screening instrument and were not administered in the present study.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Table 1: Proposed Five-Item Questionnaire for Multigenerational Social Engineering Vulnerability

Item	Proposed Questionnaire Statement	Primary Construct
Q1	I verify unexpected requests for credentials, payments, sensitive information, or account changes through an independent communication channel.	Verification behavior
Q2	I can recognize common psychological tactics such as urgency, authority, fear, familiarity, and unusual pressure in suspicious communications.	Social-engineering awareness
Q3	I am confident that I can identify suspicious emails, text messages, telephone calls, and authentication requests before taking action.	Security self-efficacy
Q4	I would report a suspected social-engineering attempt promptly even if I had already clicked, replied, approved a request, or disclosed information.	Reporting confidence
Q5	The cybersecurity training I receive reflects the communication tools, work situations, and attack methods that I encounter in my actual job.	Training relevance

Source: Proposed by the author for future empirical validation. The items were not administered to actual employees.

A future validation study should examine internal consistency, construct validity, measurement invariance across age cohorts, and relationships between self-reported responses and observable simulated-attack behavior. Reliance exclusively on self-reported security competence should be avoided because perceived ability and actual performance may differ.

4.3 Analytical Procedure and Research Ethics

The simulated analysis compares the four workforce cohorts across the five social-engineering vulnerability domains. Aggregate vulnerability was calculated as the mean of the five domain scores for descriptive comparison. Rather than deliberately creating a monotonic age effect, the simulation was designed to represent different vulnerability configurations across attack types.

This choice reflects the mixed findings of existing empirical research concerning age and phishing susceptibility. The simulated values are therefore not intended to estimate population susceptibility. Their purpose is to demonstrate how future workforce studies could avoid reducing multidimensional human risk to a single demographic variable.

No research-ethics approval was required for the synthetic demonstration because no human participants or identifiable employment records were used. Actual organizational social-engineering research requires substantially greater ethical care. Simulated phishing can create embarrassment, anxiety, employment concerns, or a perception of surveillance if implemented punitively. Future studies should therefore obtain appropriate institutional approval, provide proportionate safeguards, minimize unnecessary collection of personal data, restrict access to individual-level results, and ensure that

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

security exercises are designed primarily for learning and risk reduction rather than employee punishment.

5. Analysis

5.1 Comparative Vulnerability Across Workforce Cohorts

The simulation does not produce a single age-based hierarchy. Instead, each cohort displays a distinct risk profile. The early-career group shows a mean aggregate vulnerability of approximately 46.0%, the established-career group 40.5%, the mid/late-career group 43.3%, and the older-workforce group 46.4%. The difference between the lowest and highest aggregate means is therefore modest compared with the differences observed within individual attack vectors. This result is intentional and consistent with the theoretical proposition that social-engineering vulnerability should be understood dimensionally rather than through age alone.

Table 2: Simulated Social Engineering Vulnerability Across Workforce Cohorts

Workforce Cohort	Credential Phishing	Voice/SMS Pretexting	Authority/BEC Manipulation	MFA/Prompt Fatigue	Reporting Delay	Aggregate Vulnerability
Early-career, 21–29	50.8%	38.2%	42.7%	55.4%	43.1%	46.0%
Established, 30–44	40.6%	36.5%	43.5%	46.2%	35.9%	40.5%
Mid/late-career, 45–60	42.3%	44.8%	48.9%	38.7%	41.6%	43.3%
Older workforce, 61+	46.7%	52.4%	49.8%	35.6%	47.3%	46.4%

Source: Author-generated synthetic data; total simulated n = 320. These percentages do not represent real employees or observed attack outcomes.

The early-career cohort records the highest simulated score for MFA/prompt fatigue at 55.4%, while the oldest modeled group records the highest score for voice/SMS pretexting at 52.4%. Authority/BEC manipulation is comparatively elevated in the mid/late-career and older cohorts, while delayed reporting is highest in the oldest simulated group.

The established-career group records the lowest overall synthetic vulnerability, but this should not be interpreted as evidence that employees aged 30–44 are inherently safer. The values are generated for methodological illustration only.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

5.2 Attack-Specific Rather Than Generation-Specific Vulnerability

The most important result is the variation *within* each cohort. Early-career employees, for example, do not show the greatest simulated vulnerability across all five domains. Their highest exposure concerns authentication-prompt fatigue, whereas voice/SMS pretexting is substantially lower.

Similarly, the older-workforce group has comparatively elevated voice/SMS vulnerability and reporting delay but the lowest simulated MFA/prompt-fatigue score at 35.6%. This result illustrates why an organization using only an aggregate phishing click rate could miss meaningful differences in how employees interact with different forms of manipulation.

The same principle applies to mid/late-career employees. Their simulated credential-phishing vulnerability is moderate, but authority/BEC manipulation reaches 48.9%, which may be especially relevant for employees whose seniority gives them responsibility for financial approvals, supplier communication, personnel decisions, or access authorization.

5.3 Multidimensional Vulnerability Profile

Graph 1: Simulated Social-Engineering Vulnerability Profiles Across Workforce Cohorts

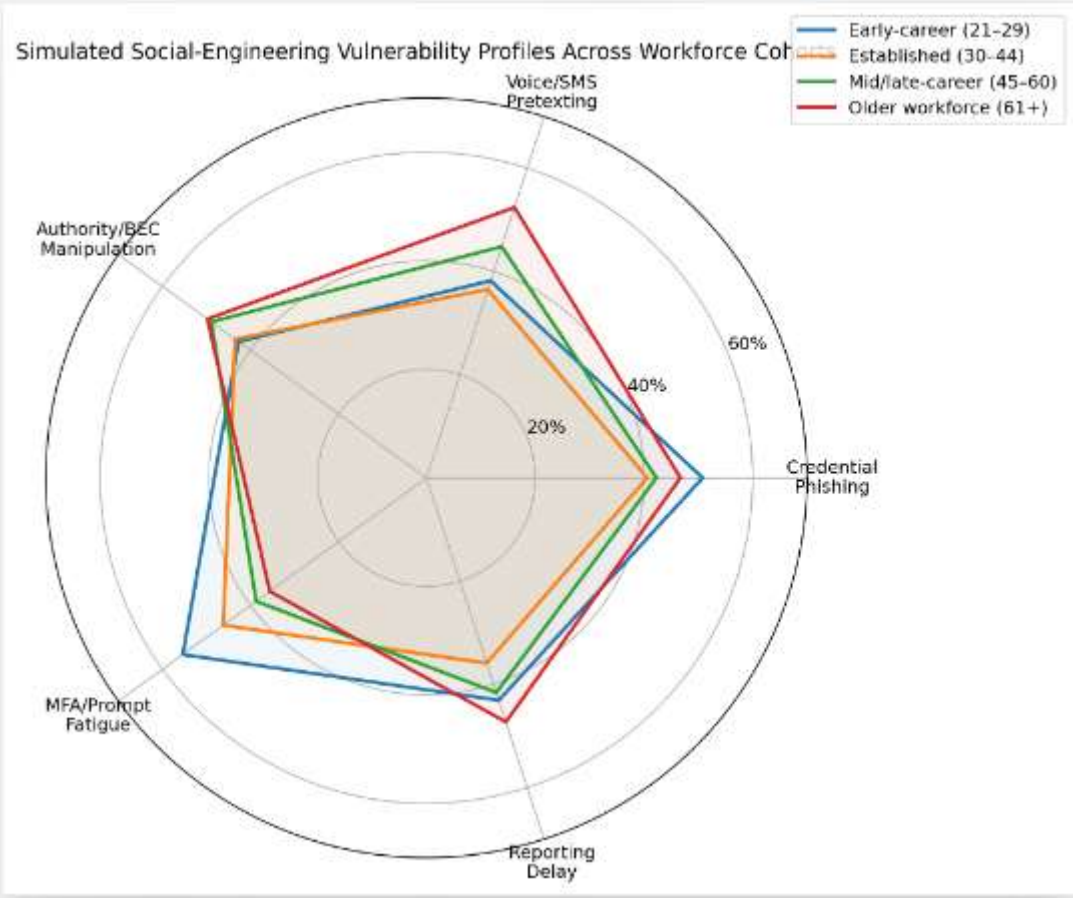

Source: Author-generated synthetic analysis.

Interpretation: The radar graph demonstrates that the four workforce cohorts have overlapping but differently shaped simulated vulnerability profiles. No cohort dominates every risk dimension.

Key Finding: The analytical value lies in the *shape* of vulnerability rather than the ranking of generations. Different attack channels may require different defensive interventions even when overall vulnerability scores appear similar.

The simulated pattern aligns with empirical literature indicating that age effects vary across studies and that context, confidence, work experience, organizational conditions, and attack characteristics materially influence susceptibility.

6. Discussion

6.1 Multigenerational Cybersecurity Requires Moving Beyond the "Weakest Generation" Assumption

The principal implication of this study is that organizations should reject attempts to identify a single "weakest generation" in cybersecurity. Existing empirical evidence does not support such a stable classification. Ribeiro and colleagues found no independent age effect on phishing susceptibility, Oner and colleagues reported greater vulnerability among older employees in a financial organization, and another recent organizational assessment identified greater vulnerability among younger workers. Greitzer and colleagues similarly documented inconsistent age effects across the broader research literature. A defensible conclusion is therefore not that one study must be wrong but that susceptibility is conditional upon context and measurement.

This distinction has practical consequences. Age-based cybersecurity stereotypes can create false confidence among employees who are assumed to be technologically competent while simultaneously stigmatizing those assumed to be digitally inexperienced. Both outcomes are undesirable. An employee categorized as "digitally native" may still respond impulsively to repeated authentication requests or malicious mobile messages. A senior employee with decades of organizational experience may identify fraudulent business requests rapidly because they understand normal operational patterns. The relevant question is therefore not whether someone belongs to a particular generation but whether the organization's current threat environment intersects with that employee's responsibilities, communication habits, cognitive workload, and available safeguards.

The simulation reinforces this reasoning by producing similar aggregate vulnerability levels for early-career and older-workforce groups while showing substantially different attack-specific patterns. Such profiles would require different interventions. Training designed around authentication fatigue may be useful for one group, while voice-impersonation verification procedures could be more relevant for another. These distinctions should be based on observed organizational behavior rather than assumptions attached to age.

6.2 Adaptive Training and Technical Controls Should Operate Together

The findings also indicate that awareness training should be differentiated according to risk context, not simply according to employee age. Reviews of cybersecurity training generally find positive effects but also demonstrate considerable variation in intervention quality and research methodology. Training remains necessary because technical systems cannot prevent every social-engineering interaction. However, an annual generic module containing identical examples for every employee is unlikely to reflect the diversity of modern attack channels and workplace responsibilities.

A more effective model would use observed attack patterns to select training scenarios. Employees frequently exposed to supplier-payment communication could receive business-email-compromise and verification exercises. Help-desk personnel could receive impersonation and account-recovery scenarios. Employees heavily dependent upon mobile communication could receive smishing, QR-code, and MFA-fatigue exercises. Executives and senior managers could receive authority-impersonation and high-value targeted phishing scenarios. These assignments would be based on role and exposure rather than assumptions that age determines capability.

Immediate and constructive feedback should also accompany simulations. Madeira and colleagues' 2022 research indicates that immediate feedback can reduce phishing susceptibility and that training effectiveness is influenced by intervention design and organizational culture. CISA likewise recommends combining awareness, simulation, results analysis, and technical safeguards within anti-phishing programs.

Training must nevertheless be supported by controls that reduce the consequences of human error. Phishing-resistant authentication, secure email controls, payment-change verification, privileged-access restrictions, independent confirmation channels, device protections, and clear escalation processes can prevent an individual mistake from becoming an organizational compromise. Human-centered cybersecurity does not mean making employees exclusively responsible for defeating attackers; it means designing systems around the reality that human judgment is fallible.

6.3 Reporting Culture and Human Risk Management Are Central to Resilience

An organization's response after an employee recognizes—or suspects—that an error has occurred can be as important as prevention. Attackers often benefit when victims delay reporting because they are embarrassed, uncertain, or afraid of consequences. A punitive security culture can therefore transform a recoverable mistake into a more serious incident.

This concern is particularly relevant to multigenerational teams because employees may differ in their expectations concerning authority, error disclosure, managerial communication, and help-seeking. Such differences should not be assumed from age, but organizations should recognize that reporting behavior is influenced by psychological safety and organizational culture. Security policies should make it clear that rapid reporting after an unsafe action is a protective behavior rather than evidence of personal failure.

The simulated results place reporting delay among the five core vulnerability dimensions for this reason. An employee who clicks a suspicious link but immediately reports the incident may enable rapid credential revocation or endpoint isolation. An employee who conceals the same action can unintentionally expand the attacker's window of opportunity.

The emerging concept of human-risk management provides a useful organizational framework. Nurse and colleagues found that practitioners increasingly conceptualize human cybersecurity risk as a broader problem involving people, behavioral data, systems, and organizational environments rather than simply an awareness-training requirement. Multigenerational cybersecurity should adopt this systems perspective. The goal should not be to identify which demographic group makes the most mistakes but to determine where work processes, technical design, communication norms, role privileges, attack exposure, and security culture combine to create preventable risk.

7. Conclusion

Social engineering remains difficult to control because it exploits ordinary human and organizational behaviors rather than depending entirely upon technical vulnerabilities. Employees must routinely trust colleagues, respond rapidly, obey legitimate authority, process large volumes of communication, and make decisions with incomplete information. Attackers deliberately manipulate these characteristics through phishing, impersonation, voice and SMS pretexting, business-email compromise, authentication fatigue, and related techniques. CISA's contemporary guidance consequently treats phishing defense as a combination of recognition, reporting, organizational preparation, and technical controls rather than a matter of employee awareness alone.

The multigenerational character of modern workforces adds complexity but should not be interpreted through simplistic age stereotypes. Current research provides contradictory findings regarding age: some empirical studies report greater vulnerability among older employees, others identify comparatively greater vulnerability among younger employees, and further research finds no independent age effect once other factors are considered. The more defensible interpretation is that social-engineering susceptibility is produced by interactions among attack design, self-efficacy, job responsibilities, organizational experience, workload, communication channel, security safeguards, training quality, and reporting culture.

The synthetic analysis developed in this study illustrates this multidimensional position. Across 320 hypothetical employee observations, aggregate vulnerability differed only moderately among the four workforce cohorts, while much stronger contrasts appeared across individual attack vectors. Early-career employees displayed higher simulated exposure to authentication-prompt fatigue, whereas the oldest modeled cohort displayed higher voice/SMS pretexting and delayed-reporting vulnerability. These numerical results are not empirical workforce findings and must not be generalized to real organizations. Their purpose is to demonstrate that security teams can obtain more useful information from attack-specific vulnerability profiles than from a single demographic ranking.

Future empirical research should test this framework using ethically governed phishing simulations, vishing or smishing scenarios, behavioral reporting measures, validated self-efficacy scales, organizational-context variables, and longitudinal training assessments. Studies should examine whether age effects disappear or change when job role, tenure, workload, digital confidence, authority level, security culture, and attack channel are controlled. Organizations should meanwhile adopt inclusive human-risk management, role-sensitive training, immediate learning feedback, psychologically safe incident reporting, and technical safeguards that limit the consequences of unavoidable human error. Multigenerational cybersecurity is most effective when employee diversity is treated as a source of different experiences and defensive strengths rather than as a hierarchy of assumed vulnerability.

References

1. Hadnagy, C. (2018). *Social Engineering: The Science of Human Hacking* (2nd ed.). Wiley.
2. Mitnick, K. D., & Simon, W. L. (2002). *The Art of Deception: Controlling the Human Element of Security*. Wiley.

3. Workman, M. (2008). Wisecrackers: A theory-grounded investigation of phishing and pretext social engineering threats to information security. *Journal of the American Society for Information Science and Technology, 59*(4), 662–674. https://doi.org/10.1002/asi.20779
4. Jakobsson, M., & Myers, S. (Eds.). (2006). *Phishing and Countermeasures: Understanding the Increasing Problem of Electronic Identity Theft*. Wiley.
5. Sheng, S., Holbrook, M., Kumaraguru, P., Cranor, L. F., & Downs, J. (2010). Who falls for phish? A demographic analysis of phishing susceptibility. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 373–382. https://doi.org/10.1145/1753326.1753383
6. Downs, J. S., Holbrook, M. B., & Cranor, L. F. (2007). Behavioral response to phishing risk. *Proceedings of the Anti-Phishing Working Groups 2nd Annual eCrime Researchers Summit*, 37–44.
7. Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L. F., & Hong, J. (2010). Teaching Johnny not to fall for phish. *ACM Transactions on Internet Technology, 10*(2), 1–31. https://doi.org/10.1145/1754393.1754396
8. Arachchilage, N. A. G., & Love, S. (2014). Security awareness of computer users: A phishing threat avoidance perspective. *Computers in Human Behavior, 38*, 304–312. https://doi.org/10.1016/j.chb.2014.05.046
9. Parsons, K., Butavicius, M., Delfabbro, P., & Lillie, M. (2015). Predicting susceptibility to social engineering attacks in organizations. *Proceedings of the 23rd European Conference on Information Systems*, 1–12.
10. Greitzer, F. L., Kangas, L. J., Noonan, C. F., Dalton, A. C., & Hohimer, R. E. (2014). Identifying at-risk employees using psychometric methods. *Journal of Information Security and Applications, 19*(4–5), 201–211.
11. Parsons, K., McCormac, A., Butavicius, M., Pattinson, M., & Jerram, C. (2014). Determining employee awareness using the Human Aspects of Information Security Questionnaire (HAIS-Q). *Computers & Security, 42*, 165–176. https://doi.org/10.1016/j.cose.2013.12.003
12. Bada, M., Sasse, A. M., & Nurse, J. R. C. (2019). Cyber security awareness campaigns: Why do they fail to change behaviour? *arXiv preprint arXiv:1901.02672*.
13. Hadlington, L. (2017). Human factors in cybersecurity; examining the link between Internet addiction, impulsivity, attitudes towards cybersecurity, and risky cybersecurity behaviours. *Heliyon, 3*(7), e00346. https://doi.org/10.1016/j.heliyon.2017.e00346
14. Egelman, S., & Peer, E. (2015). Scaling the security wall: Developing a security behavior intentions scale (SeBIS). *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2873–2882. https://doi.org/10.1145/2702123.2702247
15. Ifinedo, P. (2012). Understanding information systems security policy compliance: An integration of the theory of planned behavior and protection motivation theory. *Computers & Security, 31*(1), 83–95. https://doi.org/10.1016/j.cose.2011.10.007
16. 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.
17. Vroom, C., & von Solms, R. (2004). Towards information security behavioural compliance. *Computers & Security, 23*(3), 191–198. https://doi.org/10.1016/j.cose.2004.01.012

18. Siponen, M., Mahmood, M. A., & Pahnila, S. (2014). Employees' adherence to information security policies: An exploratory field study. *Information & Management, 51*(2), 217–224. https://doi.org/10.1016/j.im.2013.08.006
19. Boss, S. R., Kirsch, L. J., Angermeier, I., Shingler, R. A., & Boss, R. W. (2009). If someone is watching, I'll do what I'm asked: The influence of social presence on information security policy compliance. *Information Systems Journal, 19*(6), 517–537. https://doi.org/10.1111/j.1365-2575.2008.00284.x
20. Furnell, S., & Vasileiou, I. (2017). Security education and awareness: Just let them burn? *Network Security, 2017*(12), 5–9. https://doi.org/10.1016/S1353-4858(17)30111-1