

The Sociotechnical Cost of Vigilance: A Human Factors Analysis of Burnout, Security Fatigue, and the AI/ML Paradox in Cybersecurity

Calvin Nobles

*School of Cybersecurity and Information Technology, University of Maryland Global Campus, Adelphi, MD, USA
calvin.nobles@umgc.edu*

Abstract: Cybersecurity professionals operate in adversarial, rapidly changing, and technologically complex environments that demand sustained attention, rapid judgment, continuous learning, and cross-functional coordination, exposing them to heavy workloads, alert fatigue, interrupted recovery, and limited control. Although current evidence does not establish a direct causal pathway from practitioner burnout to cybersecurity incidents, cybersecurity-specific and broader occupational research identifies credible risks to cognition, coordination, retention, and human reliability. This position paper argues that organizations should integrate burnout and fatigue prevention into cybersecurity governance rather than treat prevention as an optional wellness benefit. Burnout can cascade across the enterprise by slowing response, transferring emergency labor to non-cybersecurity personnel, and encouraging restrictive controls that generate workforce-wide security fatigue. Artificial intelligence and machine learning intensify the challenge through a dual pathway: they can automate repetitive work and augment analysis, yet they can also create technostress, verification burdens, reskilling demands, productivity intensification, and accountability ambiguity. The paper reframes burnout prevention as a human-reliability and governance control, introduces two analytic constructs, the cascading human-risk loop and the verification tax, and proposes a five-layer prevention model spanning demand reduction, resource amplification, cognitive ergonomics, recovery protection, and governance, supported by an implementation pathway and a multi-domain evaluation framework.

Keywords: Cybersecurity Burnout, Fatigue, Human Factors, Artificial Intelligence, Machine Learning, Technostress, Organizational Resilience

Introduction

Cybersecurity professionals do not work in stable environments with predictable workloads; they operate in adversarial sociotechnical systems in which threats evolve continuously, vulnerabilities emerge without regard to staffing capacity, and routine work can escalate into a high-consequence incident within minutes. Analysts, incident responders, engineers, and security leaders must interpret incomplete signals while balancing operational urgency, legal obligations, and organizational risk tolerance (Dykstra & Paul, 2018; Nepal et al., 2024; Nobles, 2022). The technological environment compounds these demands: contemporary security work requires simultaneous interaction with monitoring, detection, identity, cloud, vulnerability, and threat-intelligence platforms, obliging practitioners to distinguish meaningful indicators from noise, reconstruct context across interfaces, and justify actions to stakeholders with competing priorities. Alert fatigue research identifies high alert volumes, false positives, prioritization difficulties, fragmented tools, and continuous triage as central sociotechnical problems rather than merely technical deficiencies (Tariq et al., 2025). Few cybersecurity ecosystems, however, are deliberately engineered to accommodate human cognitive limitations.

Cybersecurity-specific burnout research remains less mature than the corresponding literature in healthcare and human services. Nevertheless, Nepal et al. (2024) found that 19 of 35 cybersecurity incident responders experienced burnout associated with elevated workload, limited control, poor teamwork, inadequate recognition, poor sleep, and work exceeding 40 hours per week. Dykstra and Paul (2018) characterized fatigue, frustration, and cognitive workload as persistent concerns during cyber operations, and Nobles (2022) argued that stress, burnout, and security fatigue constitute human-factors risks that technology-centered security programs

routinely fail to address. Collectively, these findings justify examining burnout as an operational design problem at the organizational and individual levels rather than an isolated personal health concern.

This paper advances the position that organizations should integrate burnout and fatigue prevention into cybersecurity governance as a human-reliability, risk-management, and operational-resilience function. The position does not imply that every incident stems from exhaustion; cyber risk emerges from interactions among technology, processes, incentives, and human performance (Nobles, 2022; Tariq et al., 2025). The paper makes three contributions. First, it repositions burnout prevention as cybersecurity infrastructure, a human-reliability control that belongs in governance alongside technical safeguards (Maslach & Leiter, 2016). Second, it develops two analytic constructs: the cascading human-risk loop, which traces how diminished practitioner capacity transfers labor and risk to non-cybersecurity personnel, and the verification tax, naming the hidden cognitive and administrative work AI adoption can redistribute to security teams (Bainbridge, 1983; Tarafdar et al., 2019). Third, it translates these arguments into a five-layer prevention model with measurable indicators and an evaluation design that tests whether interventions reduce total demand rather than relocate it.

Position Statement and Theoretical Foundation

The central claim of this paper is direct: burnout and fatigue among cybersecurity professionals constitute an operational risk to both practitioners and the enterprise, and their prevention warrants the same governance discipline organizations apply to technical controls. Defending that claim requires precision about what burnout is and why the demand-resource structure of cybersecurity work renders both conditions predictable rather than exceptional.

Burnout and fatigue overlap physiologically, but they are not interchangeable. Fatigue may arise acutely or cumulatively from prolonged effort, insufficient sleep, circadian disruption, or interrupted recovery (Bufano et al., 2024; Sonnentag et al., 2017). Burnout, by contrast, reflects a chronic occupational pattern of exhaustion, cynicism or psychological distance from work, and reduced professional efficacy; it expresses the relationship between individuals and their work environments, not a failure of individual coping (Maslach & Leiter, 2016). The distinction matters because the corresponding interventions differ. A brief rest period may relieve acute fatigue but will not correct chronic overload or persistent role conflict. At the same time, a long-term culture initiative will not protect a responder who has worked through the night (Bakker & Demerouti, 2017). Effective programs must therefore prevent both immediate performance degradation and chronic deterioration.

The Job Demands-Resources model supplies the theoretical foundation. Job demands require sustained physical, cognitive, or emotional effort; job resources include autonomy, supportive leadership, role clarity, recognition, staffing, and recovery opportunities; and burnout risk increases when sustained demands exceed available resources (Bakker & Demerouti, 2007, 2017; Bakker et al., 2023). Cybersecurity exhibits demand characteristics that make the model particularly applicable, including continuous vigilance, threat uncertainty, limited tolerance for error, after-hours escalation, and continuous competency development (Dykstra & Paul, 2018; Nepal et al., 2024). In a systematic review of 64 studies, Bufano et al. (2024) found that irregular shifts, job stress, and extended hours each erode attention, working memory, executive functioning, and decision-making. Although those findings did not focus on cybersecurity, the analogy holds because cybersecurity work depends on the same cognitive functions.

Cybersecurity as a Complex Technological Work System

Cybersecurity professionals rarely evaluate risk within one integrated environment. They move among tools with different terminology, severity scales, data structures, and confidence indicators, and a single investigation may require correlating identity events, endpoint alerts, cloud logs, threat intelligence, and user statements (Dykstra & Paul, 2018; Tariq et al., 2025). Such

fragmentation creates hidden work: navigating systems, reconciling inconsistent data, and determining which source is authoritative escape staffing models yet consume attention and working memory, so organizations measuring only closed tickets underestimate the cognitive effort behind them (Bufano et al., 2024). Tool consolidation is therefore a workforce intervention as much as a technical one.

Security monitoring adds a signal-detection problem: analysts must sustain vigilance despite large volumes of benign, duplicate, or poorly contextualized alerts. Alert fatigue arises from repeated decisions under uncertainty, interrupted attention, and the possibility that a consequential event hides among routine notifications, encouraging rapid dismissal, over-escalation, or mechanical processing (Nepal et al., 2024; Tariq et al., 2025). Human-AI teaming may help when organizations assign routine alerts to automation and reserve collaboration for novel or ambiguous threats (Baruwal Chhetri et al., 2024), but automation does not eliminate the human burden when analysts must verify opaque classifications, manage exceptions, and correct model drift (Bainbridge, 1983; Parasuraman & Manzey, 2010). Automation thus reshapes vigilance work rather than removing it.

Unlike most demanding occupations, cybersecurity also involves an intelligent adversary that manufactures deception, ambiguity, and time pressure. Practitioners must consider that routine behavior may be malicious and malicious behavior benign, deciding whether to contain, observe, escalate, or communicate before complete information arrives (Dykstra & Paul, 2018; Nobles, 2022). Incidents likewise disregard schedules: a practitioner may work through the night and return to routine obligations without sufficient recovery, and a culture that equates constant availability with commitment normalizes fatigue as professionalism rather than preventable operational exposure (Bufano et al., 2024; Sonnentag et al., 2017). Recovery thus becomes a structural casualty rather than a personal choice.

Cybersecurity expertise also changes rapidly, and artificial intelligence adds technical, governance, and ethical knowledge demands. Learning functions as a job resource when organizations provide protected time and mentoring; it becomes a demand when leaders expect new competencies acquired off-hours while operational workload remains unchanged (Bakker et al., 2023; Hogemann et al., 2025). Role expansion adds ambiguity as security personnel advise on privacy, third-party risk, AI, data governance, and compliance without clear decision rights; when leadership treats cybersecurity as the destination for every unresolved technology risk, practitioners inherit accountability without corresponding authority or staffing (Bankins et al., 2024; Parent-Rochelleau & Parker, 2022). Expanding responsibility without matching resources is a textbook precondition for burnout.

Cascading Effects on Non-Cybersecurity Professionals

The consequences of cybersecurity burnout do not remain within the security function. Cybersecurity is an enterprise dependency, and disruptions can affect services and outcomes well beyond the compromised unit (Dameff et al., 2023; Neprash et al., 2026). This paper proposes a cascading human-risk loop in which excessive cyber demands and inadequate resources contribute to fatigue and burnout; diminished capacity weakens prioritization, coordination, communication, and recovery; operational delays or degraded decisions extend disruption; disruption transfers labor to business units, creating additional fatigue, workarounds, and exposure; and those conditions return more work to the cybersecurity team. The loop is offered not as an established causal chain but as a theoretical synthesis aligned with evidence on occupational cognition, security fatigue, incident disruption, and workload transfer (Bufano et al., 2024; Cram et al., 2021). Framed this way, the loop directs hypothesis development and program evaluation.

When a cyber event disrupts information systems, operational departments absorb much of the burden through manual procedures, reconstructed information, and additional hours. Dameff et al. (2023) examined 19,857 emergency department visits and found that a ransomware attack

was associated with disruptions at adjacent, otherwise unaffected emergency departments. Neprash et al. (2026) found that ransomware attacks reduced hospital volume by 17% to 24% during the initial week and were associated with a 34% to 38% increase in in-hospital mortality among patients already admitted at attack onset. These studies demonstrate cascading incident consequences without establishing burnout as their cause, a distinction responsible analysis must preserve. Security teams under sustained pressure may also respond by increasing restrictions, approval requirements, and policy complexity, transferring cognitive and administrative work to the general workforce (Cram et al., 2021; Malik et al., 2026). Security fatigue denotes the frustration, resignation, or disengagement employees experience when requirements feel excessive, repetitive, or obstructive (Stanton et al., 2016). Cram et al. (2021) linked security fatigue to both compliance and noncompliance behavior, and Malik et al. (2026) found among 298 employees that repeated demands deplete self-regulation capacity and foster emotional exhaustion and security-related cynicism. Exhausted teams may likewise turn adversarial toward users, encouraging shadow systems, delayed reporting, and informal workarounds (Nepal et al., 2024). Restrictive postures thus become a new source of enterprise risk.

Burnout-related turnover, moreover, removes more than headcount. It eliminates tacit knowledge of legacy systems, undocumented dependencies, and prior incidents that organizations cannot replace immediately through recruitment (Maslach & Leiter, 2016; Nepal et al., 2024). Remaining personnel inherit the departing workload, accelerating a cycle of overload, turnover, and declining organizational memory.

Artificial Intelligence and Machine Learning as a Burnout Paradox

Artificial intelligence and machine learning occupy two positions within the burnout problem. They can function as job resources that reduce repetitive work, yet also create new demands involving verification, monitoring, compliance, reskilling, and accountability. Organizational research increasingly describes this dual impact rather than uniformly positive or negative outcomes, and the central challenge is capturing AI-enabled efficiency without transferring hidden cognitive and governance costs to employees (Bankins et al., 2024; Chuang et al., 2025; Valtonen et al., 2025). Which position dominates depends on how organizations manage adoption. As a resource, AI-enabled systems can correlate events, classify alerts, summarize cases, and retrieve information, removing low-value repetition and preserving attention for ambiguous cases requiring context and judgment (Raisch & Krakowski, 2021). Frameworks distinguishing automation, augmentation, and collaboration match human involvement to task complexity (Baruwal Chhetri et al., 2024), and AI can strengthen learning across experience levels (Valtonen et al., 2025). Deployed deliberately, they function as genuine job resources.

AI does not simply eliminate tasks; it frequently redistributes them. Employees may spend less time producing initial analysis but more verifying outputs, reviewing exceptions, documenting provenance, monitoring model drift, and determining accountability (Bainbridge, 1983; Hogemann et al., 2025). This paper designates the added cognitive and administrative effort a verification tax, most consequential when systems generate fluent but incorrect outputs or when users cannot inspect the evidence behind a recommendation. Tarafdar et al. (2019) distinguished constructive and destructive technology-related stress, and in a three-wave study of 600 employees, Chuang et al. (2025) found AI efficacy and generative AI use positively associated with productivity while AI technostress increased exhaustion, intensified work-family conflict, and reduced job satisfaction. Qualitative evidence reinforces the pattern: professionals using generative AI reported stressors involving regulation, compliance, reliability, and shifting skill expectations alongside efficiency and learning benefits (Hogemann et al., 2025). The paradox is structural: tools that raise output can deepen exhaustion.

Cybersecurity teams also absorb AI-related work that implementation plans rarely count, including reviewing vendor architectures, controlling what data enters AI systems, investigating shadow AI use, and answering privacy and regulatory questions, while AI simultaneously

expands the threat environment through prompt injection, data leakage, model manipulation, and adversarial use of generative systems (Das et al., 2025). Leaders who budget for the tool but not its governance labor create a predictable workload gap.

When AI reduces task time, leadership may raise expected output rather than preserve recovery, converting efficiency into additional demand; organizations should therefore decide explicitly whether time savings will fund recovery, training, backlog reduction, or greater output (Chuang et al., 2025; Parent-Rochelleau & Parker, 2022). Overreliance on AI may also erode foundational skills. When AI consistently performs initial analysis, junior employees receive fewer opportunities to develop independent judgment, and experienced professionals lose autonomy when leaders treat AI output as authoritative (Bainbridge, 1983; Parasuraman & Manzey, 2010; Raisch & Krakowski, 2021). The appropriate objective is calibrated human-AI teaming that removes low-value demand while preserving learning, agency, and accountability.

The Case for an Integrated Burnout Prevention Program

Organizations commonly treat burnout as a human-resources concern addressed through counseling, wellness applications, or resilience education. Such resources can benefit individuals, but they do not correct a work system that continually generates excessive demands (Maslach & Leiter, 2016; Panagioti et al., 2017). A cybersecurity burnout prevention program should instead function as an integrated management system connecting cybersecurity governance, occupational health, workforce planning, human factors, AI governance, and enterprise risk management.

Evidence from other high-demand occupations supports this organizational orientation. Panagioti et al. (2017) analyzed 20 controlled physician interventions involving 1,550 participants and found small but significant reductions in burnout, with organization-directed interventions producing larger effects than physician-directed ones. Although this evidence does not prove that identical interventions transfer to cybersecurity, it supports prioritizing workload, scheduling, control, teamwork, and recovery over individual resilience. Table 1 presents a five-layer model that integrates these requirements.

Table 1. Integrated Cybersecurity Burnout Prevention Model

Layer	Primary objective	Illustrative controls	Core measures
Demand reduction	Remove avoidable workload and intensity.	Alert tuning; work-in-progress limits; meeting reduction; on-call guardrails.	Alert volume; false-positive burden; backlog age; after-hours work.
Resource amplification	Increase capacity, control, and support.	Adequate staffing; role clarity; cross-training; supportive supervision; recognition.	Job control; staffing ratio; turnover intention; training access.
Cognitive ergonomics	Align tools and workflows with human capabilities.	Integrated context; usable playbooks; interruption controls; AI provenance and uncertainty cues.	Tool switching; rework; handoff defects; verification time.
Recovery protection	Preserve physiological and psychological recovery.	Shift limits; protected breaks; compensatory rest; post-incident workload relief.	Shift length; sleep opportunity; recovery quality; fatigue.
Governance and resilience	Create accountability, privacy safeguards, and enterprise integration.	Executive sponsorship; privacy by design; nonpunitive measurement; AI workload review.	Program participation; privacy concerns; operational quality; spillover hours.

Demand reduction identifies work the organization should remove, automate, consolidate, or reassign, examining alert volume, false-positive rates, duplicate tools, meeting load, on-call frequency, and backlog growth; alert tuning should function as workforce protection as well as technical optimization (Tariq et al., 2025). On-call guardrails should limit consecutive high-intensity hours and prevent immediate return to full workload after overnight response (Nepal et al., 2024; Sonnentag et al., 2017). Demand reduction also requires prioritization authority, as leaders must determine what the organization will not do.

Resource amplification increases capacity to accomplish legitimate work safely by strengthening staffing, autonomy, supervision, role clarity, recognition, and training rather than assigning additional wellness activities (Bakker & Demerouti, 2017; Bakker et al., 2023). Managers should receive training in workload assessment, psychological safety, and fatigue recognition; cross-training requires protected time; and recognition should reward rigor, prevention, and mentoring rather than crisis performance alone.

Cognitive ergonomics aligns tools, information, and workflows with human capabilities through consistent severity information, preserved context across handoffs, reduced screen switching, and playbooks that structure work without forcing ill-fitting steps (Dykstra & Paul, 2018; Tariq et al., 2025). For AI-enabled tools, this layer requires confidence communication, source attribution, uncertainty indicators, and human override so output does not appear more authoritative than its reliability warrants (Parasuraman & Manzey, 2010). Design thus determines whether tools amplify or erode judgment.

Recovery protection designs rest in operations rather than leaving it to individual requests. Protected breaks, shift rotation, post-incident workload relief, after-hours communication limits, and compensatory rest should become operational standards (Bufano et al., 2024; Sonnentag et al., 2017). Counseling remains a valuable complement but cannot compensate for chronic understaffing (Maslach & Leiter, 2016). Recovery is thus an operational control, not an employee benefit.

Governance requires an accountable executive sponsor and a cross-functional body spanning cybersecurity, human resources, occupational health, privacy, and legal counsel, linking workforce indicators to operational risk without converting well-being data into surveillance; team-level trends take precedence, and personal health information remains confidential (Kellogg et al., 2020; Parent-Rochelleau & Parker, 2022). AI governance belongs here. Before deployment, organizations should assess accuracy, security, privacy, and total workload effects, including verification time, exception rates, and deskilling risk; leaders should not deem a system labor-saving until they measure the full work system after deployment (Bainbridge, 1983; Das et al., 2025). Governance supplies the accountability that sustains the other layers.

To illustrate, consider a hypothetical mid-sized healthcare security operations function of 14 personnel facing high alert volume, frequent on-call rotations, two senior resignations, and an AI triage tool that shifts analyst time toward verifying machine classifications. A phased, 12-month application proceeds layer by layer: alert tuning and work-in-progress limits reduce demand; backfills, supervisor training, and scheduled cross-training amplify resources; workflow consolidation and required confidence and provenance displays improve cognitive ergonomics; caps on consecutive response hours and compensatory rest protect recovery; and governance assigns an executive sponsor and requires a documented AI workload review before the tool is deemed labor-saving. Success would appear as declining alert, after-hours, and verification-time indicators alongside improving job control, reduced turnover intention, and stable operational quality. Failure would be equally visible: alert reductions coinciding with rising exception handling elsewhere indicate displaced demand, and productivity gains absorbed into higher quotas indicate intensification rather than relief (Chuang et al., 2025). The case thereby renders the cascading human-risk loop empirically testable.

Positive Aspects of an Integrated Program

The strongest argument for integration is that cybersecurity depends on reliable human judgment. Reducing excessive workload, improving recovery, clarifying roles, and redesigning alert systems protect the attention, working memory, and executive functioning on which security operations depend; a prevention program cannot eliminate error but can reduce the predictable conditions that make error, delay, or premature closure more likely (Bufano et al., 2024; Dykstra & Paul, 2018). Such a program also strengthens retention, preserves tacit knowledge, and limits the secondary workload vacancies impose, while gains in staffing, role clarity, and psychological safety improve incident coordination, which depends on teams sharing uncertainty and near misses without blame (Maslach & Leiter, 2016; Nepal et al., 2024). Human-capacity protection is therefore inseparable from operational reliability.

At the enterprise level, a resilient security workforce reduces the probability that avoidable delays transfer crisis work to other departments (Dameff et al., 2023; Neprash et al., 2026), while human-centered controls reduce friction, limit security fatigue, and preserve reporting willingness (Cram et al., 2021; Stanton et al., 2016). Embedding AI governance within burnout prevention discourages headcount-reduction logic and directs leaders to reinvest productivity gains in recovery, training, or staffing (Raisch & Krakowski, 2021). Prevention thus protects the enterprise, not only the security team.

Negative Aspects, Counterarguments, and Program Risks

Prevention is not costless. Demand reduction, staffing increases, protected recovery, and workflow redesign require investment, and implementation may temporarily reduce capacity (Panagioti et al., 2017). The objection is legitimate; leaders should weigh program costs against turnover, rework, prolonged incidents, and lost institutional knowledge rather than assume inaction is free. A poorly designed program can itself generate work through surveys, meetings, and reporting requirements, and organizations should avoid wellness-washing, where leaders publicize counseling resources while leaving workload, staffing, and management practices unchanged (Maslach & Leiter, 2016). Cost objections therefore argue for disciplined design, not inaction.

Surveillance presents a further risk. Organizations may attempt to infer burnout from email patterns, keyboard activity, biometric devices, or after-hours system use; such monitoring can damage trust, generate false inferences, and encourage employees to conceal fatigue; it requires informed consent, strict purpose limitation, independent oversight, and protection from disciplinary use (Kellogg et al., 2020; Parent-Rochelleau & Parker, 2022). Stigma compounds the problem, as employees may avoid reporting exhaustion when disclosure could affect promotion, assignments, or security clearances (Nepal et al., 2024). Programs must therefore separate health information from performance management.

Workload displacement and inequitable access to support can also undermine program integrity. Protected recovery time may burden colleagues when adequate staffing or coverage is unavailable, while reducing alert volume for one team may transfer exception handling, escalation, or monitoring responsibilities elsewhere (Bainbridge, 1983). These risks are especially consequential for night-shift, remote, contract, and on-call personnel, who often face the greatest fatigue exposure yet the least access to recovery resources and supervisory support (Dykstra & Paul, 2018; Nepal et al., 2024). Program evaluations should therefore examine intervention access, workload displacement, fatigue reduction, and operational outcomes across shifts, roles, work locations, and employment categories.

Measurement problems, moreover, can create false assurance. Burnout scores do not directly measure security performance, and faster response is not better when it reflects premature closure (Tariq et al., 2025). AI-supported prevention may also induce automation bias when employees overtrust recommendations, so organizations should preserve meaningful human review and periodically test whether personnel can perform critical

functions without AI support (Parasuraman & Manzey, 2010; Raisch & Krakowski, 2021). Measurement must therefore capture quality and resilience, not merely speed.

Implementation and Evaluation

Organizations should begin with a baseline assessment combining confidential workforce data, workflow observation, operational metrics, and interviews to identify structural demand-resource imbalances rather than classify individuals as security risks (Bakker et al., 2023; Panagioti et al., 2017). Leaders should establish executive ownership, map the work system, pilot interventions in a bounded function, and evaluate whether improvements reduce total demand without transferring work elsewhere.

Evaluation spans five domains. Human outcomes include validated measures of burnout, fatigue, recovery, job control, psychological safety, work-family conflict, and turnover intention (Maslach & Leiter, 2016; Nepal et al., 2024). Work-design indicators include shift length, after-hours activity, on-call frequency, alert burden, backlog age, and training time. Operational-quality indicators include handoff defects, rework, near misses, and documentation quality. Enterprise-cascade indicators include downtime, manual-workaround hours, overtime among non-cybersecurity teams, and post-incident recovery time (Dameff et al., 2023; Neprash et al., 2026). AI indicators include verification time, override and exception rates, model drift, unauthorized use, and ongoing governance labor (Baruwal Chhetri et al., 2024; Chuang et al., 2025; Das et al., 2025). Spanning five domains prevents improvement in one from masking deterioration in another.

Evaluation should use repeated measurements rather than a single post-program survey, comparing teams or periods while adjusting for incident volume and staffing changes, and supplementing quantitative findings with interviews because aggregate measures may conceal mechanisms (Panagioti et al., 2017). When AI or redesign saves time, leaders should again decide how that capacity will be reinvested (Raisch & Krakowski, 2021). Explicit reinvestment decisions keep efficiency gains from becoming intensification.

Limitations and Research Agenda

The evidence base does not yet justify strong causal claims about cybersecurity burnout and security outcomes. Nepal et al. (2024) studied 35 incident responders, and those findings should not be generalized across roles or organizations. Cybersecurity-specific controlled prevention trials remain scarce, and much of the intervention evidence derives from healthcare occupations sharing high workload but differing in adversarial conditions and technology (Panagioti et al., 2017). The AI evidence is also developing rapidly; studies examine different technologies, occupations, and contexts, and findings about traditional AI may not transfer to generative AI, so conclusions in both directions should remain conditional (Chuang et al., 2025; Das et al., 2025; Hogemann et al., 2025). These constraints define the evidentiary gaps future work must close.

Future research should conduct longitudinal and quasi-experimental studies connecting job demands, fatigue, burnout, cognition, and operational quality, examining whether prevention changes alert disposition, handoff quality, incident duration, and retention (Bufano et al., 2024; Nepal et al., 2024). Studies should investigate cascading effects on non-cybersecurity personnel and whether redesigned controls reduce security fatigue (Cram et al., 2021; Malik et al., 2026). AI research should measure total work-system effects, including verification burden, skill development, and work transferred among teams, clarifying when AI functions as a resource and when it becomes a demand (Bainbridge, 1983; Raisch & Krakowski, 2021). Such research would supply the causal evidence this paper anticipates.

Conclusions

Cybersecurity burnout prevention should not remain confined to employee wellness programming or individual responsibility. Cybersecurity professionals work in complex, high-consequence environments characterized by sustained vigilance, fragmented technologies, adversarial uncertainty, interrupted recovery, and rapid competency change, conditions that create a credible human-reliability concern even though direct causal evidence remains incomplete (Dykstra & Paul, 2018; Tariq et al., 2025). The organizational consequences extend beyond security teams, and artificial intelligence offers no simple remedy; the relevant question is whether AI reduces total human demand while preserving judgment, learning, accountability, and recovery (Chuang et al., 2025; Raisch & Krakowski, 2021). An effective program reduces unnecessary demand, strengthens resources, improves cognitive ergonomics, protects recovery, and establishes ethical governance while guarding against surveillance, stigma, and wellness-washing (Bakker et al., 2023; Kellogg et al., 2020). Burnout prevention becomes cybersecurity infrastructure only when it changes the conditions under which cybersecurity work occurs. Organizations that treat practitioner capacity as a governed asset, measured, protected, and deliberately renewed, will sustain reliable cybersecurity operations through the AI era more effectively than those that treat exhaustion as the price of vigilance.

References

- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775-779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- Bakker, A. B., & Demerouti, E. (2007). The job demands-resources model: State of the art. *Journal of Managerial Psychology*, 22(3), 309-328. <https://doi.org/10.1108/02683940710733115>
- Bakker, A. B., & Demerouti, E. (2017). Job demands-resources theory: Taking stock and looking forward. *Journal of Occupational Health Psychology*, 22(3), 273-285. <https://doi.org/10.1037/ocp0000056>
- Bakker, A. B., Demerouti, E., & Sanz-Vergel, A. I. (2023). Job demands-resources theory: Ten years later. *Annual Review of Organizational Psychology and Organizational Behavior*, 10, 25-53. <https://doi.org/10.1146/annurev-orgpsych-120920-053933>
- Bankins, S., Ocampo, A. C., Marrone, M., Restubog, S. L. D., & Woo, S. E. (2024). A multilevel review of artificial intelligence in organizations: Implications for organizational behavior research and practice. *Journal of Organizational Behavior*, 45(2), 159-182. <https://doi.org/10.1002/job.2735>
- Baruwal Chhetri, M., Tariq, S., Singh, R., Jalalvand, F., Paris, C., & Nepal, S. (2024). Towards human-AI teaming to mitigate alert fatigue in security operations centres. *ACM Transactions on Internet Technology*, 24(3), Article 12, 1-22. <https://doi.org/10.1145/3670009>
- Bufano, P., Di Tecco, C., Fattori, A., Barnini, T., Comotti, A., Ciocan, C., Ferrari, L., Mastorci, F., Laurino, M., & Bonzini, M. (2024). The effects of work on cognitive functions: A systematic review. *Frontiers in Psychology*, 15, Article 1351625. <https://doi.org/10.3389/fpsyg.2024.1351625>
- Chuang, Y. T., Chiang, H. L., & Lin, A. P. (2025). Insights from the Job Demands-Resources model: AI's dual impact on employees' work and life well-being. *International Journal of Information Management*, 83, Article 102887. <https://doi.org/10.1016/j.ijinfomgt.2025.102887>
- Cram, W. A., Proudfoot, J. G., & D'Arcy, J. (2021). When enough is enough: Investigating the antecedents and consequences of information security fatigue. *Information Systems Journal*, 31(4), 521-549. <https://doi.org/10.1111/isj.12319>
- Dameff, C., Tully, J., Chan, T. C., Castillo, E. M., Savage, S., Maysent, P., Hemmen, T. M., Clay, B. J., & Longhurst, C. A. (2023). Ransomware attack associated with disruptions at adjacent emergency departments in the United States. *JAMA Network Open*, 6(5), e2312270. <https://doi.org/10.1001/jamanetworkopen.2023.12270>
- Das, B. C., Amini, M. H., & Wu, Y. (2025). Security and privacy challenges of large language models: A survey. *ACM Computing Surveys*, 57(6), Article 152, 1-39. <https://doi.org/10.1145/3712001>
- Dykstra, J., & Paul, C. L. (2018). Cyber operations stress survey (COSS): Studying fatigue, frustration, and cognitive workload in cybersecurity operations. In *11th USENIX Workshop on Cyber Security Experimentation and Test (CSET 18)*. USENIX Association. <https://www.usenix.org/conference/cset18/presentation/dykstra>
- Hogemann, M., Hein, L., Britsche, J. O., & Thomas, O. (2025). Technostress and generative AI in the workplace: A qualitative analysis of young professionals. *Frontiers in Artificial Intelligence*, 8, Article 1728881. <https://doi.org/10.3389/frai.2025.1728881>
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366-410. <https://doi.org/10.5465/annals.2018.0174>

- Malik, A., Goel, S., & Sinha, S. (2026). Security fatigue: Manifestation of emotional exhaustion and cynicism by depletion of self-regulation capacity. *European Journal of Information Systems*. Advance online publication. <https://doi.org/10.1080/0960085X.2026.2621809>
- Maslach, C., & Leiter, M. P. (2016). Understanding the burnout experience: Recent research and its implications for psychiatry. *World Psychiatry*, 15(2), 103-111. <https://doi.org/10.1002/wps.20311>
- Nepal, S., Hernandez, J., Lewis, R., Chaudhry, A., Houck, B., Knudsen, E., Rojas, R., Tankus, B., Prafullchandra, H., & Czerwinski, M. (2024). Burnout in cybersecurity incident responders: Exploring the factors that light the fire. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1), Article 27, 1-35. <https://doi.org/10.1145/3637304>
- Neprash, H., McGlave, C., & Nikpay, S. (2026). Hacked to pieces? The effects of ransomware attacks on hospitals and patients. *American Economic Journal: Economic Policy*, 18(1), 256-281. <https://doi.org/10.1257/pol.20240594>
- Nobles, C. (2022). Stress, burnout, and security fatigue in cybersecurity: A human factors problem. *HOLISTICA - Journal of Business and Public Administration*, 13(1), 49-72. <https://doi.org/10.2478/hjbpa-2022-0003>
- Panagioti, M., Panagopoulou, E., Bower, P., Lewith, G., Kontopantelis, E., Chew-Graham, C., Dawson, S., van Marwijk, H., Geraghty, K., & Esmail, A. (2017). Controlled interventions to reduce burnout in physicians: A systematic review and meta-analysis. *JAMA Internal Medicine*, 177(2), 195-205. <https://doi.org/10.1001/jamainternmed.2016.7674>
- Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3), 381-410. <https://doi.org/10.1177/0018720810376055>
- Parent-Rochelleau, X., & Parker, S. K. (2022). Algorithms as work designers: How algorithmic management influences the design of jobs. *Human Resource Management Review*, 32(3), Article 100838. <https://doi.org/10.1016/j.hrmr.2021.100838>
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46(1), 192-210. <https://doi.org/10.5465/amr.2018.0072>
- Sonnentag, S., Venz, L., & Casper, A. (2017). Advances in recovery research: What have we learned? What should be done next? *Journal of Occupational Health Psychology*, 22(3), 365-380. <https://doi.org/10.1037/ocp0000079>
- Stanton, B., Theofanos, M. F., Prettyman, S. S., & Furman, S. (2016). Security fatigue. *IT Professional*, 18(5), 26-32. <https://doi.org/10.1109/MITP.2016.84>
- Tarafdar, M., Cooper, C. L., & Stich, J. F. (2019). The technostress trifecta: Techno eustress, techno distress and design: Theoretical directions and an agenda for research. *Information Systems Journal*, 29(1), 6-42. <https://doi.org/10.1111/isj.12169>
- Tariq, S., Baruwat Chhetri, M., Nepal, S., & Paris, C. (2025). Alert fatigue in security operations centres: Research challenges and opportunities. *ACM Computing Surveys*, 57(9), Article 224, 1-38. <https://doi.org/10.1145/3723158>
- Valtonen, A., Saunila, M., Ukko, J., Treves, L., & Ritala, P. (2025). AI and employee wellbeing in the workplace: An empirical study. *Journal of Business Research*, 199, Article 115584. <https://doi.org/10.1016/j.jbusres.2025.115584>